An Advanced Fast Steering Mirror for Optical Communication

by

Daniel Joseph Kluk

B.S., Mechanical Engineering
Northwestern University, 2002

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

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Author . . . .
Department of Mechanical Engineering
May 24, 2007

Certified by . . . .
David L. Trumper
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by . . . .
Lallit Anand
Chairman, Departmental Committee on Graduate Students

This work was sponsored by the Unted States Air Force under Air Force Contract FA8721-05-C-002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.

An Advanced Fast Steering Mirror for Optical Communication

by

Daniel Joseph Kluk

Submitted to the Department of Mechanical Engineering
on May 24, 2007, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

I describe in this thesis the design, fabrication, assembly, and testing of an Advanced Fast Steering Mirror (AFSM) for precision optical platforms. The AFSM consists of a mirror driven in two rotational axes by normal force electromagnetic actuators, and controlled via position feedback loops. The dynamic performance is sufficient to provide high bandwidth (approximately 5 kHz) disturbance rejection of base motion, and as such the device is particularly suited to beam stabilization tasks in laser communication, lidar, and similar optical applications. In fact, work on the Mars Laser Communication Demonstration project at MIT Lincoln Laboratory provided the original impetus for developing the subject technology.

My work on this project is divided into five distinct phases: Electromagnetic and mechanical design of the mirror itself; fabrication and assembly of the mechanical hardware; initial testing and dynamic model generation; design and fabrication of an electronic analog controller; and final closed loop performance demonstrations. I performed the first two phases on the MIT campus, and the final three phases at MIT Lincoln Laboratory. Each project phase is described in detail herein.

Ultimately, I demonstrate performance from the hardware and control electronics exceeding the original design goal of 5 kHz. As this original prototype is merely a testbed, I also describe possible evolutions of the design to optimize form factor, performance, and flightworthiness.

I wish to thank the Advanced Concepts Committee at MIT Lincoln Laboratory for providing the sponsorship that made this project a success.

Thesis Supervisor: David L. Trumper
Title: Professor of Mechanical Engineering

Acknowledgments

So many people have contributed to this small accomplishment that I will have great difficulty naming them all in this limited space. But rest assured, this work is not mine alone–not by a long shot. I owe many a debt of gratitude to all those mentioned here, who each in their own way have given me the knowledge, strength, discipline, love, and support I needed to see this task through to the very end.

First, I would like to thank my advisor, Professor David Trumper. Professor Trumper has been a source of vast knowledge and support from the very beginning of this project. Besides providing advice and insight into the design of the AFSM, he constantly motivated me to step outside of my comfort zone by challenging my ideas and expanding the scope of the project to fit my potential–despite the fact that I often didn't know what the limit of my potential was! I owe many thanks to Professor Trumper for his teaching, wisdom, practicality, and encouragement.

In addition to Professor Trumper, five other individuals provided direct assistance, in various roles, on this project. Mike Boulet was very helpful in a wide variety of areas during the testing of the AFSM–especially, but not limited to, design and layout of the wire-wrap two-channel controller board and coding of several helpful Matlab scripts. Mike's impeccable attention to detail and wizardry with Matlab were just what I needed during this time–especially given my relative weakness in both these areas.

Jamie Burnside was my primary contact at Lincoln Labs, and was instrumental in facilitating the initial proposal activity to secure ACC funding. He also set up design reviews, and provided much of the background information needed to set the AFSM performance goals. He also guided me through several key administrative tasks at Lincoln, such as purchasing.

Third is David Cuff, my labmate[1] during the design phase of the AFSM. David was instrumental as a sounding board for ideas, and we struggled through the finer points of magnetic analysis and control systems together. David also provided key insights into the design of the elastomer bearings. He has become a good friend.

The fourth key contributor is a man I haven't even met in person: Fred Sommerhalter. Fred has worked with Professor Trumper for many years, and is an expert craftsman who performs coil winding and other magnetic assembly work as a hobby. Fred performed the critical work of winding the AFSM coils and assembling the actuator halves with the level of precision demanded by the design. Fred's artisanship is self-evident in the hardware he builds, and the work he produced for the AFSM was no exception.

Finally, I owe thanks to Gerry Wentworth, the machinist in charge of the LMP machine shop. Gerry helped me in every aspect of the AFSM component fabrication, from the simplest turned parts to the most complex CNC-machined housings and base pieces. Besides being an excellent teacher of the machinist's trade, Gerry is also very easy-going and friendly, and I enjoyed chatting and sharing stories with him.

Of course, I must also thank many others, whose contributions to this project

---

[1] In a few jaded moments, I often thought of him as my cellmate!

were less direct, but yet were equally essential to its success. First, I must thank my friends and colleagues at MIT, who in one way or another provided insight, tutoring, advice, camaraderie, and commiseration. Many thanks to the members of the PMC Lab: Emre Armagan, Joe Cattell, Larry Hawe, Aaron Gawlik, Dean Ljubicic, Ian MacKenzie, Kevin Miu, and our wonderful administrative assistant, Laura Zaganjori. I would also like to thank the officers and members of the Snowriders ski club, and my teammates on my various intramural sports teams–especially the GAME and Neanderthals hockey teams for allowing me to subject them to my pitiful goals-against average in net.

Second, thanks goes to several employees at Lincoln Labs for their assistance during my time there: Ed Corbett, Annmarie Gorton, Karen Sacco in the electronic stockroom, and Jeanne Clarke in Group 76 electronics. Special thanks go to Todd Mower for his help on the elastomer bearing stiffness investigation, and to Al Pillsbury for his explanations of the Lincoln HBSM hardware, and input into the AFSM mechanical design.

Next I'd like to acknowledge my friends, both in Boston and around the country, who in their own unique styles kept me laughing, loving, relaxed, and (relatively) stress-free while I ran the mental gauntlet that is MIT: Andrea Brand, Dariusz Golda, Cathy and Daniel Hsu, Jeremy Indick, Dave Ison, Paul King, Chad Laurent, Charlie McGinnis, Alex Mettler, RaNaye Neudauer, Kevin Paul, Kurt Schoonover, Matt St. Martin, Dave and Samantha Stagney, Melissa Stenach, Tami Urman, Bill Vogt, Ruthie Wertheimer, Ryan Wiechens, Justin Verdirame, and many others. I treasure your friendship and hope to spend more time visiting you now that this task is complete!

Finally, and most importantly, I wish to thank my family for all their love and support, not only at MIT, but in all other times of my life: My father Robert, my mother Carolyn, brothers Brian and Rob, and sisters Erin, Leanne, and Andrea. Rob, Erin, and Andrea passed away in a tragic auto accident in 2002. Since that time, my pursuit of success has been motivated by a desire to honor their memory, and capture a few of the opportunities that they were not able to realize during their short lives. To them I dedicate this work.

To Andrea, Erin, and Rob

I will see you again someday...

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis I document the design and development of a novel mechanism for beam pointing and stabilization in optical systems: The Advanced Fast Steering Mirror (AFSM). My work on the device is part of a technology pathfinder for advanced optical communications systems research, led primarily by NASA and the United States Air Force. This work was sponsored by the Advanced Concepts Committee of MIT Lincoln Laboratory.

1.1 Thesis Organization and Summary

I open in Chapter 1 with a summary of my work and results, as well as an overview of the optical communication technology that provides the impetus for my research. I also review prior work in fast steering mirror technology, and how shortcomings in existing hardware have motivated the pursuit of the higher performance AFSM design. I conclude the chapter by stating the performance goals for the design, and provide an overview of the hardware configuration and operational principles. A photograph of the AFSM hardware is shown in Figure 1-1.

Chapter 2 covers the fundamental mechanical issues in the AFSM design. I first discuss the implications of the mechanical design on the control system design, and illustrate how fundamental choices made at this stage become critical for maximizing performance of the controlled system. After establishing mechanical dynamic goals,

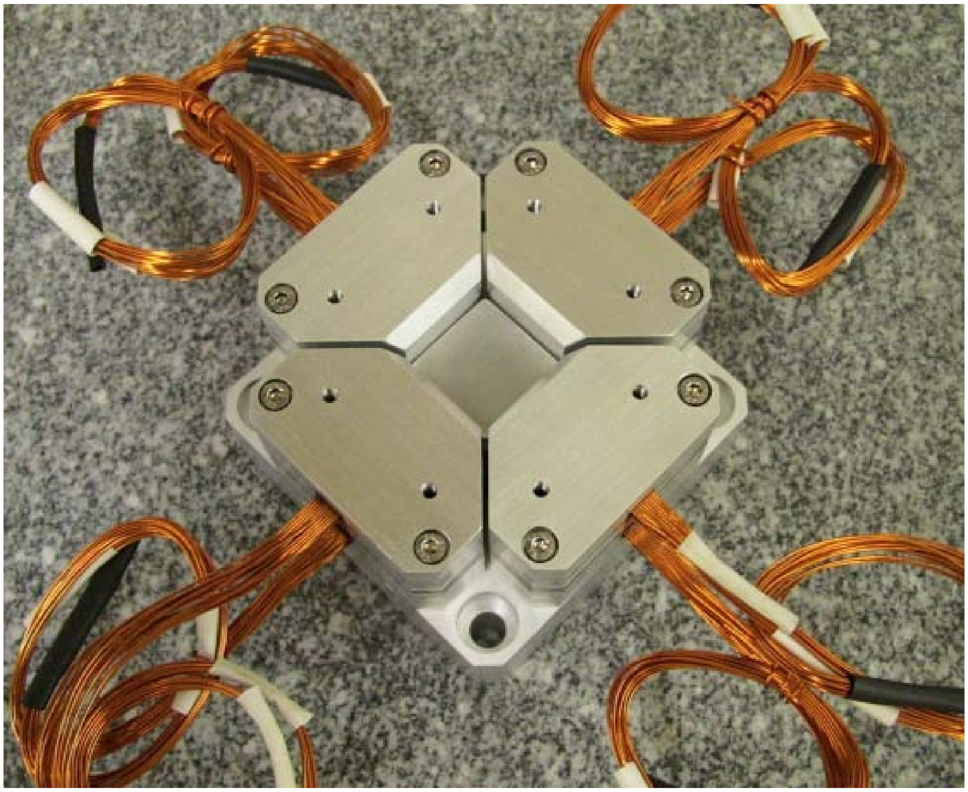

Figure 1-1: Finished AFSM base assembly, prior to electrical integration.

I discuss the tradeoffs required to meet them. In particular, I consider the moving mirror mass versus stiffness properties, the selection of inertia and suspension stiffness values to properly prescribe the mirror's rigid body resonant mode, and the selection of symmetric mass properties to mechanically decouple the two actuated degrees of freedom. Closely tied to these tradeoffs is the important problem of proper kinematic constraint of the dynamic elements. I use a combination of metallic and elastomeric flexure elements for this purpose. I also provide modal analysis results of the AFSM dynamic components using the finite element method, and discuss the acceptability of these results in terms of the closed-loop controller design. The main components in the mechanical design are illustrated in Figure 1-2.

I cover in Chapter 3 the design and modeling of the flux-steering electromagnetic actuators which are fundamental to attaining the specified performance improvements. The actuator is illustrated in Figure 1-3. I derive results using the reluctance model (magnetic circuit) approximation method, and make an argument for this

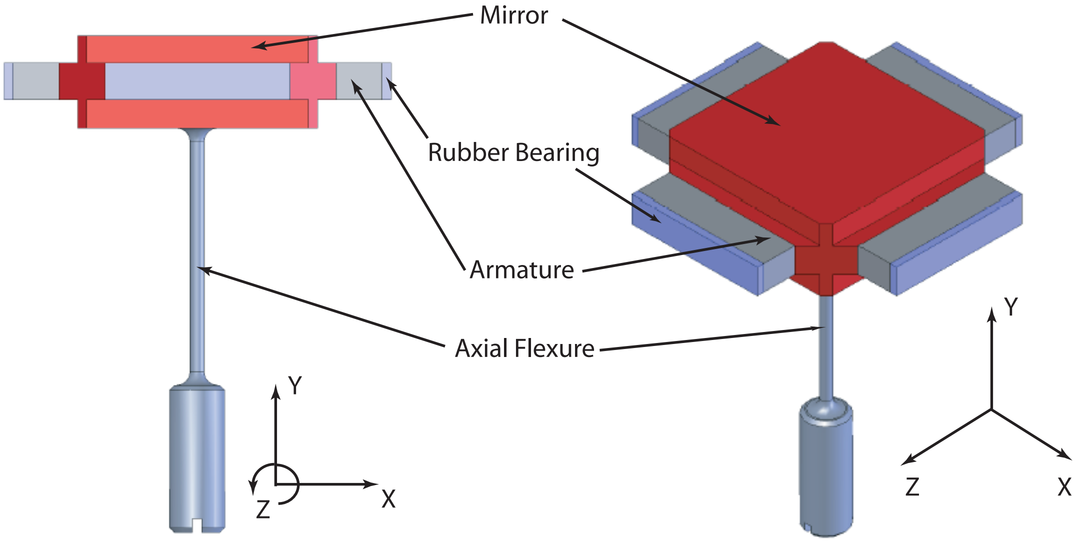

Figure 1-2: CAD model view of the AFSM moving mechanical components.

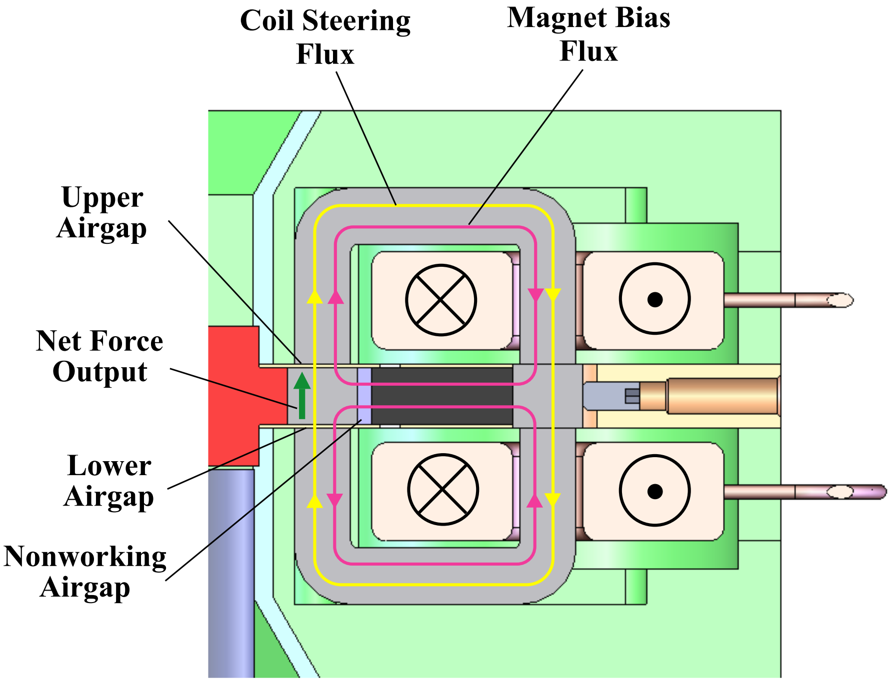

Figure 1-3: The AFSM flux-steering actuator.

method as the preferred choice in many practical electromagnetic machine design applications. To illustrate the validity of the technique, I also derive the same analytical results using Maxwell's equations, and present the analysis in Appendix A. I close the discussion in Chapter 3 by describing the physical design of the actuator component parts.

The AFSM component fabrication and assembly is described in Chapter 4. Here I detail the essential features in the physical components that ensure precision in the subsequent hardware performance. I also discuss in this chapter the power electronics and sensors (capacitance probes and an optical quad cell) that serve as the main electronic interface to the system controller, as well as the supporting electronics I designed to manipulate and scale the signals into a form suitable for use in a feedback control system. Together, the AFSM hardware, power electronics, and sensors comprise the dynamic plant around which the closed-loop compensator is designed.

In Chapter 5 I describe the system identification and controller design for both types of feedback sensors used. An example of the plant open-loop dynamics is given in Figure 1-4. Here I consider the tradeoffs between analog and digital control architectures, and ultimately make a case for the analog design that I actually implemented and tested. Also discussed are the final experimental results of the fully integrated system in both single- and dual-axis operational modes. I close the chapter by comparing the test results to the original AFSM design goals. These results are given here in Table 1.1. A Bode plot of the system closed-loop performance at 10 kHz bandwidth is given in Figure 1-5.

Table 1.1: AFSM Measured Performance Characteristics

Performance Parameter	**Design Goal**	**Measured Performance**
Bandwidth (Single Axis)	5 kHz	10 kHz
Bandwidth (Dual Axis)	N/A	2 kHz
Angular Range	± 10 mrad	± 3.5 mrad
Angular Acceleration	1×10^5 rad/sec^2	$1x10^5$ rad/sec^2
Angular Resolution	8.7×10^{-7} rad	Variable
Mirror Aperture	20 mm	20 mm

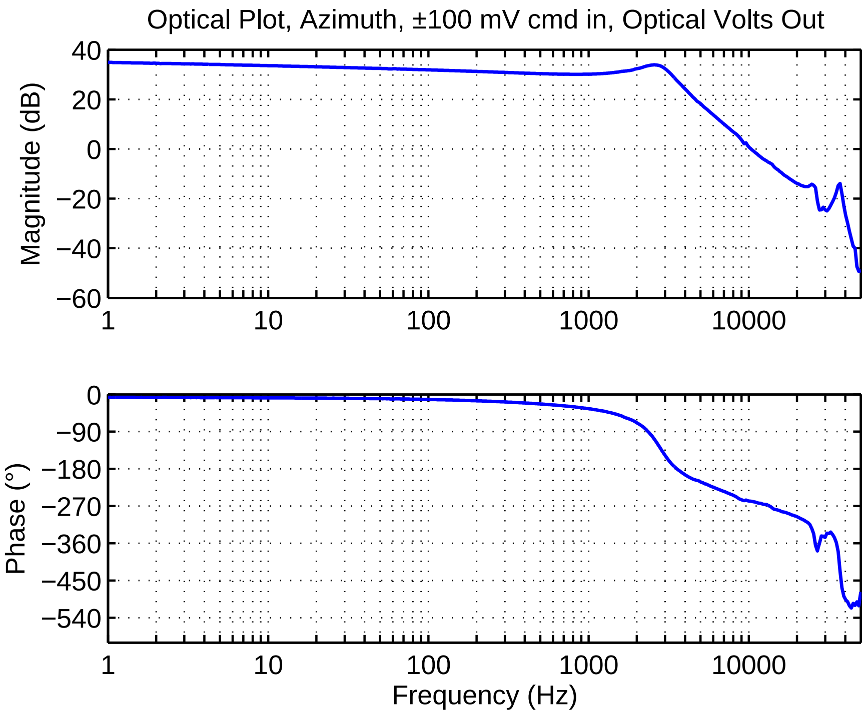

Figure 1-4: Azimuth closed-loop frequency response using the optical quad cell.

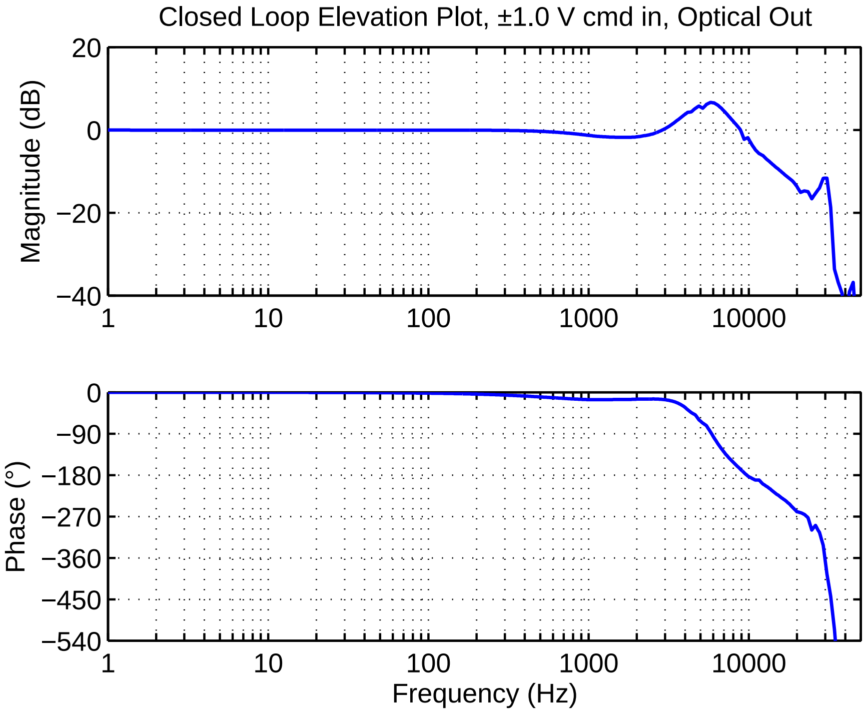

Figure 1-5: Elevation closed-loop frequency response under optical feedback.

Chapter 6 concludes the thesis by offering suggestions for future work in the form of optimized designs for form factor, small signal bandwidth, etc. These designs seek to evolve the AFSM design, which was conceived purely as a technology demonstrator, into mature concepts that provide solutions for specific applications.

The final portion of this thesis consists of three appendices containing supporting information. As previously mentioned, Appendix A provides a derivation of the AFSM actuator electromagnetic equations using Maxwell's laws. Appendix B documents the follow-up work I performed regarding the unexpectedly high stiffness characteristics discovered in the AFSM hardware during testing. Finally, Appendix C provides transfer function derivations for the analog controllers designed in Chapter 5.

1.2 Optical Communication Technology

In recent years, long-distance communication and data transfer technology through optical means has emerged as a potential replacement for traditional radio-based systems [15]. In particular, such optical communication systems show promise as a secure, power-efficient means of data transmission between aircraft, spacecraft, and/or ground stations during military and scientific-exploration missions. In the last few years, NASA and the U.S. Air Force have both initiated programs to develop this technology. The most notable effort to date has been the Mars Lasercomm demonstration program of 2005, in which MIT Lincoln Laboratory played a major development role.

Despite the system-level benefits of enhanced power efficiency and security, long-distance communication through optical means introduces significant new technical challenges that must be overcome prior to the implementation of any practical system. One of these is the problem of beam pointing, acquisition, and stabilization. Unlike radio-based systems, which even for directional signals are characterized by beam angular dispersions of several degrees or more, the lasers proposed for use in optical systems have beam dispersions only a small fraction of this size. Despite being very power-efficient, the small dispersion greatly compounds the pointing problem. For example, as determined during the Mars Lasercomm technology demonstration, in order to hit Earth from a Mars a laser beam originating from an orbiting satellite must be pointed within 400 nanoradians RMS error [11]. When contrasted with pointing error requirements on the order of tens of microradians for a radio system, it is apparent that the pointing precision for the optical beam is about ten to twenty times tighter than that required for a radio system.

Compounding the comparatively simple problem of static pointing, the main implication of reducing the pointing error budget is that previously inconsequential (i.e. "in-the-noise") disturbance sources become significant, and must be addressed in the system design. For instance, in a satellite application, minute sources of mechanical vibration transmitted through structural components to the optical system may introduce pointing errors of unacceptable magnitude. As described in Loney [15], vibration sources such as those from solar array drives and momentum wheel bearing noise vary widely in frequency and contain enough energy to introduce meaningful beam pointing errors. Hence, the optical pointing system must be robust enough to reject these disturbances over a wide frequency spectrum. As another example, a similar system installed on an aircraft experiences disturbances in the optical path arising from turbulent airflows over the aircraft exterior. Here, the disturbances are inherently random in both amplitude and frequency, and thus a robust, broad-spectrum active rejection system is required to eliminate the resulting errors.

1.3 Motivations for AFSM Development

To address the issue of broad-spectrum disturbance rejection, the system developed during the Mars Lasercomm technology demonstration at MIT Lincoln Laboratory in 2005 employed many separate components in a staged configuration. These are described in Hawe [11] and are shown in the simplified schematic given in Figure 1-6.

As designed, the system was capable of rejecting spacecraft-generated disturbances on the order of microradian amplitude from DC to approximately 1 kHz, with an overall closed-loop system bandwidth of about 2 kHz. Without delving into great detail,

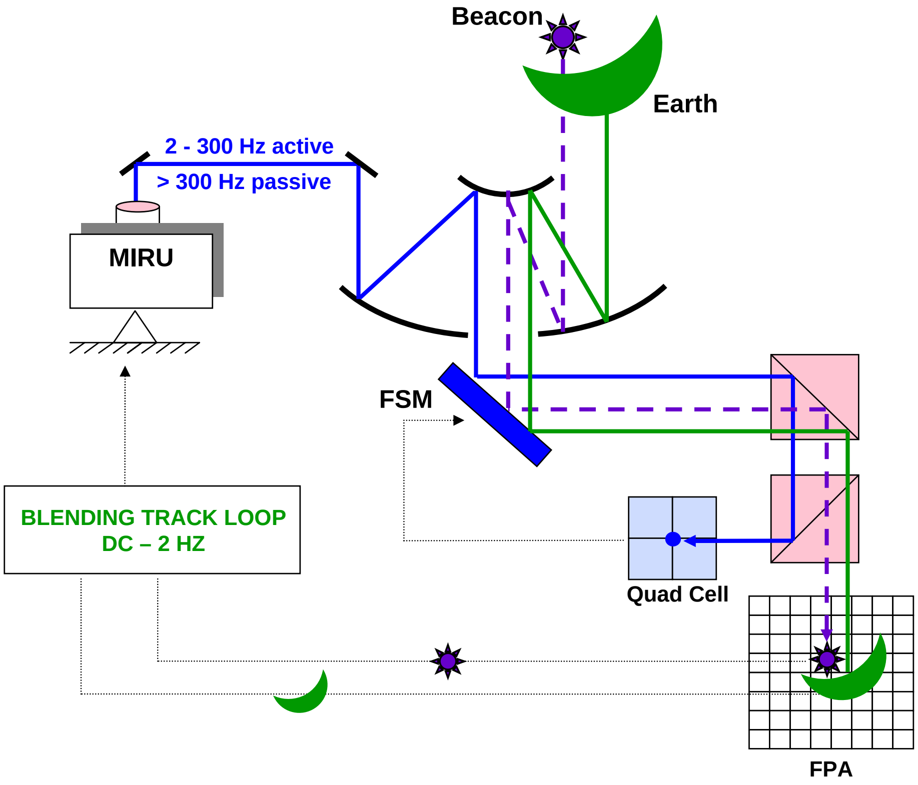

Figure 1-6: Simplified Mars Lasercomm spacecraft optical component schematic (from Hawe [11]).

the system operation is as follows: After coarse spacecraft pointing via conventional methods (i.e., on-board gyroscopes, star trackers, and attitude control), the earth-based beacon laser signal is acquired by the telescope and projected onto the focal plane array (FPA), which provides angular position information from DC to about 2 Hz. In addition, angular rate sensors on board the magnetohydrodynamic inertial reference unit (MIRU) generate feedback information from 2-300 Hz. The MIRU also generates an inertially stable reference laser beam that is injected into the same telescope optics used to detect light signals from earth. The signals originating from both earth and the MIRU are reflected off the fast steering mirror (FSM), and the reflected MIRU signal is sensed by the quad cell detector. The quad cell is the feedback sensor in a zero-reference control loop closed around the FSM, and with a sensor bandwidth of approximately 100 kHz, it is capable of measuring the remainder of the disturbance spectrum. Because the reference beam from the MIRU is inertially stable, any errors detected at the quad cell must be due to mechanical base disturbances in the spacecraft transmitted through one or more of the system's optical components (note that no atmospheric disturbances exist due to the vacuum environment). The FSM control loop actively corrects for these errors by actuating the mirror in both azimuth and elevation. Hence, even though the main optical components in the system (for example, the telescope mirrors) may experience mechanical disturbances, the incoming and outgoing optical signals remain inertially stable due to the corrections provided by the FSM. Hence, the optical system is pointed and stabilized to the required accuracy, and is ready to transmit and receive data.

Although successful as a demonstration, the Mars Lasercomm technology is still subject to certain limitations. First, due to hardware limitations the active control system does not have sufficient loop gain to adequately reject all angular errors within the required frequency band. Therefore it is necessary to provide additional attenuation by mounting the optical components on a passive isolation platform, which increases the system's cost, mass, and complexity. A second, related problem is that the system bandwidth is not sufficient to reject the entire disturbance spectrum necessary to meet angular error requirements, which again requires the addition of a passive vibration isolation system to augment the active component.

Specifically, the Mars Lasercomm investigators found that the existing FSM technology used in the demonstration was a key contributing factor to both of these shortcomings, and concluded that the availability of a higher bandwidth FSM could yield a cheaper system with less complexity, yet have the same or better performance as existing designs. The results of one study into this issue are shown in Figure 1-7. Here, the power spectral density (PSD) in angular error of a conventional isolation system design consisting of a 1 kHz bandwidth FSM and passive platform isolators with 20 Hz resonant frequency is plotted (dark blue line) against several other cases.

In the first case (green line), the existing FSM is maintained, but the resonant frequency of the passive isolators is changed to 200 Hz. Such isolators are cheaper, simpler, and more compact than a 20 Hz isolator, but provide far less attenuation of

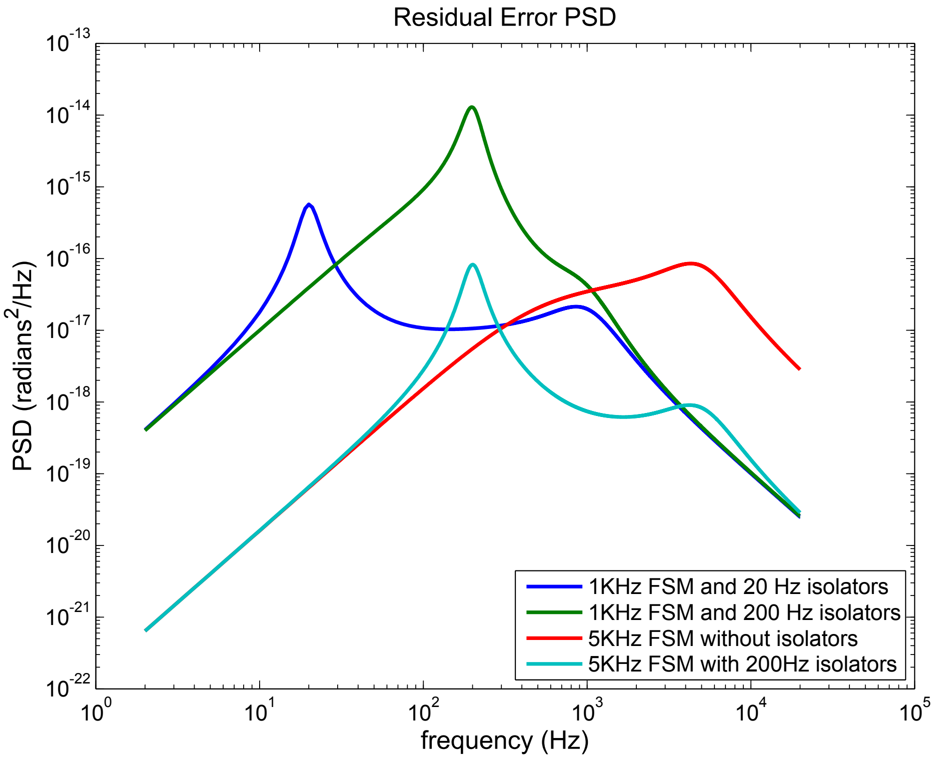

Figure 1-7: Spacecraft pointing error PSD plots for various active and passive isolation cases. Plot courtesy of Jamie Burnside.

base vibrations. The result is that the FSM is subject to larger disturbance amplitudes, which may not be sufficiently rejected. Note that relative to the original case, the mid-band error rejection is very poor–almost four orders of magnitude worse in the 100-300 Hz range.

In the second case (red line), the 1 kHz FSM is replaced with a 5 kHz FSM and the passive isolators are removed completely. This represents the simplest and cheapest case. Here, it is evident that the error performance at low frequencies is superior by about three orders of magnitude to either of the prior configurations; however, it is worse by about two orders of magnitude at high frequencies due to the absence of the passive isolators.

The final case (light blue line) takes a middle ground, using the 5 kHz FSM and the cheaper 200 Hz isolators. This system provides the same superior rejection at low frequencies due to the high gain FSM, and aside from a small band around 200

Hz, it provides equal or better high-frequency performance relative to the original system. Thus, the system displays a net performance increase, despite being simpler and cheaper.

Studies such as this one prompted the search for new, alternative technologies for use in a flight system. In addition, the investigators desired a new FSM technology with performance sufficient to be useful for many years of anticipated future missions. Along with the space and airborne laser communication technologies mentioned previously, use of the new FSM technology is foreseen in lidar systems and ground-based telescope adaptive optics, and also has application to industrial applications such as maskless photolithography for semiconductor manufacturing.

1.4 Discussion of Prior Art

The development of a proof-of-concept advanced fast steering mirror (AFSM) is the subject of this thesis. I describe in this and the following chapters the underlying technology, as well as my work designing, fabricating, and testing the prototype AFSM. Before beginning this discussion, however, it is prudent to review prior efforts in FSM design, as well as the enabling research that led to the design that I eventually built and tested.

1.4.1 Early Fast Steering Mirrors

Fast steering mirror technology has existed in various stages of sophistication for many years. The most basic technology; that is, a single axis mirror driven by a galvanometer[1], has existed for decades, and is still widely used in many scientific and industrial machines, such as bar code scanners, high speed product printing, medical imaging devices, and laser light show entertainment displays [5]. Two-axis operation

---

[1] A galvanometer is an electromagnetic actuator in which a current-carrying coil is placed in a static magnetic field. It is similar to a voice coil, but is designed to provide a torque (via alignment of magnetic dipoles) rather than a linear force. Unlike typical rotating electric motors, galvanometers are designed for only a limited range of angular motion. Typically, the electromagnetic torque is counteracted by a restoring torsion spring.

may be achieved through the use of two such mirrors in a staged configuration. A photograph of a commercially available galvanometer from Cambridge Technology, Inc. is shown in Figure 1-8.

Figure 1-8: Cambridge Technology model 6400, a single-axis galvanometer.

Galvanometer-based mirrors typically feature very high angular travel (on the order of degrees to tens of degrees), but limited bandwidth. This tradeoff may be reversed by employing piezoelectric actuators to move the mirror, rather than the Lorentz-force electromagnetic principle used in the galvanometer. Piezoelectric designs, such as those marketed commercially by Physik Instrumente Corporation, feature high stiffness, high bandwidth, and a simple mechanical configuration; however, they suffer from limited angular travel, high hysteresis losses, and require sophisticated high voltage drive amplifiers. For example, PI models such as the S-325 and S-330 have a closed-loop bandwidths of only 500 to 1000 Hz and an angular range of only 2 to 5 mrad. [23].

1.4.2 Recent Fast Steering Mirror Development

Fast steering mirrors for space applications have been developed in the past two decades, and a few designs have seen flight service. In particular, Ball Aerospace Company and MIT Lincoln Laboratory are frontrunners in this area. These designs

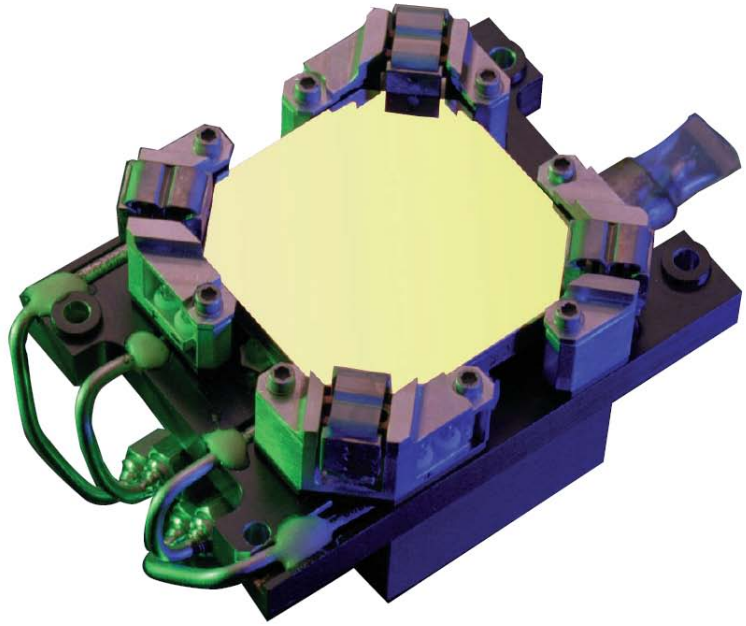

Figure 1-9: A fast-steering mirror which employs voice-coil actuation, built by Ball Aerospace Corporation.

represent the primary technology foundation that forms the basis for the AFSM design.

Ball Aerospace

Over the past several decades, Ball Aerospace Corporation has produced a wide range of devices for air and space flight service, which vary in size, functionality, and performance [2]. The heritage Ball designs typically employ electromagnetic actuation via Lorentz-force (voice coil) drives. Mirror sizes are on the order of tens of millimeters, and typical published bandwidths range from 250 to 1000 Hz, with 1.5 kHz as the highest advertised bandwidth. The devices are typically are constrained to one or two degrees of freedom, with two being the most common. A picture of a typical Ball mirror design is shown in Figure 1-9.

Through research and development, Ball has undertaken several interesting explorations of FSM technology. One such development, detailed in U.S. Patent No. 6,612,192, is a single-frequency conical scanning mirror illustrated in Figure 1-10. This devices employs a mirror and reaction mass coupled together through metal-

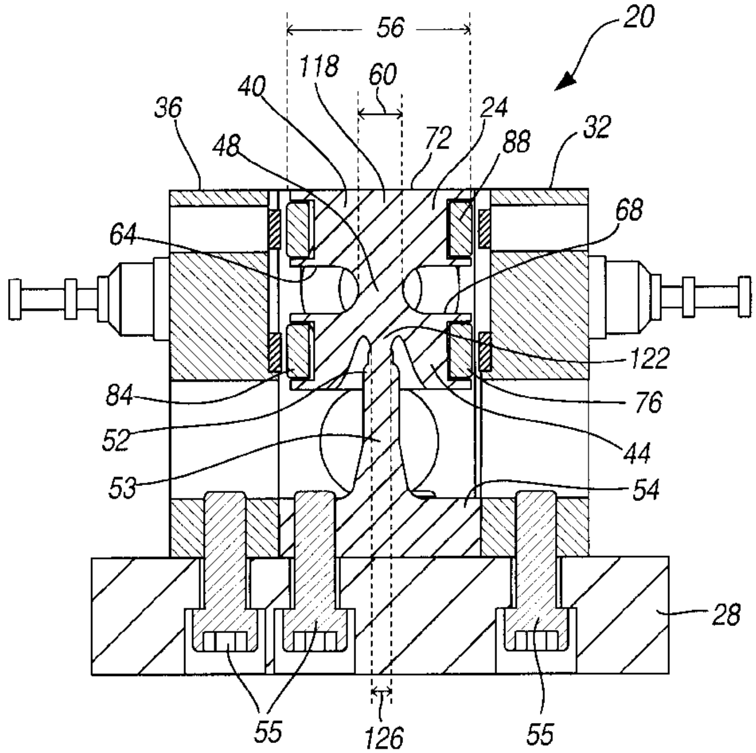

Figure 1-10: Two-axis reaction-free scanning mirror developed by Ball Aerospace (image from Patent No. 6,612,192).

lic flexures, and is designed to be driven at the mechanical resonant frequency with opposing torques via Lorentz-force actuators. The acceleration of the reaction mass cancels the mirror accelerations, thus attenuating undesirable reaction forces into the base structure. Since the mirror is driven at resonance, very little power is required to maintain the scanning action [9].

As part of a separate research project, Ball also developed a fully levitated six-degree-of-freedom mirror, which is documented in several journal articles such as [20]. This mirror is unique in that it employs no flexures for kinematic constraint, and instead relies on active forces from a suite of eight voice coil actuators to levitate, center, and position the moving element. The absence of flexures greatly improves the angular range (± 87 mrad) and positioning accuracy, but comes at the price of a very complex MIMO control architecture and relatively high power consumption, even quiescently, due to the need to actively levitate the moving mirror mass. The published

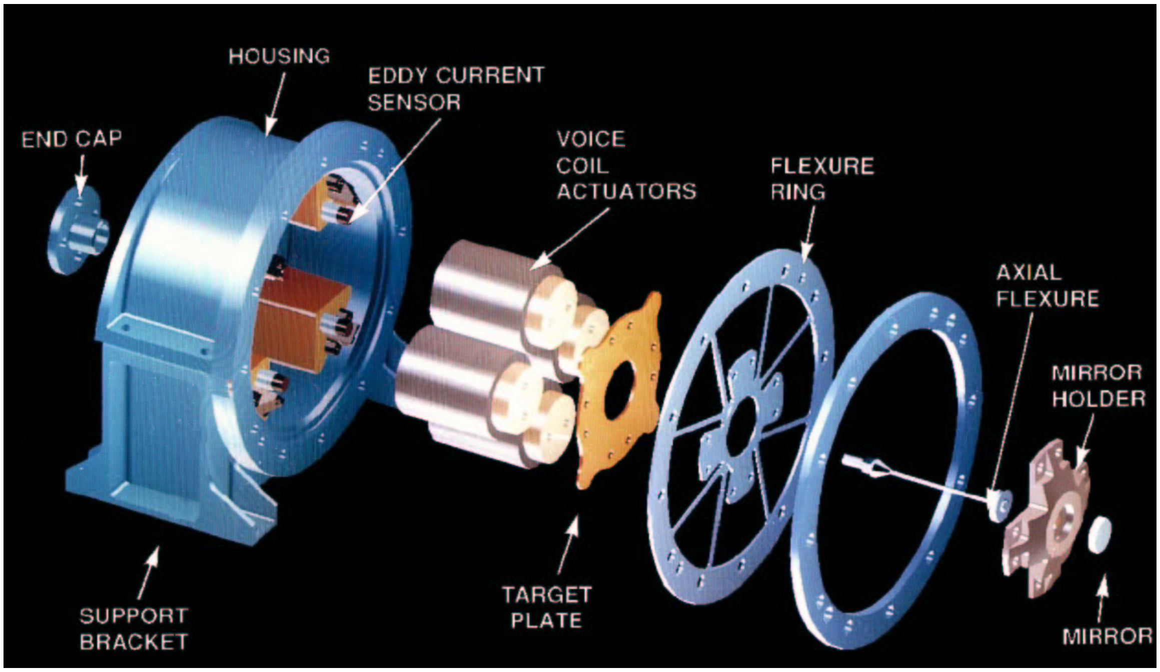

Figure 1-11: Cutaway view of the Lincoln Laboratory High Bandwidth Steering Mirror (HBSM).

small signal bandwidth was rather modest at 600 Hz; however, the prototype used a large mirror with a 5-inch aperture.

MIT Lincoln Laboratory

In contrast to Ball's wide range of developments in FSM technology, FSM developments at Lincoln Laboratory over the last two decades have been more tightly focused. Lincoln's current suite of fast steering mirrors, several of which have seen flight service on various missions, are all based on original development work performed in the early 1990s by Gregory C. Loney. The mirror design has been dubbed the High-Bandwidth Steering Mirror (HBSM), and each subsequent design variant is given a letter designation. The current variant is the HBSM-D. Loney's report, "High Bandwidth Steering Mirror Research"[15], became my primary reference for the AFSM design.

An exploded view of the Lincoln HBSM is shown in Figure 1-11. In this figure the key components which provide the mirror's functionality are illustrated. At the heart of the device are four linear voice coil actuators, which operate on the four quadrants of the moving assembly, with opposite pairs of actuators in a push-pull configuration. To reduce moving mass, the actuators are designed with stationary magnets and cores, and moving coils. The moving coil heads are bolted through a target plate and flexure ring to the mirror holder, which is the main dynamic structure and is manufactured from beryllium. The polished mirror is bonded to the mirror holder. Kinematically, the mirror holder is constrained in torsion and the two lateral translational directions by the flexure ring, and in the axial translational direction by the central axial flexure. Both the axial flexure and flexure ring are made of stainless steel.

Position sensing is accomplished via four eddy current sensors (sometimes called "Kaman sensors" in reference to their manufacturer, Kaman Measuring Systems), which measure the linear distance to the four lobes of the target plate sandwiched between the mirror holder and voice coil heads. All components are supported structurally by a common housing.

In the documentation of his original design, Loney demonstrates good performance numbers, albeit for a small mirror apertures of 10 and 16 mm. The published performance for the prototype design is given in Table 1.2. Subsequent flight designs using larger mirrors and higher design margins operated with reduced performance numbers.

Table 1.2: HBSM-B Prototype Measured Performance

Property	**HBSM-B Performance**
Bandwidth (-3 dB)	10 kHz
Angular Range	±13 mrad
Angular Resolution	0.2 μrad
Angular Acceleration	1.3×10^4 rad/sec^2
Mirror Aperture	10-16 mm
Mirror Construction	Beryllium

1.5 AFSM Performance Goals and Design Constraints

The prior art fast steering mirror designs described in the previous section, while adequate for their intended missions, may not meet the needs of future high-precision optical applications. The goal of this thesis is to demonstrate new AFSM technology that offers a significant improvement in performance over prior designs–most notably in bandwidth and acceleration–and thus improves robustness and flexibility at the system level.

To this end, the original AFSM funding proposal identified the performance goals listed in Table 1.3. As a basis for comparison, the proposal referenced the performance characteristics of the present Lincoln Laboratory HBSM-D design. Note that the D variant is designed with a bigger mirror and larger design margins than the prototype B variant, and thus represents a typical-performance heritage design. In his Master's thesis, Larry Hawe [11] provides details measurements of the HBSM-D performance and explored several different control schemes for an optical system using this mirror.

Table 1.3: AFSM Proposal Performance Goals

Property	**AFSM Proposal Goal**	**HBSM-D Performance**
Bandwidth (-3 dB)	20 kHz	1-2 kHz
Angular Range	±20 mrad	±10 mrad
Angular Resolution	$5\text{x}10^{-8}$ rad	$1\text{x}10^{-5}$ rad
Angular Acceleration	10^6 rad/sec^2	10^4 rad/sec^2
Mirror Aperture	30-60 mm	15-100 mm
Mirror Construction	Beryllium	Beryllium

1.5.1 Limitations of Existing Technologies

To assess whether such an advance in performance is feasible, it is helpful to understand the limiting characteristics of the prior art designs. As discussed in Section 1.4, these designs are categorized based on the actuation principles they employ: Electromagnetic, as in the voice coil or galvanometer designs, and piezoelectric. Other actuation principles, such as magnetostriction and electrostatics, are not in wide use and are therefore not considered here. In his PhD thesis, Lu [16] provides an excellent review of the physical limits of the two main prior art technologies, particularly with regard to force output and bandwidth. His findings are summarized below:

- **Lorentz Force Actuators:** Moving coil actuators (voice coils and galvanometers) based on the Lorentz force are inherently limited in peak force output due to thermal constraints; that is, the metallic conductors used in voice coil design (most commonly copper or aluminum) are subject to maximum current density levels above which the coil overheats. Using established empirical thermal limits, Lu calculates that an ideal air-cooled aluminum coil with no other attached mass loads can sustain maximum accelerations of 109 G (where G is the acceleration due to earth's gravity). The more common copper coil fares much worse, with a maximum acceleration of 42.7 G.

  Although the thermal limitation may be alleviated through the use of convective cooling methods, especially liquid cooling, any practical implementation of such a system will add mass to the coil and tend to counter the resulting gains in allowable current density. Another alternative is to use superconductors in place of the traditional metallic conductor, or to increase the flux density of the external magnetic field in which the coil acts; however, at the time of this writing, superconductors are technologically viable only for the most exotic and expensive applications (such as magnetically levitated trains). Another limitation of superconductors is that they are intolerant of rapid changes in flux density, and are therefore unsuitable for high-bandwidth dynamic applications. Also, flux densities are limited by currently available materials. For example, the best available permanent magnets are limited to a remanence of 1.4 to 1.5 tesla, while the best soft magnetic materials, carrying a flux induced from an electric coil, saturate at about 2.0 tesla.

- **Piezoelectric Actuators:** Mechanisms built around piezoelectric actuators

> are generally characterized by short strokes and relatively high bandwidths. However, a piezoelectric stack is inherently limited by its own internal electrical and mechanical losses, which have the effect of reducing stroke output (and hence, bandwidth) as frequency increases. Thus, direct drive devices are usually limited to bandwidths of around 1 kHz.
>
> Another problem with piezoelectric actuators is the difficulty of building long-stroke devices. Because the proportionality of strain to applied voltage is so small (about 0.1%), increasing the available stroke usually involves the use of a lever-type mechanical amplifier. PI offers a fast steering mirror design, the model S-334 that employs this technique to achieve an angular stroke of 50 mrad. However, doing so adds mass, and the overall stiffness is reduced due to the transmission ratio N^2 and the compliance of the lever. Both of these complications tend to reduce the available bandwidth, and thus, a tradeoff between stroke and bandwidth becomes necessary if the decision is made to use piezoelectric actuators. The S-334, for example, has a bandwidth of only 800 Hz.

In addition to the voice coil actuators, Loney reports in [15] that the metallic flexures used in the HBSM designs also reduced performance. In the original design iterations, the flexures were undamped, and the resonances of the long slender beam members coupled directly into the angular output of the mirror. The first of these parasitic modes occurred around 2 kHz, which limited the achievable bandwidth. The problem was mitigated by redesigning the flexure and adding a layer of viscoelastic damping epoxy to the part. Refer to Hawe [11] for additional details.

1.5.2 Enabling Technologies

As it turns out, two of the contemporary research areas taking place at the Precision Motion Control lab at MIT provided solutions for both the actuator force limitations, and also the problematic metallic flexure resonances.

First, as part of his PhD thesis, Xiaodong Lu developed a normal-force, flux steering actuator configuration that he successfully employed in a single degree-of-freedom fast tool servo system. The actuator relies on normal force principles similar to a solenoid, but has the advantage of being nearly linear in both current and stroke, thus making it easy to control. In his thesis [16], Lu performs a theoretical calculation to show that the accelerations attainable on a steel armature for practically attainable flux densities is up to 4000g, which is 100 times that of the copper voice coil (40g). Also, the actuator configuration lends itself well to stroke ranges of about 10 to 1000 microns, which falls between a piezoelectric actuator (micron-order) and voice coils (millimeter-order). This combination of stroke and force is ideal for the fast steering mirror application.

The second area of research, spearheaded by Augusto Barton and David Cuff as part of their Master's theses ([17], [6]), was the use of viscoelastic materials (specifically, elastomers such as neoprene and silicone) as kinematic bearings in place of traditional metallic flexures. The use of such bearings involves many tradeoffs, as explained in Chapter 2; however, they do provide immediate solutions to the problems encountered by Loney with his metallic flexure design–namely the extra modes and lack of damping.

1.6 Final Performance Goals

The initial proposal goals listed in Table 1.3 were based on the theoretical analysis performed by Trumper and Lu. As part of the proposal, I assessed the performance of a practical device relative to the theoretical numbers and came to the conclusion that although achievable, the project scope to attain the full proposal performance was too large for both the allocated time schedule (approximately 9 months) and available funding. My assessment was based on the fact that considerable design optimization through numerical analysis was necessary, coupled with the fact that the resulting design would require exotic materials and expensive manufacturing techniques. I concluded that the funding and schedule would be better suited to the reduced performance specification given in Table 1.4. By demonstrating this basic level of performance, the project would serve as a demonstration of the fundamental technology from which future optimized designs could be based.

The AFSM hardware described in the remainder of this thesis is designed to meet the performance goals listed in Table 1.4.

Table 1.4: AFSM Prototype Performance Goals

Performance Parameter	**Design Goal**
Bandwidth	5 kHz
Angular Range	± 10 mrad
Angular Acceleration	1×10^5 rad/sec^2
Angular Resolution	8.7×10^{-7} rad
Mirror Aperture	20 mm
Mirror Construction	Aluminum

1.7 AFSM Operational Concepts

Before discussing the mechanical and electromagnetic details of the AFSM design, I provide in this section an overview of the basic hardware configuration and operational concepts. This discussion may be used as a point of reference for the analysis presented in subsequent chapters.

A CAD rendering of the final AFSM design iteration is shown in Figure 1-12. The mirror which comprises the heart of the optical system is located in the center, and is drawn in red. The two principal actuated degrees of freedom, θ_x and θ_z, are indicated by the dashed lines running through the mirror center.

Arrayed around the mirror are four identical actuators, which act on magnetic armatures bonded to the mirror edges. Each actuator applies a force to the armatures (and hence, the mirror edges) in either the positive or negative Y direction. By applying equal and opposite forces at opposing actuator edges, a torque is developed which moves the mirror in one or both actuated DOFs. Both the mirror and actuators are mounted to the machine base, which provides structural support for the entire assembly. The actuators are rigidly mounted to the base, while the mirror is

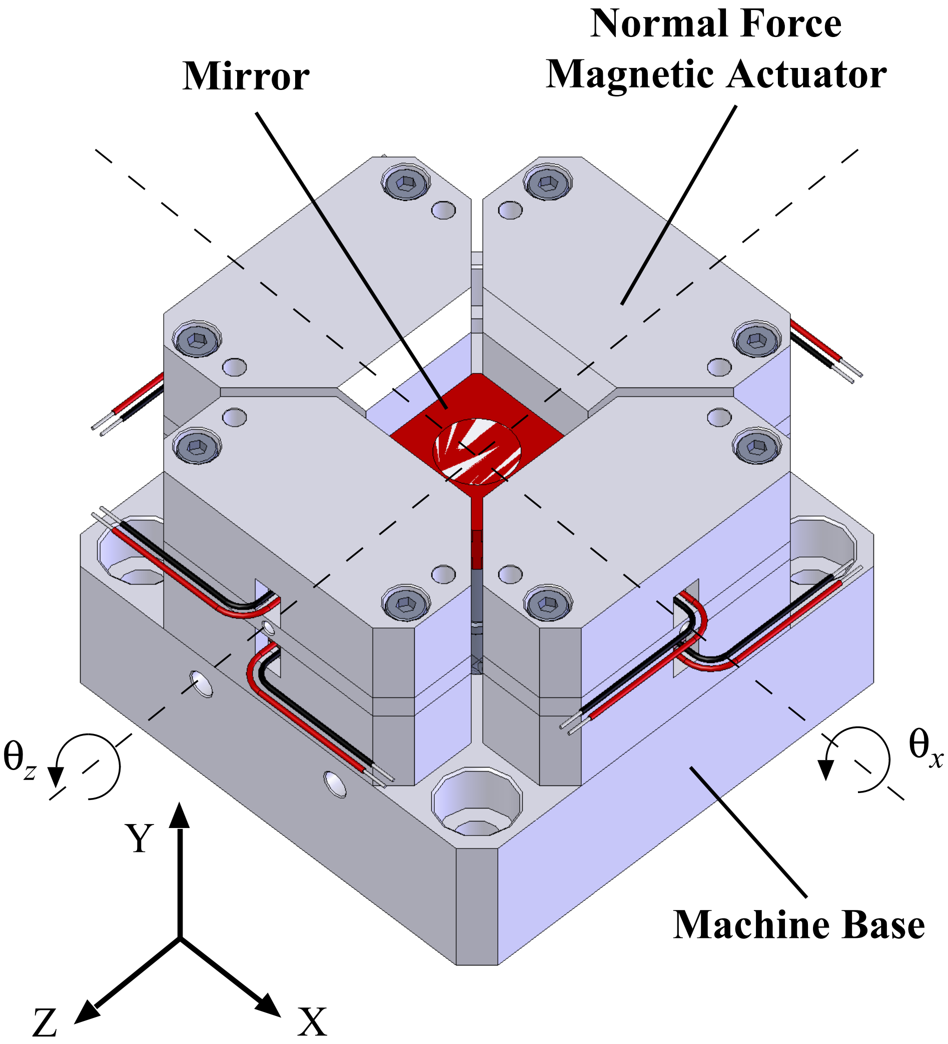

Figure 1-12: Isometric CAD rendering of the AFSM showing the actuators arrayed around the central mirror. The dashed lines indicate the rotation axes of the mirror.

suspended relative to the base by a set of flexures.

A clearer understanding of the AFSM's internal operation can be obtained by viewing Figure 1-13, which illustrates the AFSM hardware in cross-section. This view is obtained by conceptually slicing Figure 1-12 vertically through either the θ_x or θ_z axis. Here the internal components of the actuators are revealed, as well as the flexure support structure for the mirror. There are two types of flexures, which in concert constrain the four non-actuated DOFs. The axial flexure constrains the mirror in the Y translation DOF, while the four elastomeric flexures constrain X and Z translation and θ_y rotation (note that two of the four elastomeric flexures are out-of-plane and are therefore not shown). Also illustrated are the capacitance probes that measure the mirror displacements relative to the housing, and the associated clamping mechanism that holds them in place. The mechanical assemblies and design are covered in greater detail in Chapters 2 and 4.

The architecture of one of the AFSM actuators is displayed in Figure 1-14. The actuator consists of two symmetric halves, upper and lower, each of which contain an electrical coil wrapped around a magnetically permeable core. Each coil and core is rigidly potted in place to the actuator housing, which is displayed in green in the figure.

Between the two actuator halves is a backiron piece which mounts flush against the rear pole faces of the cores. This piece provides a return path for the magnetic fluxes generated in the cores. At the heart of the actuator is a powerful neodymium-iron-boron (NdFeB) permanent magnet. The magnet provides a biasing magnetic flux to the armature. The DC flux from the magnet is steered to one core half or the other by the superposed flux from the electrical coils. For a more detailed treatment of the actuator operation, see Chapter 3.

With the basic layout and operational principles in hand, I now proceed with the detailed design treatment. The mechanical design elements are discussed in the following chapter, while the electromagnetic actuators are analyzed in Chapter 3. The full mirror assembly and supporting hardware are discussed in Chapter 4.

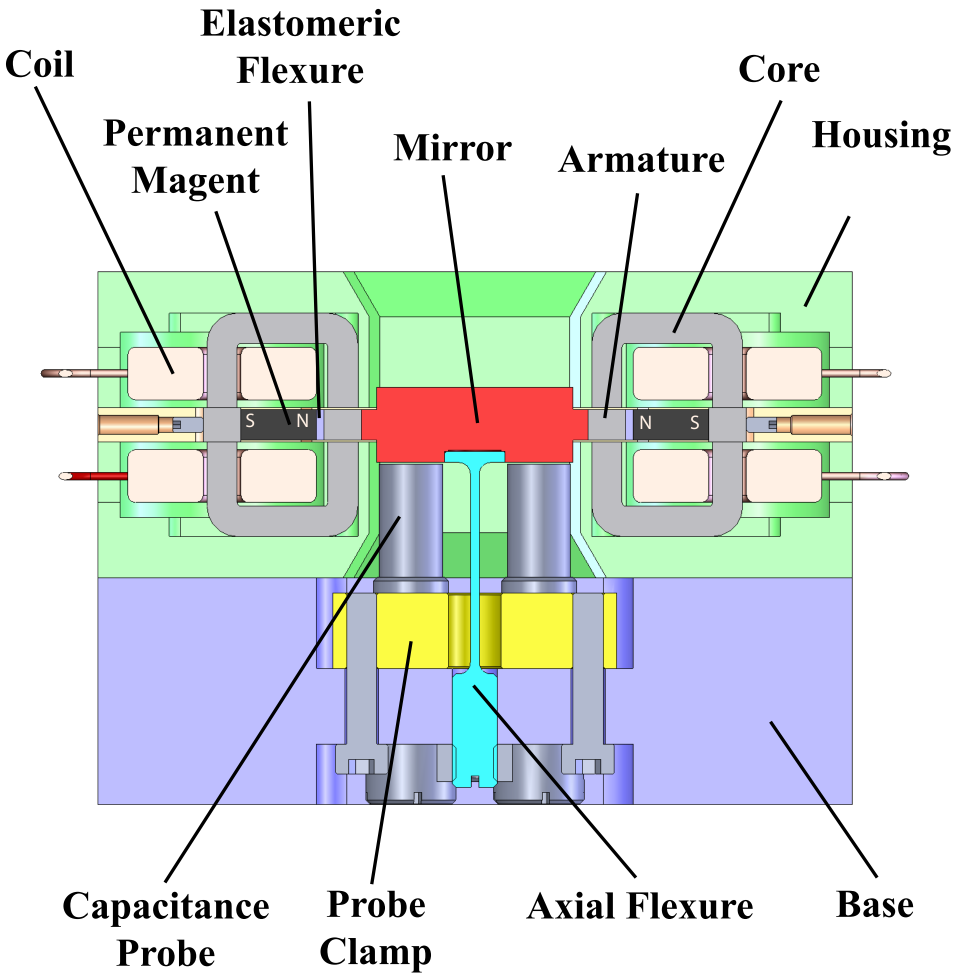

Figure 1-13: Cross-sectional view of the AFSM showing the mirror, flexure supports, capacitance probes, and actuators.

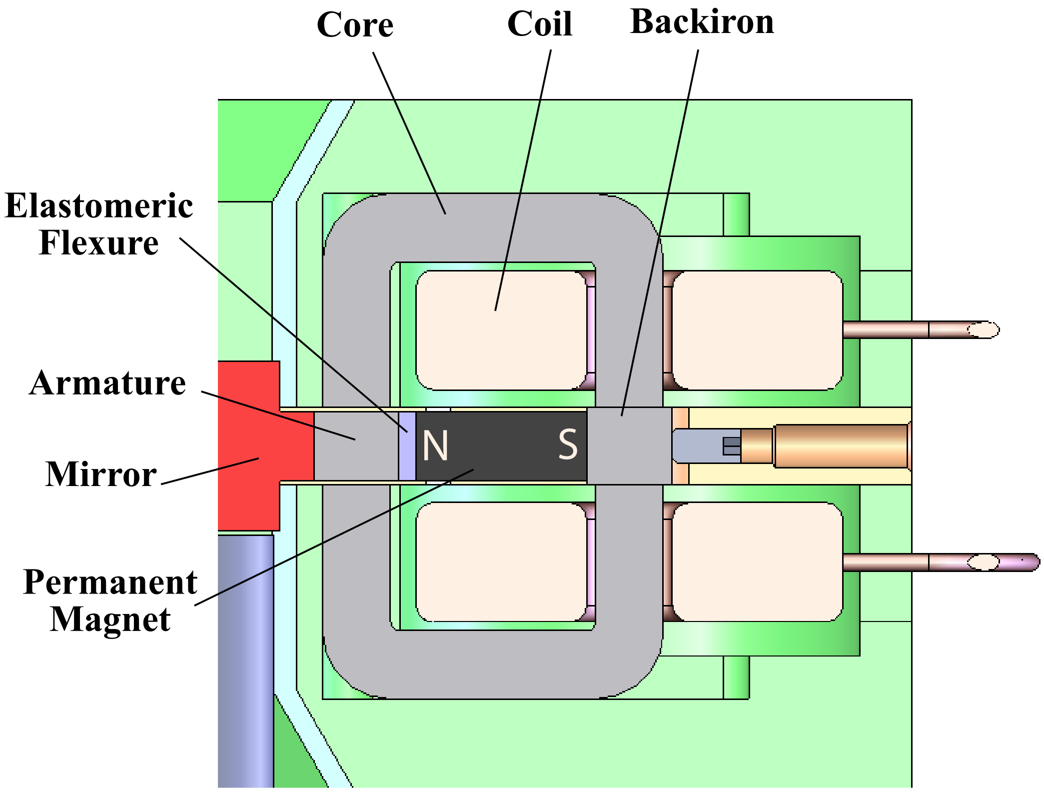

Figure 1-14: Cross-sectional view of the magnetic actuator showing component detail.

Chapter 2

Mechanical Design

A successful mechanical design is the fundamental factor which determines the ultimate performance of the AFSM system as a whole. The primary design emphasis must be mechanical, because the system output in terms of optical beam direction and stability is directly determined by the speed and precision with which the mechanical components are positioned. This may seem like a trivial statement until one considers the large dynamic range under which the mechanical system must perform. When considering precision mechanical operations through frequencies spanning tens of kilohertz, the simplifying assumptions of rigid body dynamics employed in most machine designs become inadequate. The true flexible body dynamics not only contribute to high-frequency positioning errors, but may also destabilize the control system if not identified and addressed in the system design. Also, since the actuators, sensors, and electronics comprising the remainder of the system function solely to drive, measure and compensate the mechanical elements, the AFSM design must begin here[1].

2.1 Control System Considerations

Before delving into the mechanical design details, however, it is prudent to take the broader viewpoint of the control systems engineer, and consider the influence

---

[1] Moreover, since I am a mechanical engineer by training, I am motivated by my misguided sense of vocational superiority to make the mechanical design my primary focus.

of the mechanical plant on the remainder of the AFSM system. The accuracy and bandwidth performance requirements for the AFSM design necessitate the use of closed-loop control to improve the natural dynamics of system. A critical step in a successful controller design is the mathematical derivation, and eventual experimental verification, of an accurate open-loop dynamic model of the mechanical plant to be controlled. However, given that the design is brand new, an opportunity exists not only to simply identify the natural dynamics, but to actually prescribe them in an advantageous way, such that a simple controller can be used to achieve the required performance. This is done, of course, by performing the mechanical design with the desired dynamics as initial design goals.

A key insight gained from studying control systems theory and systems modeling is that as the disparity between the natural plant dynamics and the desired open-loop dynamics increases, the difficulty of designing a compensator to achieve the desired performance increases proportionally. Furthermore, even though the designer may arrive at a controller design that appears tractable mathematically, physical hardware limitations (for example, amplifier saturation) may make practical implementation of the design difficult if not impossible. Conversely, however, if the natural plant dynamics are very close to the desired dynamics, designing and implementing a successful control architecture becomes a straightforward task.

As an example, consider the plant described by the following equation, and the Bode plot given in Figure 2-1.

$$G(s) = 1 \times 10^5 \frac{s^3}{(s+10)^3(s+100)^2} \tag{2.1}$$

This plant represents an AC-coupled system; that is, a system with zero response at DC (zero frequency). An example of a physical AC-coupled system is an electrical transformer. Since its output relies on the establishment of a time-varying magnetic field within the input and output windings, applying a DC signal to the input produces no output. In the example above, the plant dynamics contain the combined output of several AC-coupled systems, which together produce the three zeros at the origin[2].

---

[2] The transfer function used in the example is derived from an active isolation system studied by

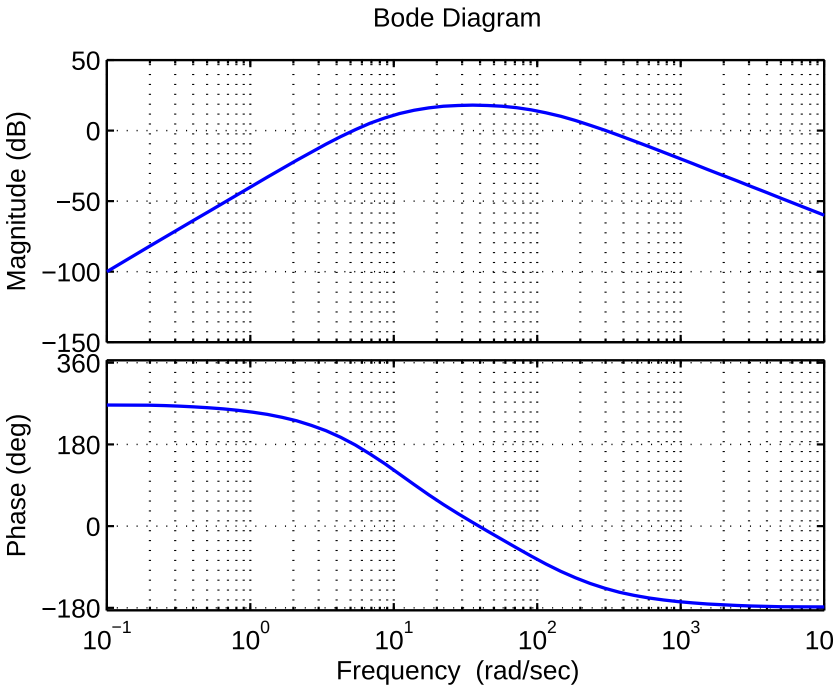

Figure 2-1: Example of an AC-coupled system

Suppose we wish to control this system across a bandwidth spanning from DC to 1000 rad/sec. From a pure mathematical standpoint, this may be accomplished without great difficulty. We simply use three poles at the origin to cancel the plant zeros, and three zeros in the mid-band to boost the system gain around these frequencies. Finally, we add a lead network for high frequency phase margin, and adjust the controller gain for the appropriate crossover. The transfer function for one such controller is

$$C(s) = 23\frac{(s+20)^2(s+50)(s+200)}{s^3(s+2000)} \tag{2.2}$$

which when combined with the plant model results in the compensated system loop transmission shown in Figure 2-2. The compensated system displays excellent gain in all frequencies from DC through crossover, and a robust phase margin of about 55 degrees at 1000 rad/sec.

Unfortunately, physically implementing the controller as designed is quite impos-

---

Wouter Monkhorst [21] as part of his master's thesis.

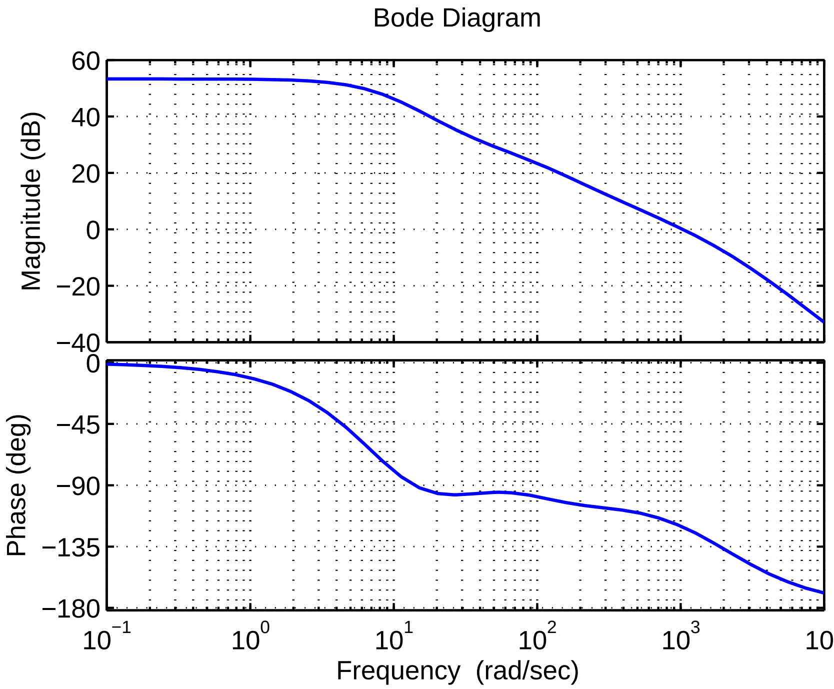

Figure 2-2: A mathematically tractable, but physically impossible control scheme for the AC-coupled system

sible! The trouble lies in our attempt to cancel the three zeros at the origin with controller poles. Remember that physically, the fact that the system is AC-coupled means that it produces *zero* output at DC, regardless of the magnitude of the input. Here, by using the three integrator poles at the origin, we have designed a controller that attempts to apply an *infinite* control effort to the system to produce a finite output–which is, of course, impossible. In terms of control theory, we have made the plant poles unobservable in the feedback measurement. In a real system, applying a DC reference command would quickly drive the three integrators into saturation, most likely at the power amplifier. Since the saturation level is finite, the system output would be zero.

This example also illustrates the prior point about the physical difficulties of using a controller to modify the plant dynamics. In reality, since we cannot use pure integrators to simply cancel the plant zeros, we would instead use a triple-lag controller, with poles at low (but nonzero) frequencies, to widen the controlled frequency band as much as possible. However, as the lag poles are moved lower in frequency, the control effort required to produce meaningful output at that frequency increases. Eventually we will reach a point where the physical limitations of the control hardware do not allow further reduction in the band pass frequency–and even if physically possible, the controller (specifically, the power amplifier) would likely be very expensive and consume a large amount of power.

2.2 Mechanical Design Goals

With these insights in mind, we return to the AFSM mechanical design. Here, to make the plant as controllable as possible, it is desirable to achieve several goals. First, since the mirror is controlled in two degrees of freedom, creating a mechanical design that decouples these degrees of freedom allows the use of two independent single-input, single-output (SISO) controllers rather than a more complex multi-input, multi-output (MIMO) architecture. The advantages of SISO controllers are numerous, but their most important characteristic is the fact that they can be readily designed using experimental frequency response model data. In addition, it is far easier to design and verify a robust controller using SISO techniques.

If we consider the situation from a purely mathematical standpoint, it may be tempting to ask why such an emphasis on mechanical decoupling is warranted. After all, if the controlled system is linear, it should be possible to apply a transformation matrix to the two signals in order to decouple them and thus treat them separately. However, the downfall again comes when trying to implement such a transformation practically. First, for the transformation to be valid, the mathematical model of the plant must be very accurate–which may not be true for many real systems, or systems subject to varying plant parameters. Second, implementing the transformation requires additional complex circuitry in an analog system, or more computation time over the sampling interrupt in a digital system. More importantly, the outputs from the transformation itself may require a large percentage of the system's limited control authority, and therefore in turn limit the overall control performance. It is far better then to perform the decoupling mechanically, so that we can take advantage

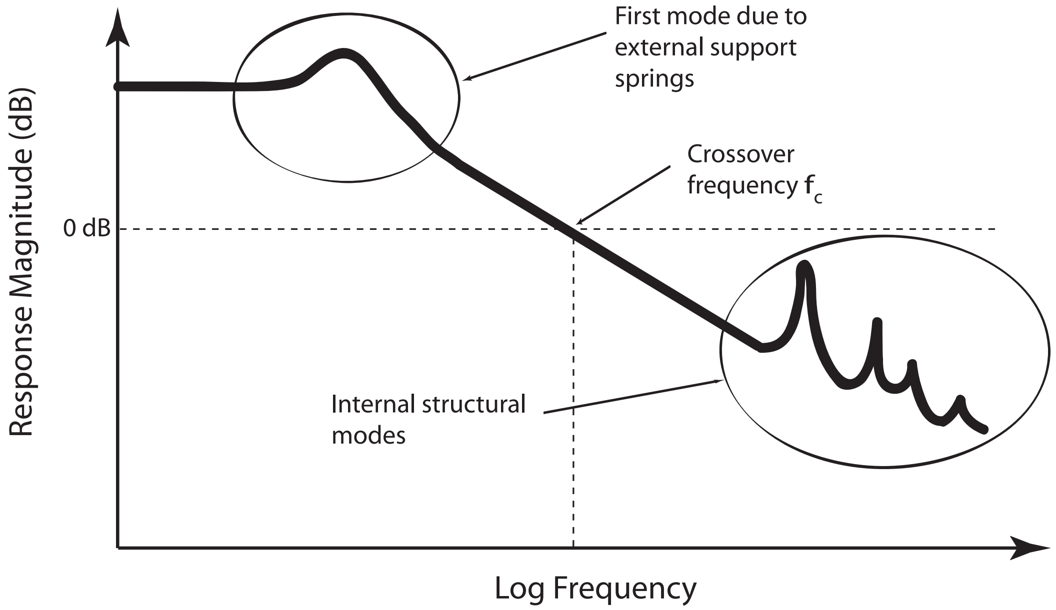

Figure 2-3: Idealized Open-Loop Plant Frequency Response

of its benefits on the linear system behavior “for free”.

As a second goal, it is desirable to maximize bandwidth. For this, we desire a plant frequency response similar to the one sketched conceptually in Figure 2-3. Note that here we are making the assumption of a decoupled plant, and thus consider a single rotational degree of freedom only. By maintaining the spring rates of the mirror flexure supports at low values in the actuation direction, we can place the suspension mode[3] of the mirror and armatures (indicated by the first resonant peak in the Bode plot) below the crossover frequency f_c. This results in high gain values throughout the majority of the controlled frequency band and a -40 dB/decade slope through f_c, ensuring a well-defined crossover point. Then, through proper design of the mirror and armature structure, we will attempt to place the resonant frequencies of the structural (non-suspension) modes well above the crossover point. As the difference in frequency between the first structural mode and f_c increases, the control loop will acquire additional gain margin and increased robustness.

---

[3] I originally referred to this first mode of the mirror resonating on its flexure supports as the “rigid body” mode, because the mirror and armature structure does not displace relative to itself during excitation at this resonance. However, Rankers [24] suggests that these should be termed “suspension modes” to avoid confusion with true rigid body modes (simple translations and rotations in 3D space), which have resonant frequencies of zero Hz.

Of course, the prototype requirements also specify high angular accelerations (which is an implicit part of the bandwidth specification). To meet this goal, it is necessary to design the actuators with high peak force, while simultaneously minimizing the inertia of the moving mirror. Since the actuator force cannot be increased arbitrarily, minimizing inertia is a primary goal. However, this in turn raises the natural frequency of the suspension mode, which is contrary to the objective of the previous paragraph. In order to maintain the suspension resonant frequency at a given value, any decrease in inertia must be accompanied by a commensurate decrease in support stiffness.

The above discussion leads to one of the classic problems in kinematic flexure design; that is, the challenge of obtaining compliance in the actuated degree of freedom while maintaining sufficiently high kinematic stiffness in the constrained degrees of freedom. I discuss such tradeoffs for the AFSM in Section 2.4.

2.3 Sizing for Desired Mechanical Performance

With the previously discussed goals in mind, the next task is to conceptualize the individual AFSM mechanical components to meet the stated objectives.

2.3.1 Mirror and Armatures

The design starts with the mirror itself. From an optical standpoint, the mirror must have sufficient frontal surface area to accommodate the required beam aperture over the desired incident and reflected angles. Secondly, its construction must be able to accommodate the kinematic supports that surround it, as well as the actuators that drive it and the sensors that measure its position. At the same time, minimizing the rotational inertia and mass of the mirror while maximizing its flexural stiffness is critical, as doing this will meet the goal of placing the suspension mode at a low frequency and the structural resonances at high frequencies, as described above. Finally, it is also desirable to design the mirror with symmetrical geometric and mass properties about the two controlled rotational axes. Such symmetry aligns the

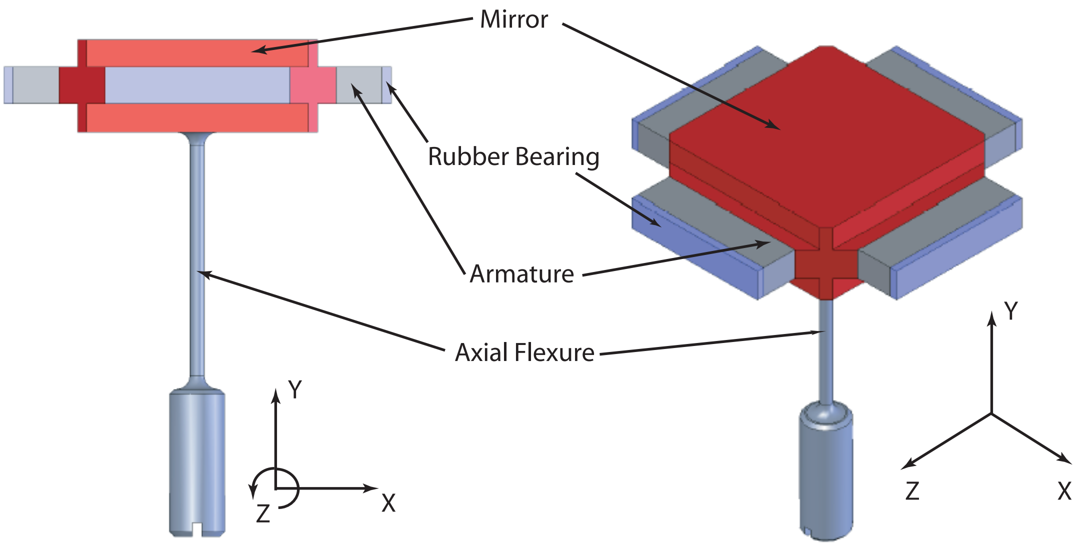

Figure 2-4: CAD rendering of the AFSM moving mirror and armature design, along with its flexure supports.

geometric center of rotation with the center of mass, which accomplishes the original goal of decoupling the two rotational degrees of freedom of interest (in mathematical terms, such a design diagonalizes the mirror's inertia matrix and places its principal axes coincident with the actuated degrees of freedom).

A configuration that addresses these considerations is the simple square prismatic geometry shown in Figure 2-4, representing the final AFSM mirror and flexure design. In this view the mirror, armatures, and kinematic supports (axial flexure and rubber bearings) are shown. The basic mirror envelope is 30x30x10 mm in the X, Z, and Y axes, respectively. The 30 mm lateral dimensions allow for a large optical beam aperture, and the 10 mm normal dimension imparts the required structural stiffness in all deformation modes. Here of course, the primary modes of concern are bending along the X and Z axes, because they act directly on the controlled degrees of freedom. However, it is reasonable to expect that other excitable or observable modes (for example, a torsional mode along the X or Z axis) may couple into the system outputs as well, and therefore it is prudent to design for a general level of stiffness throughout the entire mirror.

In order to meet the dynamic performance requirements, choosing a low density, high modulus material for the mirror will minimize its rotational inertia and, along with the stiff geometry, maximize its flexural stiffness. Several good candidate materials are identified immediately by consulting Figure 2-5 from Ashby [1]. The figure plots density ρ versus elastic modulus E for a wide variety of materials. Since the optimal stiffness-to-mass ratios are obtained by minimizing ρ and maximizing E, the best materials are located in the upper left area of the chart. Here it is easy to see that elemental beryllium, certain composite materials, and ceramics such as silicon carbide and silicon nitride are excellent candidates. However, since this is a fatigue application, the ceramic materials are a poor choice due to their brittleness and low fracture toughness. Also, the composite materials are expensive and difficult to manufacture. Beryllium is readily machined, but due to toxicity concerns it is also very expensive and must be outsourced. Due to funding and schedule concerns, I made the decision to proceed with a less ideal, but readily available and inexpensive material: 6061-T651 aluminum alloy. Future designs using beryllium would enjoy an increase in stiffness-to-mass ratio by about a factor of five over aluminum, yielding additional performance and robustness gains[4].

Arrayed around the edges of the mirror are four rectangular prismatic armatures. Although simple in geometry, the design of these parts is critically important for two reasons. First, because they are acted upon magnetically by the normal force actuators, the faces and volume of the armatures must be large enough to sustain the actuators' magnetic field without saturating. This problem is discussed further in Chapter 3. Second, from a dynamic standpoint the armatures are the most significant contributor to the moving structure's moment of inertia, and as such this contribution must be minimized. Large armature inertias are unavoidable, however, because their density and volume are constrained to minimum values by the magnetic

---

[4] It is worthy of note that the Lincoln Laboratory flight qualified HBSM-D heritage design used a beryllium mirror optimized via FEA. Although not performed here for programmatic reasons, it is reasonable to expect that similar measures applied to the AFSM would result in substantial performance increases relative to those published in this thesis.

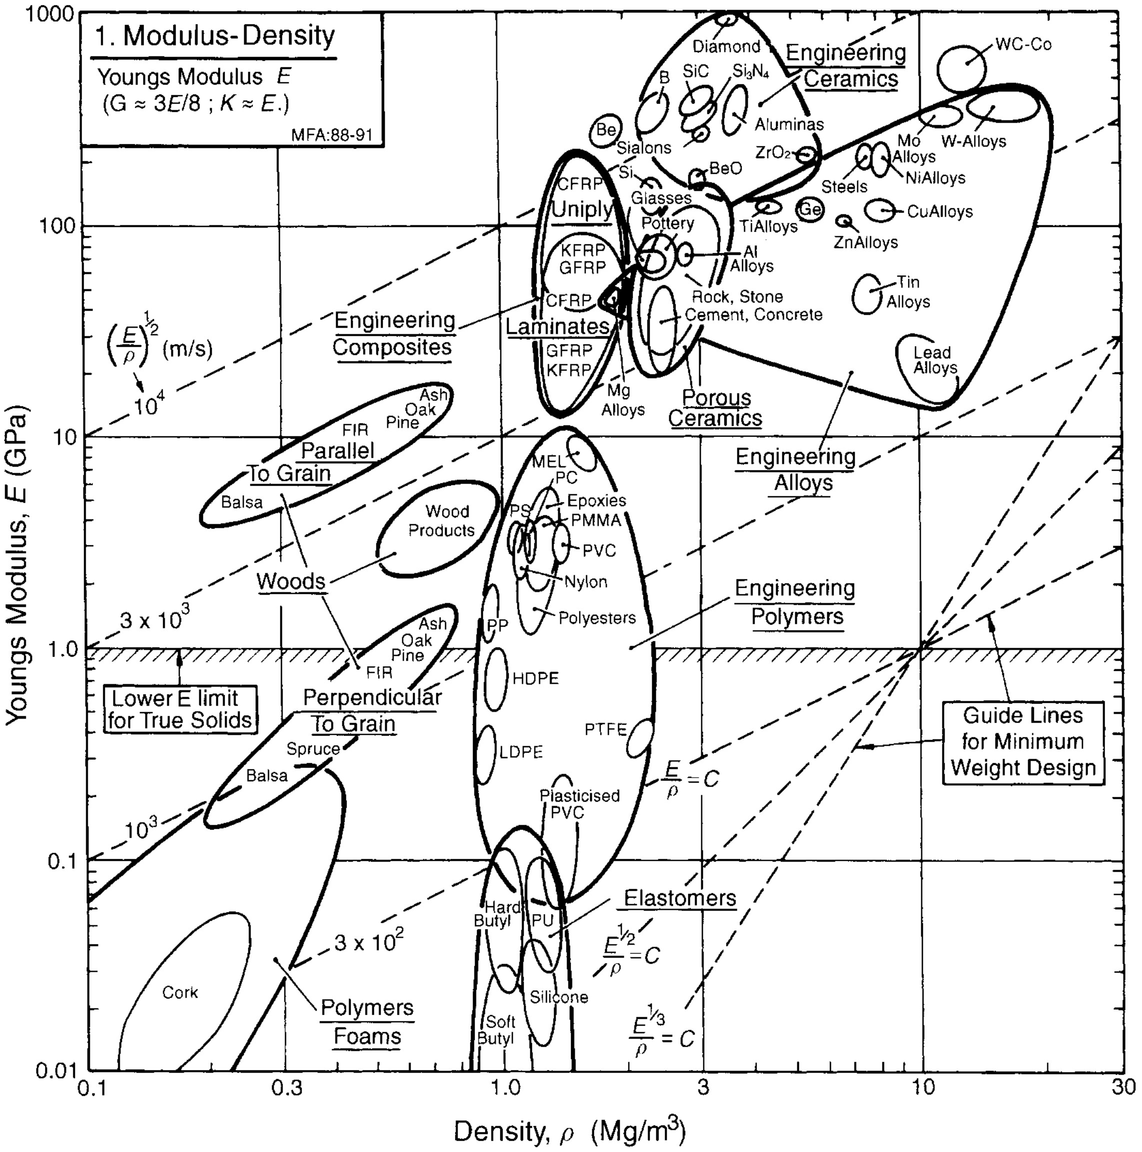

Figure 2-5: Elastic modulus E plotted against density ρ (from Ashby [1]).

material properties, as well as the required actuator force output. Also, the armatures must be placed at large radii relative to the center of rotation of the mirror. This is to accommodate the optical beam aperture, while ensuring that the torque produced by the actuators acts through the mirror's center of mass to produce decoupled motion. Despite these complications, it is still advantageous to minimize the armatures' distance from the center of rotation as much as possible, because although the torque output increases with increasing radius, the moment of inertia due to the armatures increases as the *square* of the radius. As a simple verification, for a given actuator force output F and armature mass m, Newton's second law as a function of the radius r is

$$\tau = I\alpha \implies Fr = mr^2\alpha \implies \alpha = \frac{F}{mr} \tag{2.3}$$

where the resulting angular acceleration is denoted by α. Thus, it is easy to see that keeping the radius small results in larger angular accelerations.

To minimize the radius, I chose a long and narrow geometry for the armature of 20x4x5 mm. I chose SM-2 HB sintered powder alloy for the armature material due to its magnetic permeability, low magnetic power loss, easy machinability and availability in our lab[5]. It also provides a small inertia advantage because its density is slightly lower than that of solid steel (7.5 g/cc as opposed to 8.0 g/cc for steel).

The normal force magnetic actuators described in Chapter 3 produce force along the Y-axis of each armature (using the coordinate system of Figure 2-4). Applying opposite forces to opposing armature pairs produces accelerations in the desired degrees of freedom, i.e. θ_X and θ_Z.

A photograph of the fabricated mirror and armature is shown in Figure 2-6. Note that the critically important dynamic analysis of the moving components has not yet been discussed; I will cover this in section 2.5.

---

[5] SM-2 is a specialized powder metal alloy produced by Mii Technologies, LLC (www.miitechnologies.com). The alloy consists of iron particles encased in a polymer binder and sintered into bars or net shapes. The binder insulates the individual particles, reducing eddy current power losses. Unfortunately, Mii is no longer in business; however, at the time of this writing the material is still available in the Precision Motion Control lab.

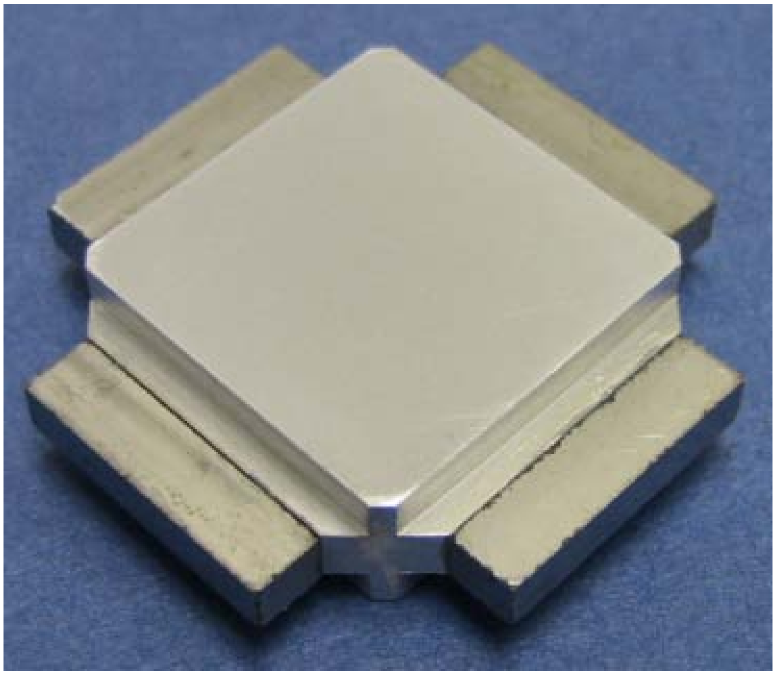

Figure 2-6: Fabricated AFSM mirror and armature assembly.

2.4 Flexure Design

With the moving element design established, we now turn to the problem of designing flexures to kinematically constrain the mirror in the non-actuated degrees of freedom, and to provide compliance in the actuated degrees of freedom.

A rigid body in three-dimensional space possesses six degrees of freedom, yet the desired mirror output acts in only two (i.e., θ_X and θ_Z). Accordingly, it becomes necessary to "ground" the remaining four degrees of freedom using stiff couplings to the machine base, which is modeled as rigid. Note that in the ideal case, the torques resulting from the actuator forces would act purely on the controlled rotational degrees of freedom, making the constraining forces from the couplings unnecessary. However, no system can be made so perfect, and therefore the kinematic constraints become essential for preventing unwanted motion.

As shown in Figure 2-4, I employ two types of flexures in the AFSM design. The metallic axial flexure constrains the axial translation DOF in the Y axis, while remaining relatively compliant in the remaining five DOFs. Meanwhile, the elastomeric rubber bearings arrayed around the edges of the armatures constrain the X and Z translational DOFs, as well as rotation in yaw (θ_Y), but are compliant in Y-translation and θ_X and θ_Z rotation.

The details of the flexure designs are discussed below. I begin with the axial flexure.

2.4.1 Axial Flexure

The design of the metallic axial flexure is heavily influenced by the success of a similar design used in the Lincoln Laboratory HBSM series fast steering mirrors, developed and flown throughout the 1990s. The design of this flexure is of great importance due to the fact that the degree of freedom it constrains (Y-translation) is parallel to the line of action of the actuators. Any mismatch in force output between pairs of actuators will produce a net translational force on the mirror that must be reacted by the axial flexure with minimum resultant deformation. Since such axial deformations introduce errors into the optical system, a stiff constraint in this direction is critical to the AFSM's performance.

The main difference between the AFSM axial flexure and the heritage design is that the AFSM actuators produce much more force than the HBSM voice coils, and therefore the flexure must be proportionally stiffer to keep unwanted deflections low. I started the design by assuming a 10 percent mismatch between actuator pairs and allowing a maximum 1 μm deflection in response to this force. At maximum current, the actuators are designed to produce about 40 newtons, and therefore the axial flexure stiffness must be $4/1 \times 10^{-6} = 4 \times 10^{6}$ N/m.

Competing with the axial stiffness requirement is the necessity to keep the stresses in the flexure low while undergoing bending excursions through the mirror's angular travel–that is, the flexure must be compliant in bending. To frame the problem, we first begin with the basic equations for a slender elastic beam of uniform cross section under tensile (or compressive) loading and bending. For verification purposes, we will also calculate torsional effects. These equations may be found in any strength of materials text (see for example, [4]).

The equations for axial stiffness k_{ax}, bending stiffness k_b, and torsional stiffness

k_t are:

$$k_{ax} = \frac{AE}{L} \tag{2.4}$$

$$k_b = \frac{IE}{L} \tag{2.5}$$

$$k_t = \frac{JG}{L} \tag{2.6}$$

where A is the cross-sectional area normal to axial loading, E is the elastic modulus of the material, L is the flexure length parallel to the axial load, I is the area moment of inertia normal to the neutral axis, J is the polar moment of inertia of the cross section, and G is the material shear modulus. Examining the equations, it is immediately apparent that the *ratio* of bending to axial stiffness depends only on the ratio between A and I. Thus, since the goal is to maximize axial stiffness while minimizing bending stiffness, the design must minimize the area moment of inertia while keeping the cross-sectional area as high as possible.

For bending in orthogonal axes, this goal is optimized by using a flexure of solid circular cross section of diameter d, in which case the area A is of course $\frac{\pi d^2}{4}$. Also, we have $I = \frac{\pi d^4}{64}$, and $J = \frac{\pi d^4}{32}$. The ratio of A to I is then $\frac{16}{d^2}$, which indicates that the value of d should be chosen large enough to meet the minimum axial stiffness requirement, but no greater.

For a given angular displacement θ, the bending moment M and maximum stress σ_{max} are given by

$$M = \theta k_b \tag{2.7}$$

$$\sigma_{max} = \frac{K_t M d}{2I} \tag{2.8}$$

where K_t is a stress concentration factor.

Because the flexure will be subjected to compressive axial loads it is also wise to consider the possibility of buckling. For a column subjected to such a load, buckling becomes possible when the load exceeds a certain value F_{cr}, given by

$$F_{cr} \geq \frac{\pi^2 EI}{L^2} \tag{2.9}$$

Finally, since the AFSM was to be subject to an unknown number of cycles (and since I suspected I might drive it unstable during verification of my controller designs–which of course I did) I thought it wise to design the flexure for infinite fatigue life. Choosing steel as the flexure material allows for this; other materials such as aluminum may be subject to a finite lifespan. Also for fatigue purposes, a corrosion-resistant material with high fracture toughness is desired.

An ideal material for this application is one that I have used successfully in prior spring designs: 17-4 PH corrosion resistant (i.e. stainless) steel. 17-4 enjoys good fatigue strength, high fracture toughness, excellent corrosion resistance, and is precipitation hardened to a number of different strength values through an aging heat treatment. Consulting a good materials reference [25], we find the relevant material properties, which are listed in Table 2.1[6].

Table 2.1: 17-4 PH Material Properties (Condition H900)

Property	**Value**
Elastic Modulus E	196.5 GPa
Shear Modulus G	77.2 GPa
Fatigue Strength σ_a	276 MPa

With the design equations established and a material selected, the final task was to iterate on the flexure length and diameter to arrive at an acceptable design. Note that I deliberately took a conservative approach here by assuming a stress concentration factor of 2.0, and an angular deflection of 20 mrad rather than the actual 10 (which

---

[6] The fatigue strength listed in Table 2.1 assumes a fully reversed stress amplitude and $K_t = 2.0$.

effectively imparts a safety factor of 2.0). The final flexure parameter values are indicated in Table 2.2 below. The design has excellent fatigue stress margin, and the critical buckling load is almost 200 times the worst case expected axial load, which effectively makes it a non-issue.

A photograph of the fabricated axial flexure is shown in Figure 2-7[7].

Figure 2-7: AFSM fabricated axial flexure

2.4.2 Elastomeric Bearings

Although commonly in use for decades in bridge bearings, automobile suspensions, and the like, the use of elastomers in precision machine applications is a relatively recent development. In their theses, Cuff [6] and Barton [17] provide detailed treatments of the history, properties, and the precision engineering applications associated with elastomer flexures and bearings, and as such I refer the interested reader to them. Here, I wish to provide a qualitative comparison of elastomers to the more traditional metallic flexure in order to explain why I chose them for the AFSM design. Following this I will discuss the selection and design details of the AFSM neoprene flexure.

---

[7] As with most of the AFSM parts, I machined the axial flexure myself. This part was the most challenging of the set, and I am rather proud of how it turned out.

Table 2.2: Axial Flexure Design Values

Property	**Value**
Length L	25 mm
Diameter d	1.5 mm
Stress Concentration Factor k_t	2.0 mm
Axial Stiffness k_{ax}	1.39×10^7 N/m
Bending Stiffness k_b	1.95 N-m/rad
Torsional Stiffness k_t	1.54 N-m/rad
Maximum Bending Moment M	0.039 N-m
Maximum Stress Amplitude σ_{max}	235.8 MPa
Design Stress Margin	0.17
Critical Buckling Load	771 N

The most common bearing element choices in precision machine designs are metallic flexures such as the one discussed in the previous section. Many stage positioning systems (for example, lithography or atomic force microscope stages) are supported entirely by such flexures. The Lincoln Laboratory HBSM-D fast steering mirror is also an all-flexure design. The advantages of metallic flexures are numerous, which of course explains their use in precision machine designs. They include:

- **Well-defined material properties.** Steels and other metals have deterministic, predictable material properties (in particular, density and elastic modulus) which allows the design of flexures with precise and reproducible spring rates and masses.

- **Low friction and hysteresis.** Metals are inherently lightly damped, and flexures are easy to weld or otherwise clamp to prevent slipping. This allows for accurate, repeatable positioning of the controlled system, even in open loop.

- **Resistant to environmental changes.** Metallic flexures maintain their properties even in extreme and rapidly changing environments, such as temperature, pressure, vacuum, exposure to harsh chemicals, and even radiation.

- **Can be designed for infinite life.** With proper design, steels and other metals have demonstrated infinite fatigue life in both testing and field service.

When considering high-frequency dynamic operations, fatigue cycle counts build quickly, and therefore designing for infinite life is often important.

Despite these advantages, however, metallic flexures do have drawbacks, many of which can be remedied through the use of elastomeric bearings. Among them are:

- **Isotropic material properties.** In structural designs it is generally desirable to employ materials with isotropic properties; however, considering the design of flexures, the goal is often to provide low stiffness in one or more directions while maintaining high stiffness in all the others. If the material is isotropic, only the flexure geometry can be used to obtain the desired stiffness ratio.

- **Large working volume.** Because the stiffness properties of metallic flexures are so geometry dependent, in order to achieve the required ratios between stiff and compliant degrees of freedom, flexures are often designed as long thin blades, similar in profile to a diving board (for examples, see [15], [16], or [22]). The compliant direction is along the thinnest dimension, and thus the flexure moving through its working range requires a large volume. As an example, out of the approximately four-inch frontal diameter of the HBSM-D, about three inches were required to accommodate the bipod ring containing the radial flexure supports for the mirror.

- **Low internal damping.** Although the low damping property of metallic flexures was listed above as an advantage in terms of positional precision, it is also may be a detriment when considering the high-frequency system dynamics. While not difficult to accommodate within the control bandwidth, the effect of lightly damped spurious modes outside the bandwidth may be devastating due to their potential to de-stabilize the control system. A long, thin metallic flexure has high-frequency sinusoidal beam bending modes, which if left undamped could couple into the control loop. Indeed, in the HBSM design, Loney reports that the lightly-damped secondary beam bending modes of the metallic flexures limited the bandwidth of the closed-loop system [15]. It proved necessary to add a layer of damping material to the flexures in order to increase system performance.

The use of elastomers in place of metallic flexures can overcome these limitations. In particular, they have the following advantages:

- **Orthotropic properties in shear and compression.** Elastomers are nearly incompressible, with a Poisson's ratio of nearly 0.5. Because of this, bearing designs which place a given volume of the elastomer material in a partially constrained state cause the effective compression modulus of the material to increase dramatically, while leaving the shear modulus at a constant value. Effectively, this causes the material to exhibit orthotropic characteristics, which can be exploited to obtain the desired compliance in the actuated DOF, while maintaining high stiffness in constrained DOFs. Thus, bearings designed using elastomers are far less geometry-dependent than their metal counterparts, because the actuated and constrained DOFs may be chosen independently[8].

- **Compact, low mass design.** As a result of the above property, it becomes possible to design flexures with stiffness properties identical to that of a metallic flexure, but at a fraction of the volume and mass. In addition, an elastomeric flexure is generally placed in a shear orientation which is conformal to the actuated DOF, while a metallic flexure is normally aligned with its long dimension perpendicular to the actuated DOF, as shown in Figure 2-8. As a result, the bearing design has a more compact form factor and less complicated support structure (due to its ability to be simply bonded to a surface rather than supported in, say, a cantilever configuration which would require a complex clamping arrangement).

- **High Internal Damping.** Unlike metal, elastomer compounds exhibit viscoelastic stress-strain behavior, which means they posses a combination of elas-

---

[8] This is not to say that elastomeric bearings have *no* geometry dependence. On the contrary, in order to obtain the desired compression behavior it is necessary to use a geometry with a high shape factor (see Equation 2.13). The assertion here is that elastomers allow greater flexibility in geometry than a metal flexure for a given stiffness ratio.

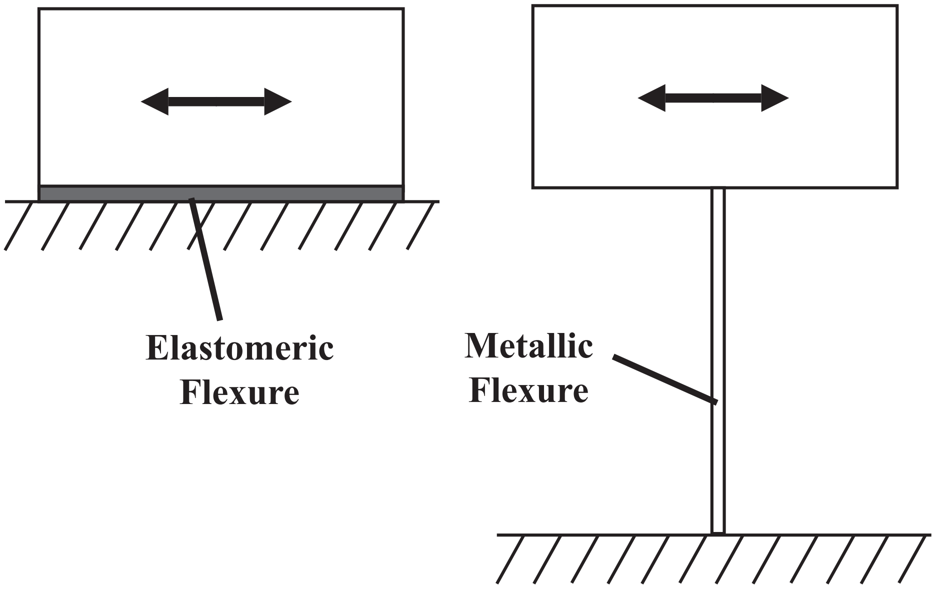

Figure 2-8: Comparison of elastomeric and metallic flexure form factors.

tic (stress proportional to strain) and viscous (stress proportional to strain rate) properties. The amount of damping in a given material is expressed in its loss factor η, which is defined as the ratio of the dissipated to stored energy in a stress-strain cycle. In a typical second-order system representation[9] with damping ratio ζ, $\eta = 2\zeta$.

Elastomeric materials typically have a frequency-dependent loss factor in the 0.1 to 1.0 range. As a consequence, when used as a flexure they provide damping to the suspension modes of the moving machine element, and exhibit small or no internal resonances.

Of course, the use of elastomer flexures involves certain compromises. The most important disadvantages are:

- **Environmental dependency.** This is arguably the most significant limitation of elastomeric materials, and one that requires careful attention during design. The material properties of elastomers exhibit wide variation as environmental

---

[9] The transfer function of a generalized second-order system with unity DC gain is $\frac{1}{(1/\omega_n^2)s^2+(2\zeta/\omega_n)s+1}$, where ω_n is the natural frequency in rad/sec and ζ is the damping ratio.

parameters change, most notably temperature. An example of this is the well-known ‘glass transition’ phenomenon whereby elastomers rapidly become hard and brittle when cooled to low temperatures. In addition, elastomers are susceptible to degradation in chemical and oxidizing environments, and many do not perform well under vacuum. Heavy radiation environments easily disintegrate them as well.

- **Variable material properties.** Because elastomers are easily modified chemically, many variations of a given compound exist. Each may contain different levels of fillers, plasticizers, or other additives which may result in a wide variation in properties. Also, many elastomers have material properties which are highly sensitive to the particular manufacturing processes used in production. Samples of the same material from different manufacturers may not perform identically.

AFSM Elastomeric Bearing Design

I chose to use an elastomer for the AFSM radial bearing design in order to realize the benefits of inherent damping, compact form factor, and orthotropic properties. Also, a key factor in this decision relied on the fact that variations in material properties with environment were not a concern. This is because the AFSM is designed to be a component in an optical system whose performance is also highly dependent on environment. Thus, such systems are tightly controlled in temperature, pressure, etc. in order to mitigate these effects. Consequently, it is reasonable to expect the AFSM environmental excursions to be relatively benign.

The choice of a thin strip of elastomer provided an elegantly compact flexure design solution, as I was able to place it in the nonworking airgap of the flux-steering actuators. Not only did this provide the proper mechanical constraints and compliances, but it also established a self-aligning, uniform, repeatable airgap between the actuator permanent magnet and the mirror armature. In fact, despite the critical functionality provided by the flexure, no additional space was needed to incorporate it into the AFSM design! Lu also took advantage of this in his actuator design [16].

As described by Lindley [14], the calculations for determining the static stiffness of simple shapes of elastomer under various loadings are relatively straightforward, provided the applied strains are moderate (roughly 10% for compression and 50% for shear). In fact, they are practically identical to the familiar equations for a pure elastic material. The compression stiffness k_c, shear stiffness k_s, and torsional stiffness k_{tor} are given by

$$k_c = \frac{AE_c}{t} \tag{2.10}$$

$$k_s = \frac{AG}{t} \tag{2.11}$$

$$k_{tor} = \frac{J_r G}{t} \tag{2.12}$$

where A is the pad loaded area, E_c is the *effective* compression modulus, t is the pad thickness, G is the shear modulus, and J_r is the pad polar moment of inertia.

The only complicating factor in the analysis is the determination of the effective compression modulus E_c. As discussed above, it is a function of the material elastic modulus, Poisson's ratio, and geometry. A mathematical derivation of the relationship of E_c to these properties is provided by Lindley; however, he also provides a more direct graphical method to find E_c based on the hardness properties of the elastomer as well as a geometric quantity called the *shape factor*. This dimensionless quantity is defined simply as the ratio of the loaded area of a given elastomer volume to the unloaded area. For the simple shapes of a thin rectangular block and a thin circular disk, the shape factors are defined, respectively, as

$$\begin{aligned} S_{rect} &= \frac{wl}{2t(w+l)} \\ S_{circ} &= \frac{\pi d^2}{4\pi t d} = \frac{d}{4t} \end{aligned} \tag{2.13}$$

where w and l are the rectangular block's width and length, t is the block thickness, and d is the circular block's diameter. Here, the compressive loading is applied along the thickness dimension t.

Once the shape factor is known, it may be used along with the elastomer hardness value to determine the effective compressive modulus using the chart in Figure 2-9. Note that in the AFSM design I used Shore A hardness, which is equal to the IRHD units used in the chart.

For the AFSM elastomer bearing design, I chose a rectangular pad geometry which fit nicely in the magnetic actuator nonworking airgap. Since the airgap length (which is equal to the bearing thickness) has only a small effect on the magnetic performance (see Chapter 3), I fixed the rubber length and width at values which matched the armature and magnet pole areas, and varied the elastomer hardness and thickness to achieve the desired stiffness properties. Again, this approach is similar to Lu's.

Since the radial forces on the armature due to the permanent magnet flux are about the same as the expected normal forces, I used the same criteria for the radial stiffness as I did for the axial stiffness; i.e. less than 1 micron deflection for a 10% force mismatch. Since two pads are in compression and two in shear for radial loading, the resultant stiffness is

$$k_{rad} = 2k_c + 2k_s \tag{2.14}$$

I derived the shear and torsional stiffness requirements for the bearings based on the desired dynamic suspension natural frequencies of the mirror in the actuated degrees of freedom. Since my goal was 5 to 10 kHz bandwidth, I wanted to place the natural frequency in the 100-1000 Hz range in order to meet the dynamic goals stated

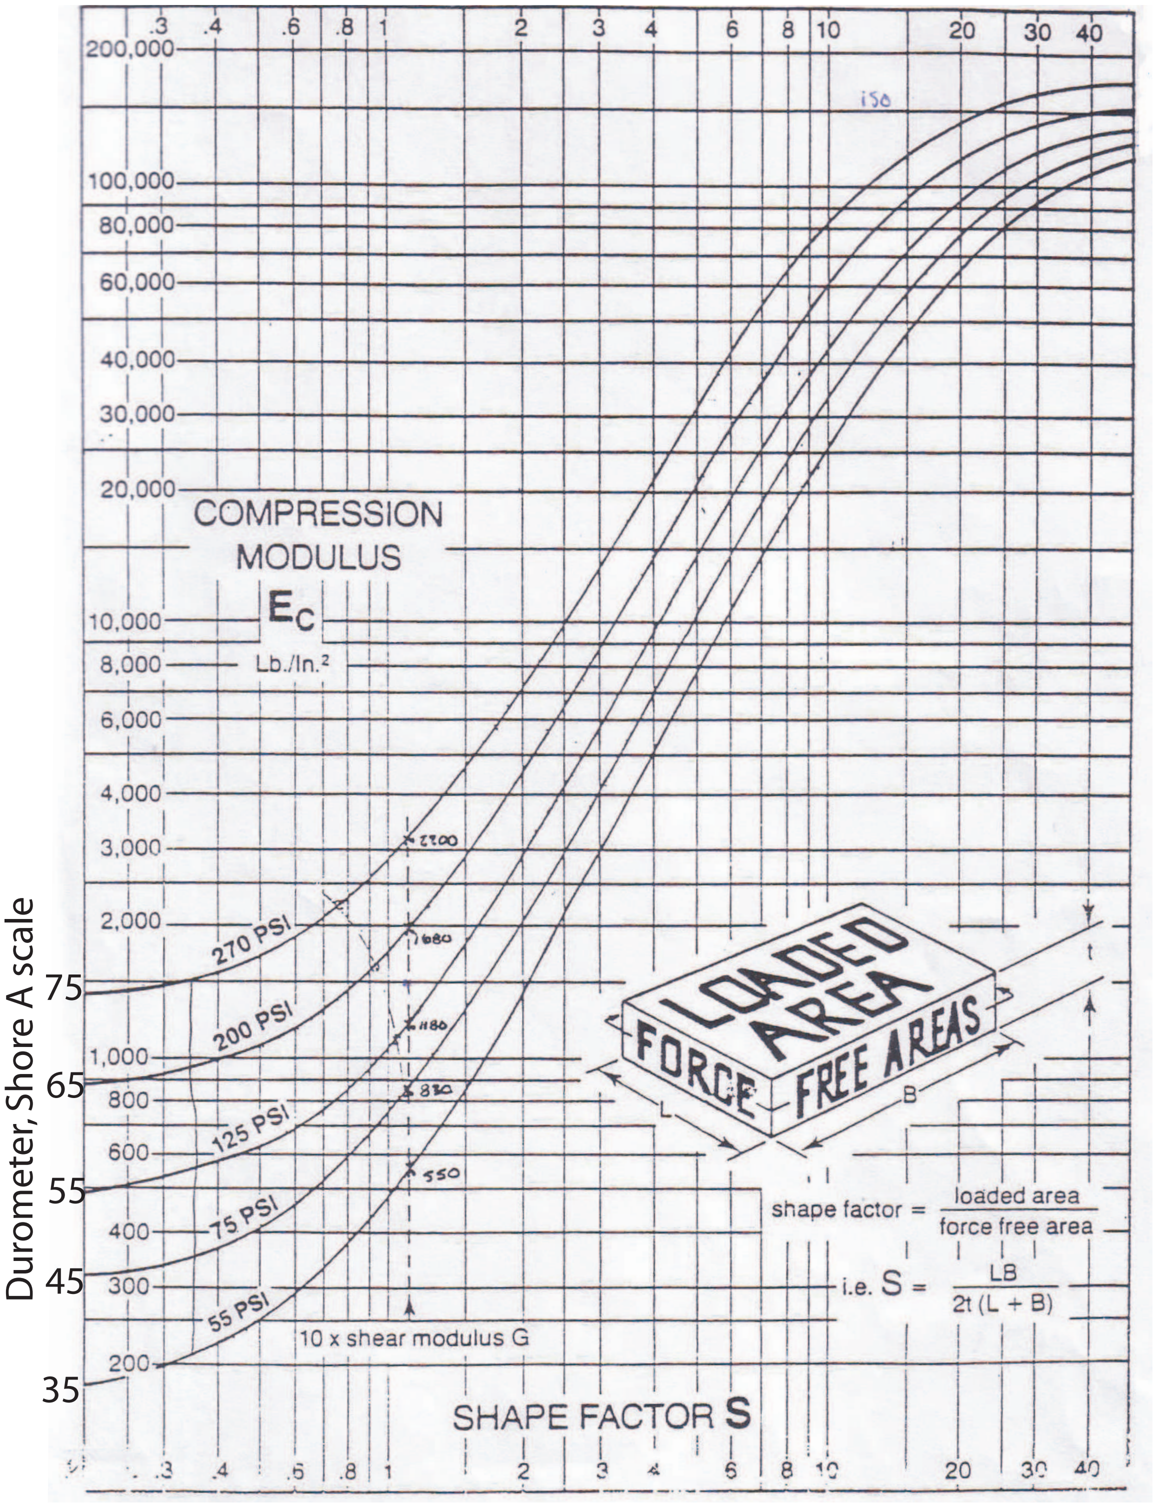

Figure 2-9: Chart used to determine effective compression modulus E_c given elastomer hardness and shape factor. Reprinted from Lindley [14].

at the beginning of this chapter. At the time I designed the bearings, the rotational inertia of the mirror and armatures was already set at 4.091×10^{-6} kg-m^2, which allowed for a fairly large stiffness range of 1.6 to 161 N-m/rad.

There are many different contributors to total rotational stiffness of the assembly along one of the two actuated DOFs. First, there are two shear components: A positive stiffness from the elastomer bearing, and a negative stiffness k_{mag} from the actuator magnetics (see Chapter 3). These produce a linear force at the end of the armature, so we must derive the resultant rotational stiffness accounting for the linear action through the moment arm r. Using the approximation $\theta \approx x/r$, the equivalent rotational stiffness k_{eq} is

$$k_{eq} = \frac{M}{\theta} = \frac{Fr}{x/r} = \frac{Fr^2}{x} = k_s r^2 \tag{2.15}$$

where M is the applied moment, θ is the rotational angle, F is the linear force, and x is the linear displacement.

The two remaining contributors to the stiffness are from the torsional loading of the off-axis elastomer bearings k_{tor}, and from the axial flexure bending stiffness k_b. Along with the shear stiffnesses, these sum to give the total rotational stiffness as

$$k_{rot} = 2(k_s - k_{mag})r^2 + 2k_{tor} + k_b \tag{2.16}$$

Note that the negative stiffness provided by the actuator magnetics allows for a certain degree of flexibility in the actuator design. Normally, the difficulty encountered in designing flexures is associated with establishing low stiffness in the actuated degree of freedom, while simultaneously providing high kinematic stiffness in the non-actuated DOFs. With the negative magnetic stiffness in place, the shear stiffness of the elastomer flexures can be designed higher in the knowledge that the overall net stiffness will be small.

Iterating through the stiffness calculations for various elastomer hardness material thickness values, I arrived at the bearing design given in Table 2.3, below.

Table 2.3: Elastomer Flexure Design Values

Property	**Value**
Length l	20 mm
Width w	4 mm
Thickness t	1 mm
Hardness	70 Shore A
Effective Compression Modulus E_c	30 MPa
Shear Modulus G	1.7 MPa
Shape Factor S_{rect}	1.66
Shear Strain at Full Deflection	20%
Compressive Stiffness k_c	2.4 N/μm
Shear Stiffness k_s	0.136 N/μm
Torsional Stiffness k_{tor}	4.7 N-m/rad

Combined Mirror Stiffness

The last step in solving Equation 2.16 is to determine the negative stiffness of the magnetic actuator. I did this numerically using the relations derived in Chapter 3. The final values for the axial and elastomer flexures and the actuator are given below, along with the total rotational stiffness, inertia, and expected natural frequencies.

Table 2.4: Consolidated Mirror Stiffness and Natural Frequency

Property	**Value**
Total Flexure Stiffness	120.2 N-m/rad
Total Actuator Magnetic Stiffness	-99.7 N-m/rad
Net Mirror Rotational Stiffness	20.5 N-m/rad
Total Mirror Axial Stiffness	14.4 N/μm
Mirror and Armature Inertia	4.091 x $10^{-}6$ kg-m^2
Mirror and Armature Mass	36.3 g
Calculated Rotational Mode Frequency	356.3 Hz
Calculated Axial Mode Frequency	3,174 Hz

2.5 Modal Analysis of the Dynamic Structure

With the dynamic components of the AFSM designed, the final mechanical task is to perform a modal analysis of the moving structure, which includes the magnetic armatures, mirror, and portions of the flexures.

The discussion in Section 2.2 shows that the modal frequencies are important parameters in the controller design, but does not consider the effect of the mode shapes. The shape of a particular vibrational structural mode relative to the placement of the sensors and actuators that measure and drive the system directly determines the degree to which the mode affects the system response [28]. For example, if the feedback sensors measuring the mirror displacement happen to be located along the nodal line of a mode shape (i.e. a line of zero net modal displacement), they will return zero output. Thus, even though the mode is excited, its displacements do not enter the feedback loop, and thus cannot destabilize the control system. In a similar manner, we can show that actuators located at such points cannot excite the mode in question. Hence, we avoid destabilization in this case as well. Clearly, foreknowledge of the structural mode shapes will allow us to advantageously place the sensors and actuators in order to gain additional system robustness, and may even allow for an increase in the controller gain to improve the system's disturbance rejection performance. Weng and Lu both provide good theoretical treatments of this type of strategy in their PhD theses [28] [16], while Rankers provides the same conclusions from an engineering point of view [24].

Of course, our success or failure at achieving the ideal frequency response and sensor and actuator placement depends on the ability to accurately predict the required mode frequencies and shapes. In particular, we would like to be able to determine the flexure-supported suspension mode and the first few structural mode frequencies very precisely. Hence we require a mathematical model that emulates the behavior of the physical structure with high fidelity. The construction of such a model is presented in this section.

2.5.1 Modal Analysis Using the Finite Element Method

I used finite element analysis to develop a valid model of the FSM structure. The goal is to select a finite element modeling strategy (i.e. element type, mesh size and structure, and physical component representation) that results in accurate modal frequencies and shapes without undue computational effort (in terms of solution time and memory usage).

I originally performed most of this analysis as part of a class project, and much of the following text is excerpted from my project report[13]. It contains extensive derivations from continuum theory and validation of the finite element model used for the analysis. For these details, I refer the interested reader to that report. Here, I wish to cover the important aspects and results of the AFSM model itself.

I considered three types of finite element formulations for use in the analysis: three-node plate elements, four-node shell elements, and eight-node three-dimensional elements. I compared these to a baseline model composed of a fine mesh of 27-node three-dimensional elements representing the "exact" solution, and also to two different analytical models derived from plate theory. The results, given in Figures 2-10 through 2-13 for two mode shapes show that for structures such as the AFSM mirror having a thickness ratio[10] of 0.33, an eight-node three-dimensional element was the best compromise in terms of accuracy and efficiency (the solution times for the eight-node element were faster than the 27-node element by a factor of 100 or more).

Finite Element Analysis of the AFSM Structure

With the element type chosen, a mesh must now be constructed that faithfully represents the given geometry, materials, and boundary conditions. The construction of the model is complicated somewhat by the problem of modeling the springs (i.e. the elastomer bearings and axial flexure) that suspend the structure in a fully three dimensional environment. In order to properly apply the spring forces to the model

---

[10] For plate structures, the thickness ratio α is defined as the ratio of the plate thickness to its least plan dimension; i.e. a 20x10x1 mm plate has $\alpha = 0.1$.

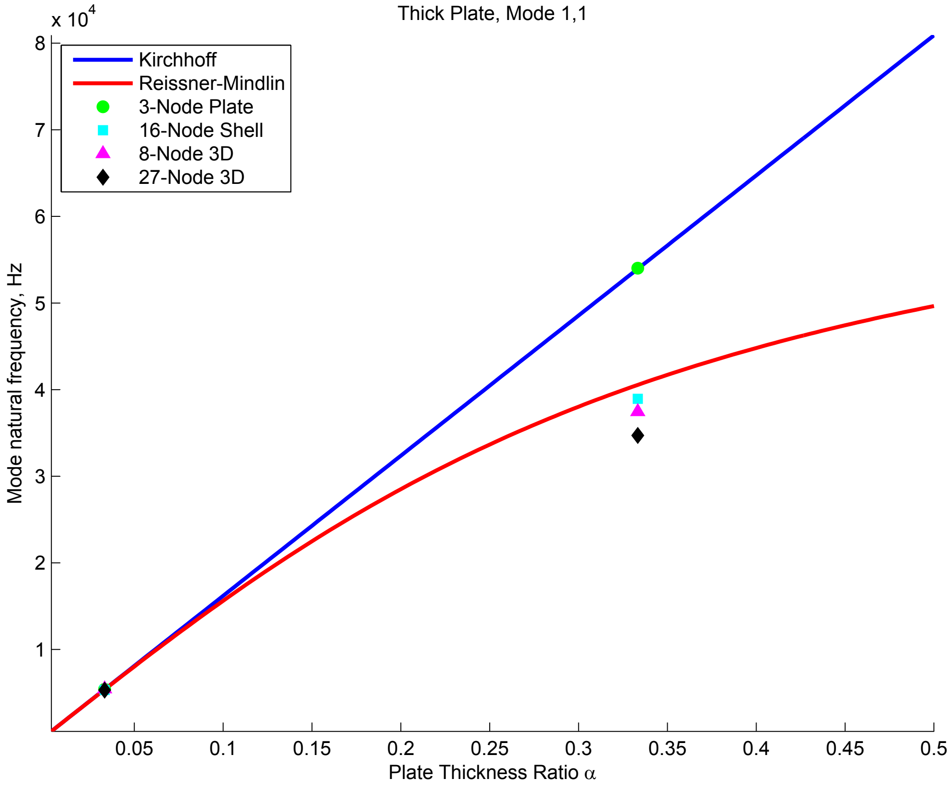

Figure 2-10: Mode (1,1) Analytical solutions and finite element results for thin ($\alpha = 0.03$) and thick ($\alpha = 0.33$) cases. The black diamond is considered to be the "exact" solution.

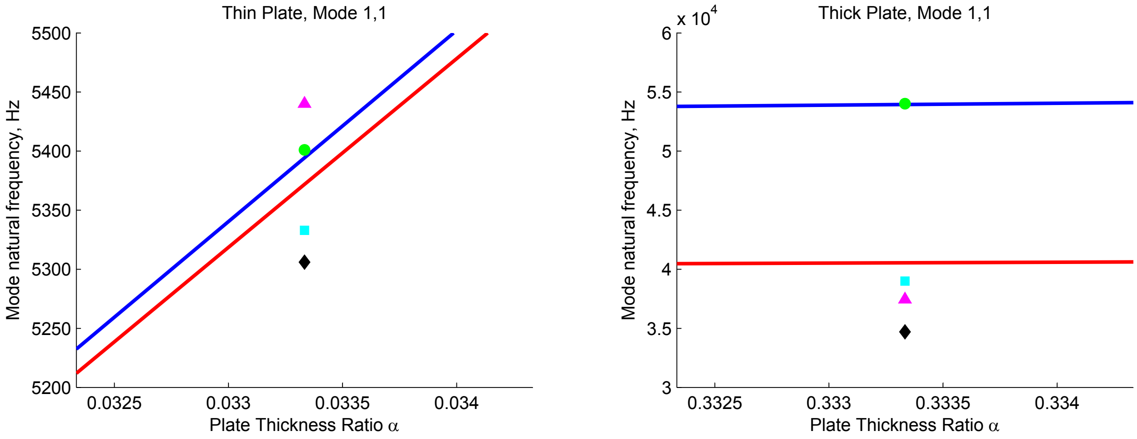

Figure 2-11: Expanded view of the results in Figure 2-10: Thin (left); thick (right)

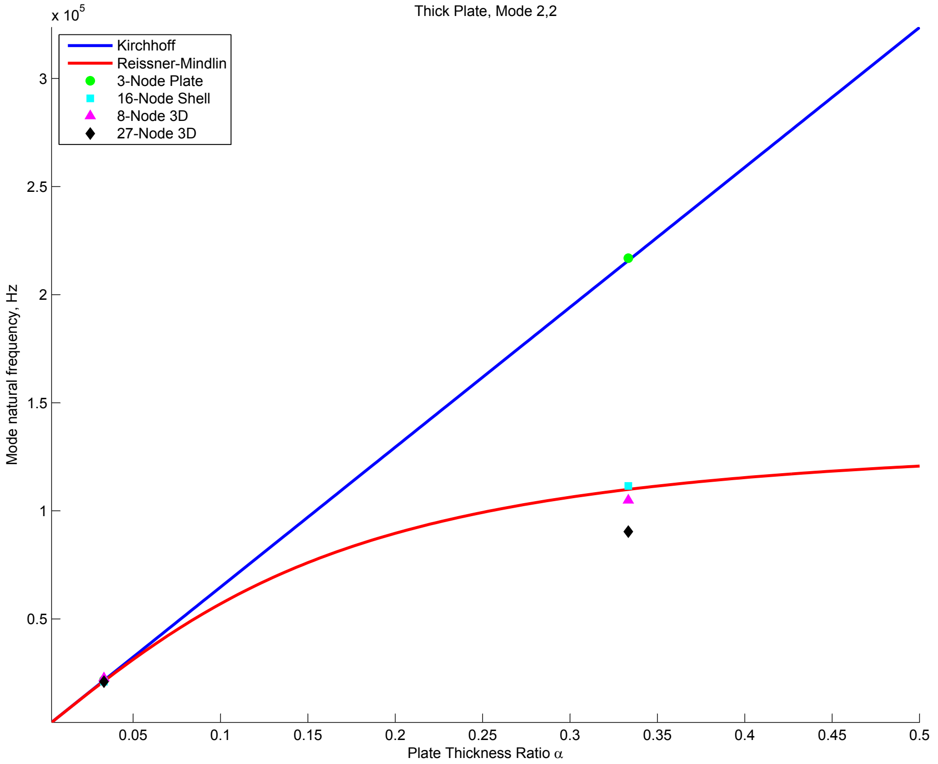

Figure 2-12: Mode (2,2) Analytical solutions and finite element results for thin ($\alpha = 0.03$) and thick ($\alpha = 0.33$) cases. The black diamond is considered to be the "exact" solution.

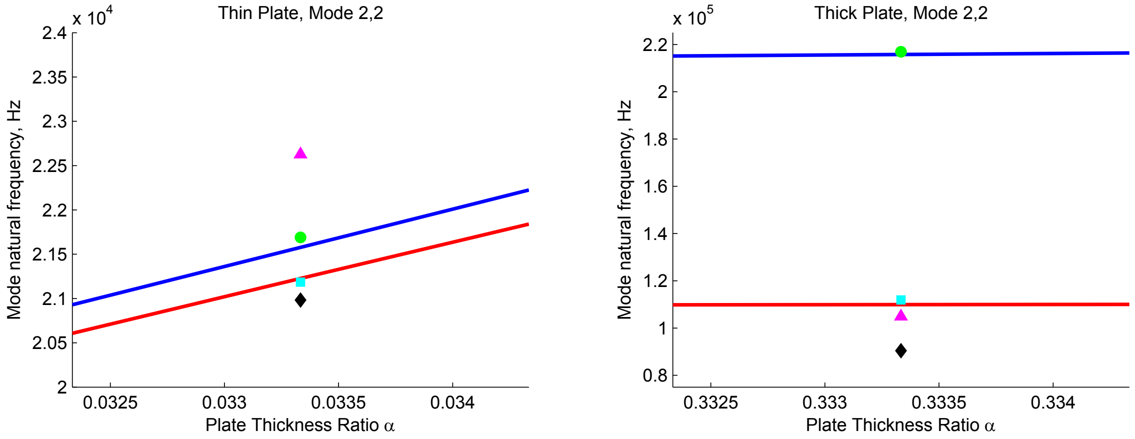

Figure 2-13: Expanded view of the results in Figure 2-12: Thin (left); thick (right)

in all three dimensions, the contribution from each spring must be broken up and distributed among the interacting structure nodes. This is particularly problematic when modeling the torsional spring constants of these components, because the three-dimensional structural model has no rotational degrees of freedom. Thus, a lumped torsional spring element would have no effect on the results, so torsional effects must be approximated via a set of equivalent translational spring elements.

We can therefore proceed in one of three ways. The first is to attempt to model the springs through individual translational elements, as described above. However, this must be done manually for each affected node, and is thus cumbersome and time consuming. Furthermore, any refinements to the structure mesh (say, for a convergence analysis) will require reallocating the spring elements accordingly.

The second recourse is to fully model the rubber bearings and axial flexure using three-dimensional finite elements, and incorporate these models into the overall modal analysis. In my original project I did not choose this option due to time constraints and instead proceeded with the third option described next. However, I returned to a full model later to verify the results of my original approach were valid.

As a compromise, a third option is simply to rely on the differences in stiffness between the structure and its supports, and consider the structure to be unsupported (i.e. to have completely free boundary conditions) within the frequency domain of the structural modes. Referencing Figure 9.3 on page 794 of Bathe [3] showing the dynamic load factor D versus normalized frequency, we surmise that such an assumption is valid provided that

$$\frac{\hat{\omega}}{\omega} < \frac{1}{4} \tag{2.17}$$

where in the present case $\hat{\omega}$ is the highest spring-supported mode frequency and ω is the lowest structural mode frequency. In this case, the supports behave as though they were static, and have no dynamic effect on the higher modes. Furthermore, we can deduce that the structural mode frequencies calculated under free boundary conditions are in fact the lowest possible frequencies for that structure. This is true because any support or constraint added to the free model will effectively increase the stiffness of the model, driving the structural mode frequencies higher.

Therefore, if we can calculate the highest spring-supported mode and lowest structural mode frequency and show that Equation 2.17 holds, we can then assume that a modal analysis with free boundary conditions will give mode shapes and frequencies that are very close to the spring-supported case, and will in fact be conservative (i.e., the free boundary analysis mode frequencies will be slightly lower than those of the actual structure).

Fortunately, determining the spring supported modes is easy to perform with a simple hand calculation using a lumped-parameter approximation. For translational and rotational modes, we use

$$\begin{aligned} \hat{\omega}_t &= \frac{1}{2\pi}\sqrt{\frac{k_t}{m}} \\ \hat{\omega}_r &= \frac{1}{2\pi}\sqrt{\frac{k_r}{I}} \end{aligned} \tag{2.18}$$

respectively, where k is the lumped spring rate, m is the lumped mass, I is the lumped inertia, and the subscripts t and r denote translation and rotation, respectively. These values are conveniently found in Table 2.4. The values of $\hat{\omega}_t$ and $\hat{\omega}_r$ are 3,174 and 356.3 Hz, respectively.

The AFSM mirror structural mesh (without supporting springs) for the modal analysis is constructed as shown in Figure 2-14. All elements are of the eight-node three dimensional formulation with incompatible modes[11]; however, because the mirror plate and armatures are made of different materials as described in Chapter 1, two different element groups are used, with the proper material model assigned to each (SM-2 for the armatures, aluminum for the mirror). Because there are no supports,

---

[11] For a discussion of incompatible modes in the element formulation, see [3].

it is necessary to change the eigensolution algorithm options to accept rigid body modes. Accordingly, the first meaningful structural mode calculated is the seventh mode in the solution output, after the six zero-frequency rigid body solutions.

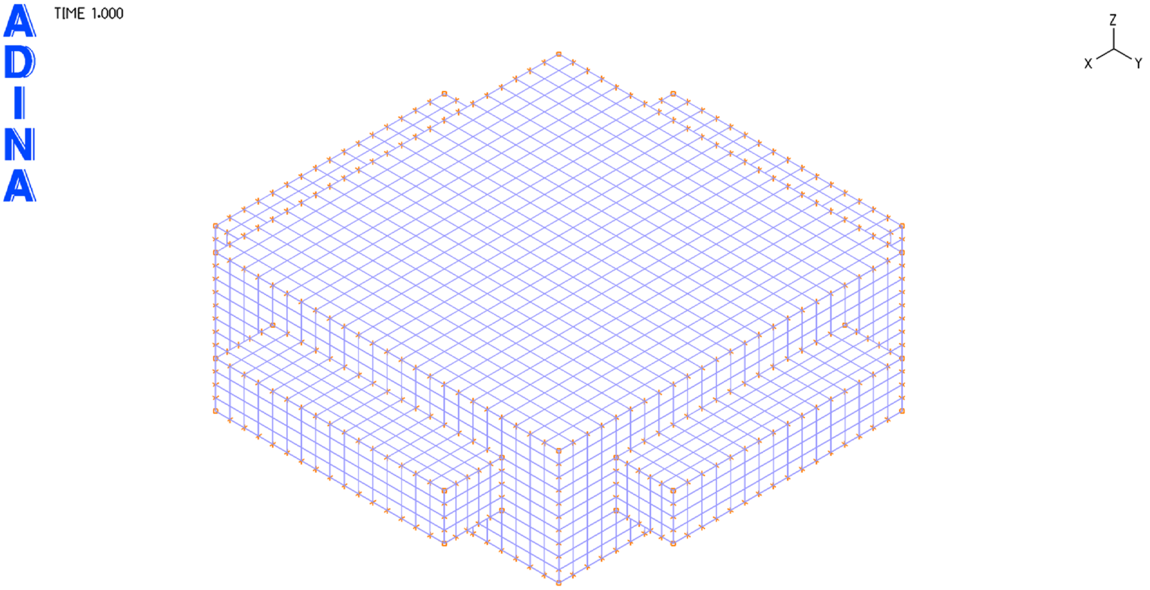

Figure 2-14: AFSM mirror structural mesh with free boundary conditions.

Modal Analyses and Convergence Study

We will now perform a modal analysis of the FSM structure with free boundary conditions to determine the lowest value of ω. If Equation 2.17 holds, we will consider the analysis results to be valid. Beginning with a coarse mesh, I conducted a convergence study on the model of Figure 2-14 to obtain an accurate modal solution. The results of this study are shown in Table 2.5 for the first and sixth structural modes. Here, I defined an acceptable convergence value as less that 2% change between solutions of varying mesh density.

As shown in the table, the final mesh with 21,200 elements converged to an acceptable solution. Here the decision to use the 8-node three-dimensional element (as opposed to the higher-order 27-node formulation) was well-rewarded, as the final solution took more than 20 minutes to complete, and required the use of out-of-core memory.

Table 2.5: FSM Structural Mesh Convergence Results

Mesh Density	**Mode 1 Frequency**	**% Change**	**Mode 6 Frequency**	**% Change**
768 elements	22,327 Hz	N/A	54,937 Hz	N/A
7,040 elements	21,997 Hz	-1.5%	51,981 Hz	-5.7%
21,200 elements	21,912 Hz	-0.4%	51,349 Hz	-1.2%

Analysis Results and Conclusions

The frequency results from the final analysis are given in Table 2.6 for the first six structural modes, and the corresponding mode shapes are shown in Figures 2-15 through 2-19. Note that the fourth and fifth modes have repeated eigenvalues, so the mode shape is symmetric.

We are now in a position to determine whether Equation 2.17 holds. The lowest structural mode under free boundary conditions is $\omega = 21,912$ Hz, and the highest spring-supported mode is the translational mode $\hat{\omega}_t = 3,510$ Hz. We therefore easily find that

$$\frac{\hat{\omega}_t}{\omega} = \frac{3,510}{21,912} = 0.16 < \frac{1}{4} \tag{2.19}$$

so the constraint on dynamic load factor is satisfied, and the analysis results are valid. Also, since the target bandwidth goal is 5 kHz, the first mode frequency of 22 kHz indicates that the structure has adequate stiffness to prevent structural mode complications in the control design.

Again keep in mind that these modal frequency results would be drastically improved by changing to a beryllium design. Beryllium has a stiffness-to-mass ratio of about 6.4 times that of aluminum, which would effectively increase the modal frequencies by more than a factor of two.

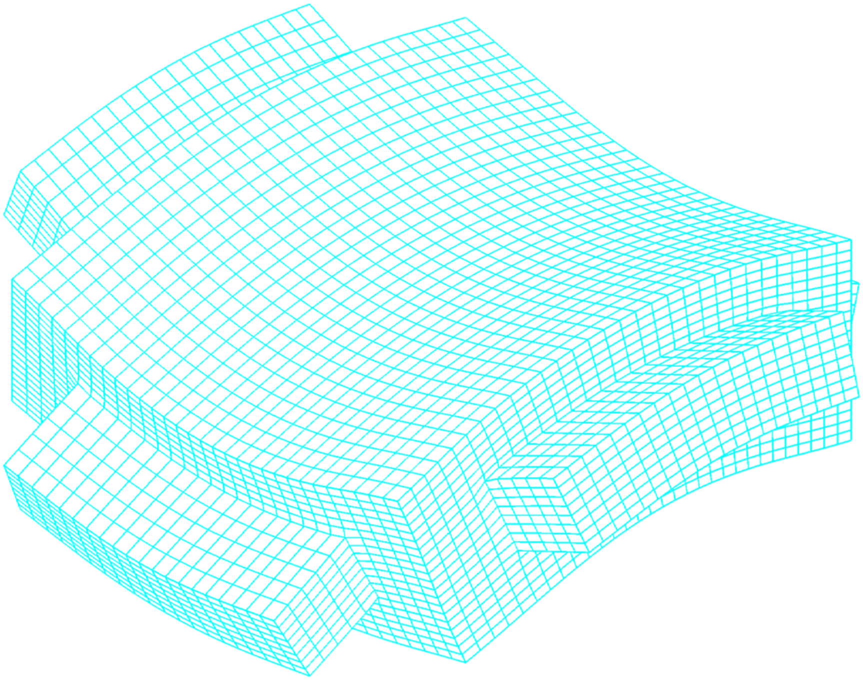

Figure 2-15: FSM first structural mode shape, free boundary conditions

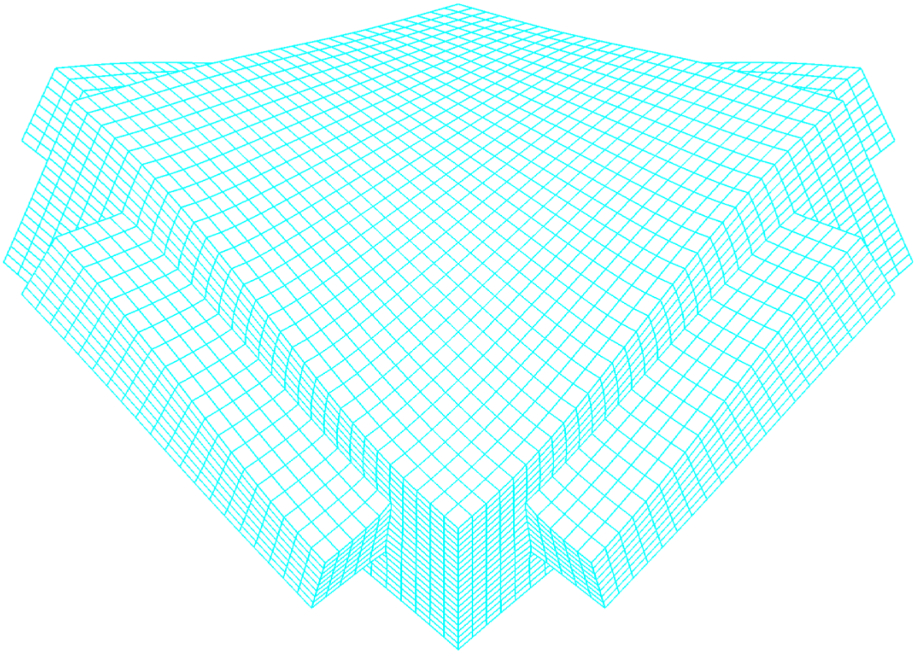

Figure 2-16: FSM second structural mode shape

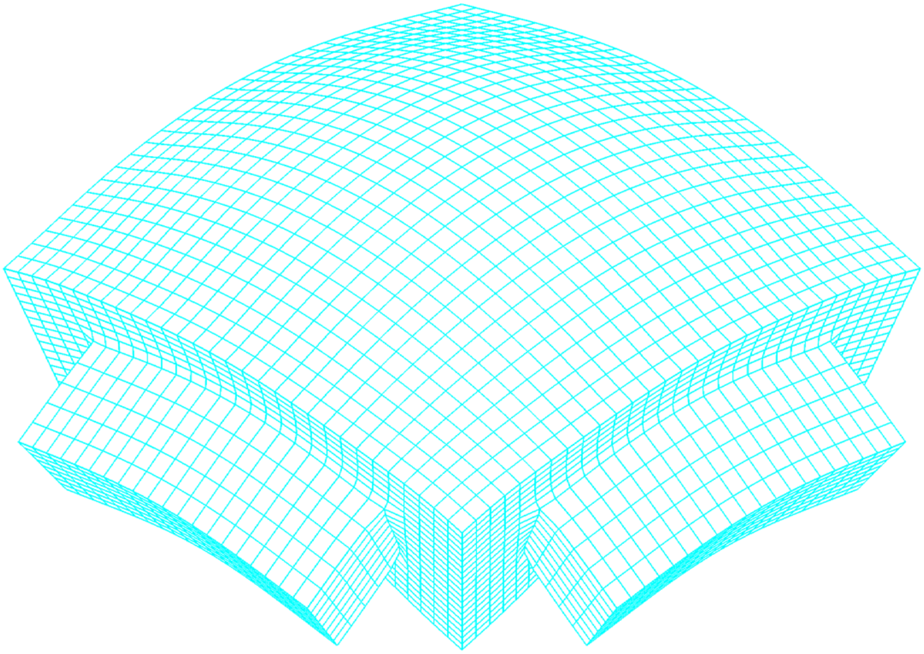

Figure 2-17: FSM third structural mode shape

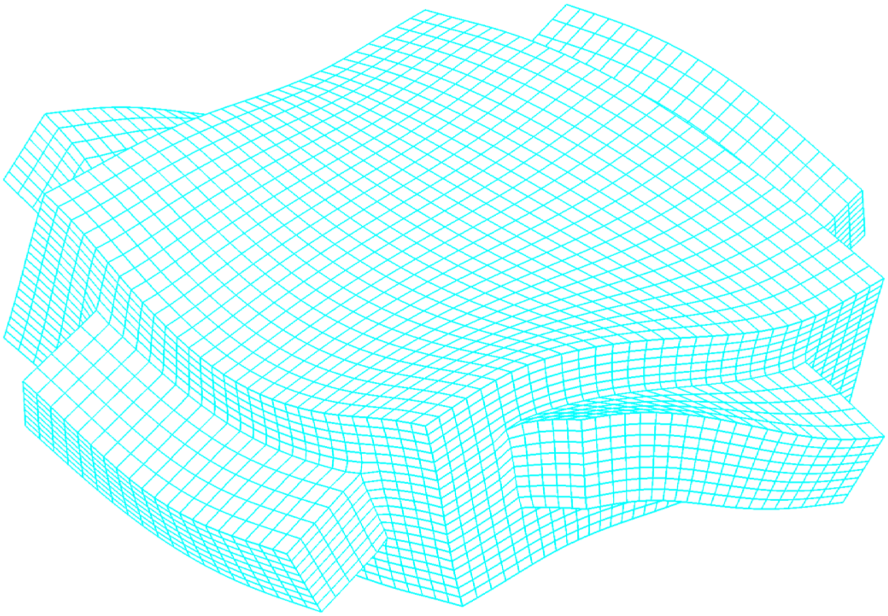

Figure 2-18: FSM fourth and fifth structural mode shapes (the fifth mode is symmetric with the fourth mode depicted here)

Table 2.6: FSM Modal Natural Frequencies, Final Mesh

Structural Mode	**Mode Frequency**
1	21,912 Hz
2	22,431 Hz
3	38,602 Hz
4	39,669 Hz
5	39,669 Hz
6	51,349 Hz

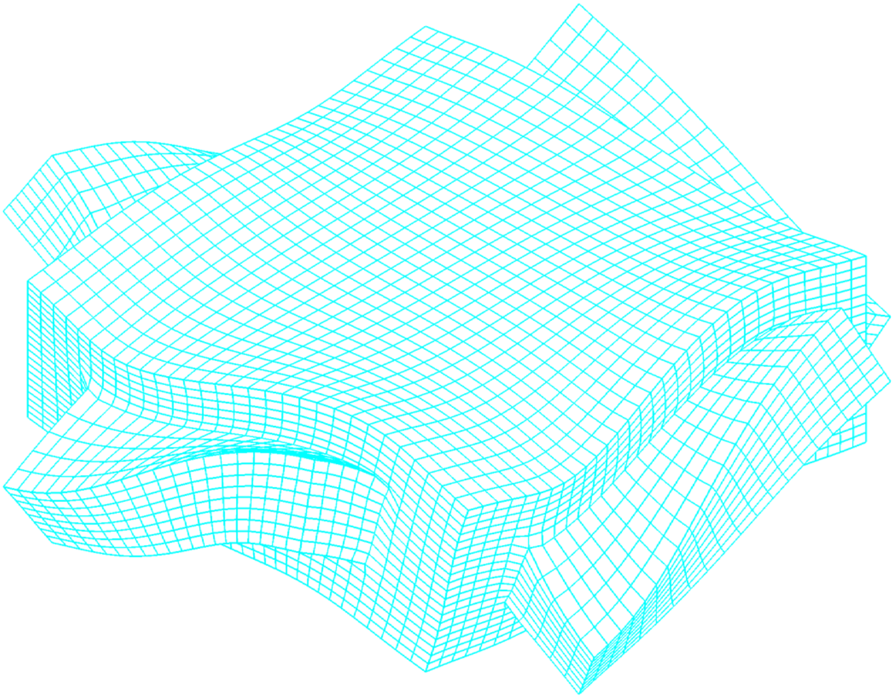

Figure 2-19: FSM sixth structural mode shape

Chapter 3

Electromagnetic Actuator Design

With the mechanical basis from Chapter 2 in mind, I now turn to the design of the electromagnetic system that supplies force to the moving assembly. The flux-steering actuator used in the AFSM design was developed in its present form by doctoral student Xiaodong Lu of the MIT Precision Motion Control laboratory in the early 2000s [16]. However, the basic electromagnetic layout of the actuator is not new, having been employed as a torque motor for hydraulic control since the late 1940s [18], and as a motor for engraving heads since the 1930's [22]. Figure 3-1 shows an example of the former: a jet-pipe servovalve. Note the biasing magnets and drive coils which provide AC flux against the bias flux.

Lu's (and concurrently, Montesanti's) contribution involved designing and implementing the actuator as a direct drive force source in a closed-loop, single DOF precision machine system [16], [22]. I employ it again for this purpose in the AFSM, but here I use four independent actuators in a two-DOF configuration. In this chapter, I present some fundamental electromagnetic concepts, and then discuss the design of the AFSM actuator based on the developed principles.

3.1 Fundamentals of Electromagnetism

The basic physics underlying the AFSM actuator is well documented, and in fact, the specific electromagnetic analysis for the AFSM flux steering actuator is virtually

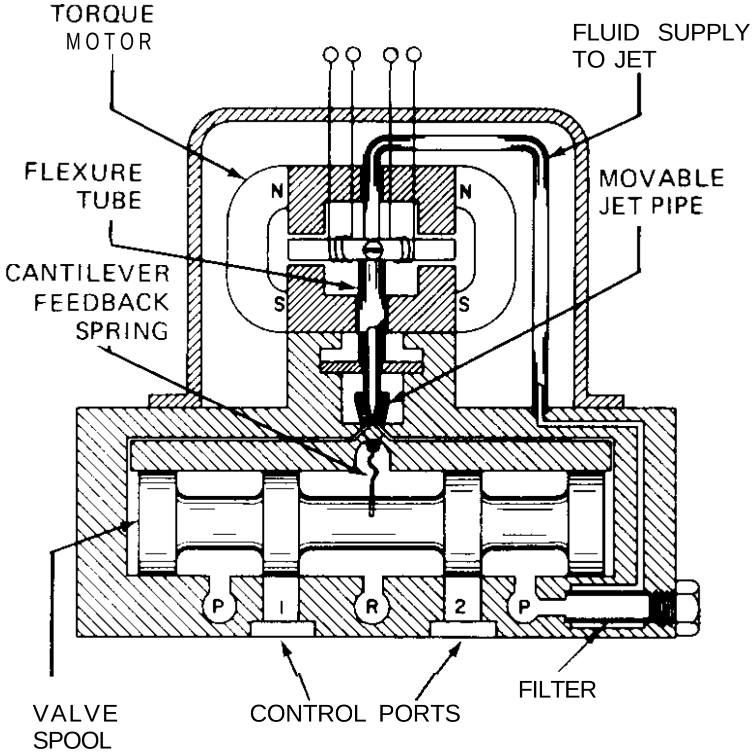

Figure 3-1: A jet-pipe servovalve cross-section showing the electromagnetic torque motor.

identical to Lu's. As such, my goal here is not to present the AFSM electromagnetic design as a novel innovation, but rather to use it to demonstrate the utility of magnetic circuit analysis for general machine design purposes. My hope is to demystify the treatment of magnetic analysis for these common cases, such that it becomes more accessible to the general engineering community–and especially to those wishing to build practical machinery. To do this, I first discuss current methods of teaching electric circuit theory, and then argue that analogues of these methods in the magnetic domain yield the most straightforward results.

3.1.1 Electrical Circuit Analysis

In the present state of engineering education, introductory physics courses present concepts such as electricity, mechanical dynamics, thermodynamics, etc. via their fundamental physical laws. For example, mechanics is introduced through a rigorous examination of Newton's second law, $\mathbf{F} = \frac{d(m\mathbf{v})}{dt}$, while electricity and magnetism begin with the laws of Maxwell and Lorentz. In later, specialized courses, the fundamental laws are modified through the use of simplifications and approximations in order to make the analysis tractable and useful for specific applications.

As an example, consider a vibrations course, where Newton's law (and others) are applied to the specific case of a lumped-parameter mechanical model. By making appropriate simplifications, a very useful engineering tool–in this case, modal decomposition–is derived, which can be directly applied to real problems. Provided the engineer remains cognizant of the underlying assumptions, the approach yields an appropriate balance between practical utility and analytical rigor.

This methodology also holds true for electric circuit analysis. In introductory physics texts, the fundamental laws governing electricity, magnetism, and electromagnetic waves are given by the famous Maxwell's equations, listed below in integral form[1]. The four equations respectively are: Ampere's law, Faraday's law, Gauss's law for electric fields, and Gauss's law for magnetic fields [10].

$$\oint_C \mathbf{H} \cdot d\mathbf{l} = \int_S \mathbf{J} \cdot \mathbf{n}\, da + \int_S \frac{\partial \mathbf{D}}{\partial t} \cdot \mathbf{n}\, da \tag{3.1}$$

$$\oint_C \mathbf{E} \cdot d\mathbf{l} = -\frac{d}{dt} \int_S \mathbf{B} \cdot \mathbf{n}\, da \tag{3.2}$$

$$\oint_S \mathbf{D} \cdot \mathbf{n}\, da = \int_V \rho dv \tag{3.3}$$

$$\oint_S \mathbf{B} \cdot \mathbf{n}\, da = 0 \tag{3.4}$$

In addition, the Lorentz force law is given as

$$\mathbf{F} = q\left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right) \tag{3.5}$$

and the Maxwell stress tensor is

---

[1] An equivalent representation of Maxwell's laws is the derivative form, which uses the gradient (∇) operator. I prefer the integral form, as I find it easier to visualize the contours and surfaces of integration.

$$\sigma_{ij} = \epsilon_o E_i E_j + \frac{B_i B_j}{\mu_o} - \frac{1}{2}\left(\epsilon_o E^2 + \frac{B^2}{\mu_o}\right)\delta_{ij} \tag{3.6}$$

The above equations use the following definitions and conventions, in SI units. Note that variables in bold type denote vector quantities, and that in Equation 3.6 the vector fields are broken into their scalar tensor components, denoted by the subscripts i and j.

- $\mathbf{E}$ is the electric field, in volts per meter.
- $\mathbf{D}$ is the electric flux density, in coulombs per square meter.
- $\mathbf{H}$ is the magnetic field, in amperes per meter.
- $\mathbf{B}$ is the magnetic flux density, in tesla (equivalently, webers per square meter).
- $\mathbf{J}$ is the current density, in amperes per square meter.
- $\mathbf{F}$ is the mechanical force, in newtons.
- $\mathbf{v}$ is the charge velocity, in meters per second.
- q is the electric charge, in coulombs.
- C is a line contour enclosing surface S, with differential element $d\mathbf{l}$.
- S is a surface with differential element $\mathbf{n}\,da$ ($\mathbf{n}$ denotes the unit vector normal to the surface).
- V is a volume enclosed by surface S, with differential element dv.
- ϵ_o is the permittivity of a vacuum, in coulombs per volt-meter.
- μ_o is the permeability of a vacuum, in newtons per ampere squared.
- δ_{ij} is the Kronecker delta.

From these fundamental laws, most electrical circuits courses immediately begin applying simplifications to establish the practical groundwork for future study. However, in this situation the assumptions made are applicable to such a wide range of real-world cases that much practical engineering work is possible without a more rigorous consideration of the fundamental laws. In fact, students may be entirely unaware that such simplifications have occurred–a likely event, since many textbooks do not bother to provide a derivation (for instance, see [12]).

As an example, consider one of the basic tools employed in electrical circuit analysis: Kirchhoff's Current Law (KCL), which states very simply that "the algebraic sum of the currents entering any circuit node is zero" [12]. This law applies to the specific case of electrical circuit junctions (nodes) composed solely of discrete ideal conductors, which carry uniform current densities and have no capacitance (i.e., no ability to accumulate charge). We can derive it from two of Maxwell's equations by applying these assumptions. To do so, first consider Ampere's law (Equation 3.1) for the special case of a closed surface (this is equivalent to placing a "control volume" around the circuit node). In this case, the left-hand side goes to zero (since the perimeter contour no longer exists) and we have

$$\oint_S \mathbf{J} \cdot \mathbf{n}\, da + \oint_S \frac{\partial \mathbf{D}}{\partial t} \cdot \mathbf{n}\, da = 0 \tag{3.7}$$

Next, take the time derivative of Gauss's law for electric fields (Equation 3.3) to obtain

$$\oint_S \frac{\partial \mathbf{D}}{\partial t} \cdot \mathbf{n}\, da = \int_V \frac{\partial \rho}{\partial t} dv \tag{3.8}$$

Combining Equations 3.7 and 3.8, we obtain the relation

$$\oint_S \mathbf{J} \cdot \mathbf{n}\, da + \int_V \frac{\partial \rho}{\partial t} dv = 0 \tag{3.9}$$

which is the equation for conservation of charge for an enclosed volume. With the further assumption of no accumulating charge within the volume (i.e., $\frac{\partial \rho}{\partial t} = 0$), we have

$$\oint_S \mathbf{J} \cdot \mathbf{n}\, da = 0 \tag{3.10}$$

Finally, if we consider the special case of n discrete conductors carrying currents $i_k = J_k a_k$, and substitute into Equation 3.10, we arrive at a mathematical expression for KCL:

$$\sum_{k=1}^{n} i_k = 0 \tag{3.11}$$

KCL is so widely applicable that it may be used, along with Kirchhoff's Voltage Law (KVL) and other similarly derived relations, to perform a lifetime's worth of circuit designs without considering underlying foundations at all. The caveat, of course, is that the designer remain aware of the assumptions under which the relations hold. For instance, attempting to apply KCL in the form of Equation 3.11 to a capacitor produces meaningless results[2]. However, if Equation 3.9 is used, the correct solution is obtained.

3.1.2 Magnetic Circuit Analysis

Comparing Equations 3.1 through 3.4 to 3.11, it is easy to see why KCL and its counterparts like KVL are so much more commonly used in engineering practice than the Maxwell relations. In short, they are lumped-parameter representations that are easily understood and simple to use. Full field solutions obtained through the general application of Maxwell's equations, on the other hand, are complex, time-consuming, and error-prone.

Given this, one would expect that designers of magnetic devices would follow a similar educational path to their counterparts in the electrical discipline. Strangely enough, this is not the case. While electrical circuit analysis is a rigorously developed subject in any engineering program, complimentary subjects in magnetic circuit analysis are far less common. Despite the presence of hundreds of types of mag-

---

[2] And in fact, it was this very dilemma that prompted Maxwell to add a correcting term involving the displacement current, $\int_S \frac{\partial \mathbf{D}}{\partial t} \cdot \mathbf{n}\, da$, to Ampere's Law when he published his work in 1864.

netic machines in use today, very few engineering programs cover the practical design fundamentals for these machines. Instead, many curricula give preference to highly theoretical full-field treatments of complex wave phenomena, such as that used for radar or magnetic resonance imaging–subjects which are certainly interesting and beneficial, but which have narrower practical applicability.

As such, it is generally up to the student or practicing engineer who wishes to design a practical machine to search the available literature, talk to other experienced engineers, and arrive at a preferred design approach. Going through this process myself for the AFSM, I came to the conclusion that magnetic circuit analysis, using similar simplifications to those used in electric circuit analysis, is the best choice for the design of practical magnetic machines. By 'best choice' I mean the method that produces a mathematical model which is sufficiently accurate for design in a minimum amount of time. Although several textbooks do cover magnetic circuit analysis with varying emphasis (see, for example, [27] or [19]), I feel the concepts are obscure enough to warrant a detailed presentation here. To do this, I will establish and draw comparisons between electric and magnetic analyses, state the assumptions involved, and then proceed with the AFSM actuator design. As a final step, I will show that the magnetic circuit solution produces identical results to a solution derived directly from Maxwell's equations. This last discussion I reserve until Appendix A for purposes of organization and brevity.

3.1.3 Basic Magnetic Circuit Equations

In an electric circuit analysis, the engineer's primary interest is to determine the voltages and currents present at any point in the circuit at a given instant in time–in other words, the *state* of the circuit. With the state known, useful quantities such as power output or signal properties may be calculated.

Similarly, the goal of a magnetic analysis is to determine the state of a magnetic circuit of interest. Here, the two quantities which compose the state are the magnetomotive force (MMF) Ψ and the magnetic flux Φ. With the magnetic state known, we can apply coupling relationships to determine the mechanical force and velocity output of the magnetic system as a function of the electrical input current and voltage. Once accomplished, the system may be treated as electromechanical in nature, requiring little further consideration of the magnetic properties.

Underlying Assumptions

To this end, we seek a straightforward set of equations, similar to the set of network and constitutive equations used in electric analysis, to determine the state of the magnetic system. Before deriving the circuit equations however, it is crucial to state the assumed conditions under which they hold. Understanding and respecting these assumptions is very important, since they restrict the magnetic designer to very specific configurations. Where practical machinery is concerned, the assumptions apply to a much narrower range of configurations than those of electrical design. The designer of a magnetic machine has the added burden of ensuring that the characteristics of the physical device resulting from the analysis are consistent with the original assumptions to an acceptable level of accuracy–a task which may not require much consideration in an electrical design, because most practical electric circuit components very closely mimic their ideal mathematical representations.

The conditions for a valid magnetic circuit analysis are:

1. **The 'conductors' in a magnetic circuit are assumed to be infinitely permeable**; that is, they sustain no magnetic field (**H**) while carrying a given flux. This is analogous to an ideal electrical conductor, which sustains no electric field while carrying a given current[3].

2. **Leakage flux is assumed to be zero.** This ensures that all flux is carried only by the conductors specified in the circuit. The real leakages present in a physical circuit are modeled by lumped-parameter reluctances, whose values are determined empirically or through a numerical analysis.

---

[3] This condition may be relaxed if the designer wishes to consider the real permeability of the conductors to gain additional accuracy in the circuit analysis. This is similar to considering the finite conductivity of a wire in an electric circuit. Engineering judgement must dictate such a decision.

3. **All flux is assumed to be uniform and normal to the circuit conductors or airgap pole faces.** Thus, no fringing is possible, and fluxes may be defined by scalar values with a reference direction. This is similar to the practice in electrical circuits of defining current as a scalar referenced to a positive direction arrow.

   To account for fringing in a more detailed analysis, one may resort to the finite element method, or to a flux-tube closed form analysis such as that described in Roters [26].

4. **Magnetic circuits do not dissipate energy.** This is the key difference between electrical and magnetic circuits. Whereas an electrical circuit contains may types of passive lumped-parameter elements, some that store energy (inductors, capacitors) and some that dissipate (resistors), a magnetic circuit contains only one type of passive lumped parameter element: A reluctance, which *stores energy* in a manner analogous to an electrical capacitor (not an inductor).

Further discussion of the last assumption is warranted. Because the energy storage property of a magnetic reluctance is not analogous to the energy dissipation property of an electrical resistor, the interpretation of a magnetic circuit may be confusing. The reason for this is that the usual schematic symbol for a magnetic reluctance is the same as an electrical resistor. The choice of symbol is simply an unfortunate consequence of convention. Still, it is important to remember that unlike an electrical circuit, whose state variables are an effort (voltage) and a flow (current), the magnetic circuit state variables are an effort (MMF) and a *displacement* (magnetic flux). The constitutive relation for a reluctance relates effort to displacement, and thus parallels the behavior of a mechanical spring or an electrical capacitor, not of a dissipative resistor[4]. This is in contrast to most of the available texts on the subject, such as [27], which incorrectly assert that magnetic flux and reluctance are analogues of electrical current and resistance. Correct electrical and magnetic representations of

---

[4] Additional discussion on this point may be found in Chapter 11 of [7].

a mathematically equivalent circuit are shown in Figure 3-2.

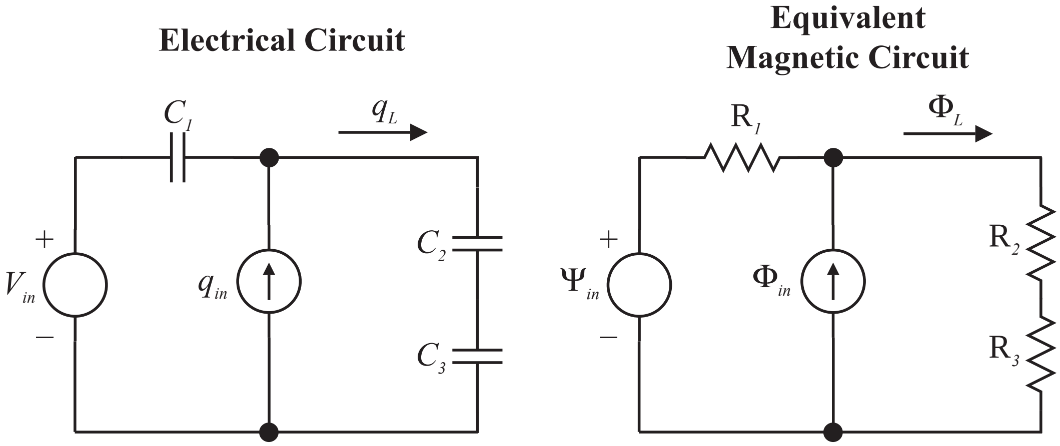

Figure 3-2: Electrical and magnetic circuit analogues with identical energy storage properties. Note that q_{in} and q_L represent charges, not currents.

This is not to contend, however, that physical magnetic systems do not dissipate energy. In fact they do, mostly through hysteresis and eddy current effects. The assumption of energy conservation in a magnetic circuit analysis is simply a convenient analytical tool. When modeling real systems, the losses resulting from dissipation of magnetic energy are often accounted for either in the electrical domain as an equivalent resistance, or in the mechanical domain as an equivalent damper. The former is the most common approach.

Circuit Equations

With the assumptions established, we are now in a position to derive the magnetic circuit equations. We require constitutive laws for the lumped-parameter circuit elements, both passive (reluctances) and active (coils and permanent magnets). We also require relations to describe the networking of the elements; i.e. node and loop equations analogous to KCL and KVL. I derive the latter first.

The magnetic equivalent to Kirchhoff's Current Law may be found directly from Gauss's law for magnetic fields, Equation 3.4. First, define the flux Φ as

$$\Phi \equiv \oint_S \mathbf{B} \cdot \mathbf{n}\, da \tag{3.12}$$

By assuming uniform normal flux through n magnetically permeable "conductors" penetrating various areas a_k of the contour surface S, such that $B_k a_k = \Phi_k$, we arrive immediately at the desired result:

$$\sum_{k=1}^{n} \Phi_k = 0 \tag{3.13}$$

The equivalent to Kirchhoff's Voltage Law is similarly direct. This time, we start with Ampere's law, assuming no time-varying electric displacement field ($\frac{\partial \mathbf{D}}{\partial t} = 0$). With a lumped parameter assumption, we assume m discrete changes in magnetic field potential (MMF), or $H_h l_h = \Psi_h$, around a circuit loop. Also assume that all sources of electrical current density are comprised of p coils, each with N_j turns and carrying current i_j. Using Equation 3.1, the result is then

$$\sum_{h=1}^{m} \Psi_h = \sum_{j=1}^{p} N_j i_j \tag{3.14}$$

Note that the form of this equation is slightly different than the electrical KVL, which states that the sum of the voltages around a closed circuit loop is zero. The reason for this is that one of the circuit source models–that of the coil–is captured on the right-hand side of Equation 3.14.

Constitutive Relations

The constitutive relation for a *linear* lumped-parameter reluctance is equivalent to Hooke's law for a linear mechanical spring, and is given by

$$\Psi = \mathcal{R}\Phi \tag{3.15}$$

where the proportionality constant $\mathcal{R}$ is the magnetic reluctance. For a magnetic material of relative permeability μ_r, cross-sectional area A, and length l, it is defined as

$$\mathcal{R} = \frac{l}{\mu_r \mu_o A} \tag{3.16}$$

For an airgap, the relative permeability is unity; for many soft magnetic materials, μ_r can be several thousand (which justifies their treatment as perfect conductors in the circuit model). Note that in initial analysis we often treat soft magnetic materials as perfect conductors, which have infinite permeability–and hence, zero reluctance.

Finally, the constitutive relation for a *linear* permanent magnet[5] is given by

$$\mathbf{B} = \mu_o \mathbf{H_m} + \mathbf{B_m} \tag{3.17}$$

where $\mathbf{H_m}$ is the magnetic field intensity applied to the magnet, and $\mathbf{B_m}$ is the residual induction of the magnet in tesla (this term may also be expressed as $\mu_o \mathbf{M_o}$, where $\mathbf{M_o}$ is the magnetization density in amperes per meter). Multiplying by the magnet pole face area A_m allows us to write Equation 3.17 in a form suitable for circuit analysis (note that here we use Assumption 3 and drop the vector notation):

$$\Phi = \mu_o A_m H_m + B_m A_m = \Phi_s + \Phi_m \tag{3.18}$$

This is a "Norton equivalent" representation of the energy supplied by the magnet. Here the magnet is modeled as an ideal flux source Φ_m in parallel with a shunt flux Φ_s due to the magnet's own internal reluctance, $\mathcal{R}_m = \frac{l_m}{\mu_o A_m}$. The value of the shunt flux depends on the the external magnetic field applied to the magnet, H_m. In most machine applications this value is negative, such that it acts to demagnetize the magnet[6]. See figure 3-3.

An alternative approach is to multiply 3.18 by the magnet reluctance to obtain

$$\Psi = H_m l_m + M_o l_m = \Psi_s + \Psi_m \tag{3.19}$$

which is a "Thévenin equivalent" representation of the magnet source, as shown in

---

[5] The assumption of linearity is valid for the powerful rare-earth (neodymium-iron-boron or samarium cobalt) permanent magnets in typical operation (second quadrant of the B-H curve), but not for weaker alnico or ferrite types.

[6] Notice that the relation for the magnet internal reluctance uses a relative permeability of one. This interesting result is due to the fact that once the magnet is fully charged, the material becomes saturated and cannot sustain additional magnetic flux in response to an applied magnetic field. Thus, its permeability after charging is identical to that of a vacuum, even though it sources a large magnetic field!

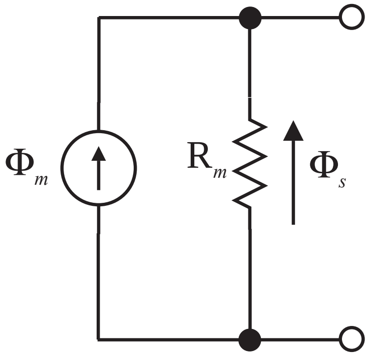

Figure 3-3: Norton equivalent circuit of a linear permanent magnet source.

Figure 3-4. Here, the model is of an ideal MMF source Ψ_m in series with the magnet's internal reluctance, such that a potential drop Ψ_s appears in response to an externally applied field.

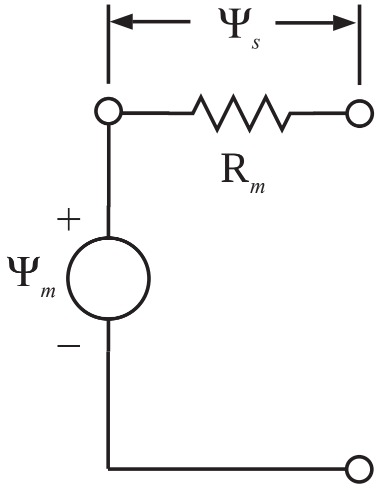

Figure 3-4: Thévenin equivalent circuit of a linear permanent magnet source.

Electromechanical Coupling

The final equations in our analysis are used to couple the magnetic state of the machine to the electrical and mechanical subsystems around it. Note that while the following equations are not required to solve for the magnetic state, we are generally not interested in the magnetic state by itself. Rather, we seek to determine its influence on the electrical or mechanical portions of the system. In most actuators the electrical signals are the input while the mechanical signals are the output; however, in electrical generation and mechanical sensing applications the opposite may be true.

To determine the effect of the magnetic state on the electrical system, we must find the voltage at the coil terminals. The coil current is already specified by Equation 3.14, so once the voltage is found the electrical state will be known.

The following derivation is based on Section 2.1.1 of [29]. First, consider a single open loop of a perfect electrical conductor, through which a magnetic field **B** passes. Assuming the ends of the loop are in a region with no magnetic field, we can use Faraday's law (Equation 3.2) to determine the voltage at the ends. We define the closed contour C to lie along the conductor, with a line connecting the open ends. We then have

$$\oint_C \mathbf{E} \cdot d\mathbf{l} = -v \tag{3.20}$$

which results in

$$v = \frac{d}{dt} \int_S \mathbf{B} \cdot \mathbf{n} \, da \tag{3.21}$$

Defining the *flux linkage* λ as

$$\lambda = \int_S \mathbf{B} \cdot \mathbf{n} \, da \tag{3.22}$$

gives the terminal voltage as simply

$$v = \frac{d\lambda}{dt} \tag{3.23}$$

The definition of flux linkage is included here because of its use in most magnetics texts. However, it becomes an auxiliary term once we recognize that, given our fundamental assumptions, the integral on the right-hand side of Equation 3.22 is simply the flux Φ passing through the coil loop! Since we are considering a coil with N turns (loops), the final terminal voltage equation is simply

$$v = N \frac{d\Phi}{dt} \tag{3.24}$$

The only tricky part here lies in determining the variation of Φ with time. In most machines, Φ is a function of an electrical variable–namely, the current in the coil(s)–and a varying mechanical displacement, typically due to a changing airgap. The latter is introduced into the magnetic equations as a variable reluctance. For such a system, we can take partial derivatives of Equation 3.24 to obtain

$$v = N \frac{\partial \Phi}{\partial i} \frac{di}{dt} + N \frac{\partial \Phi}{\partial x} \frac{dx}{dt} = L \frac{di}{dt} + K_m \frac{dx}{dt} \tag{3.25}$$

where i is the coil current, x is the mechanical displacement, $L \equiv N \frac{\partial \Phi}{\partial i}$ is the *inductance*, and $K_m \equiv N \frac{\partial \Phi}{\partial x}$ is the *motor constant*. Note that in assuming L and K_m to be constants, we are implicitly making the assumption that Φ varies linearly in i and x, which holds true for many designs when operating away from saturation. Determining L and K_m analytically can be computationally messy, but is mathematically straightforward.

The final required equation relates the magnetic state to the mechanical output of the system, and in particular the applied force. For this relation we make use of the Maxwell stress tensor[7], Equation 3.6. With no applied electrical field, it immediately simplifies to

$$\sigma_{ij} = \frac{B_i B_j}{\mu_o} - \frac{1}{2} \left(\frac{B^2}{\mu_o} \right) \delta_{ij} \tag{3.26}$$

Per our fundamental assumptions, we simplify by considering only normal magnetic fluxes, so $i = j$ and $\delta_{ij} = 1$. Also, we have $B = \frac{\Phi}{A}$, where A is the cross-sectional area of the flux path. This gives

$$\sigma_n = \frac{F_n}{A} = \frac{\Phi^2}{2\mu_o A^2}$$

---

[7] Many texts prefer to derive magnetic force results using using magnetic energy and/or co-energy relations in lieu of the stress tensor approach. See, for example [19].

$$F_n = \frac{\Phi^2}{2\mu_o A} \tag{3.27}$$

Note that in certain applications that involve shear stresses (such as a typical rotating electric motor), the forces resulting from the stress tensor become more difficult to calculate, because the flux vectors have two- or three-dimensional geometry that must be resolved into components to calculate the contributions in i and j.

3.2 AFSM Magnetic Design

Equipped with the analytical tools derived above, we are now in a position to design the AFSM actuator. For the configuration, I chose the normal force flux steering arrangement used by Lu [16] for two reasons. First, it provides a thermally-limited steady-state force per unit actuator volume about 100 times higher than a Lorentz force (voice coil) actuator. Clearly, a high force output is critical to achieving acceleration and bandwidth performance. Second, the actuator force output is approximately linear in both current and mechanical stroke. The linear actuator properties allow the use of a simple linear controller, which is consistent with the design goals for the control system. Although designing a nonlinear controller is possible, it would have to be done digitally and may impose limits on the system dynamic performance, for the same reasons outlined in Chapter 2.

The flux-steering actuator component detail is shown in Figure 3-5, and its magnetic operational principles are detailed in Figure 3-6. At the heart of the actuator is a permanent magnet, which generates a DC biasing flux that passes through the armature and returns through the upper and lower core halves. This flux is indicated by the magenta-colored arrows in Figure 3-6. As will be shown mathematically, the amount of flux passing through one core half or another is determined by the position of the armature, as well as any steering flux from the coils. With the armature in the neutral position and no coil current, the magnet flux is evenly divided between the upper and lower core halves, and thus there is no net force on the armature.

The two actuator coils are wired in series, and thus behave as if they were one

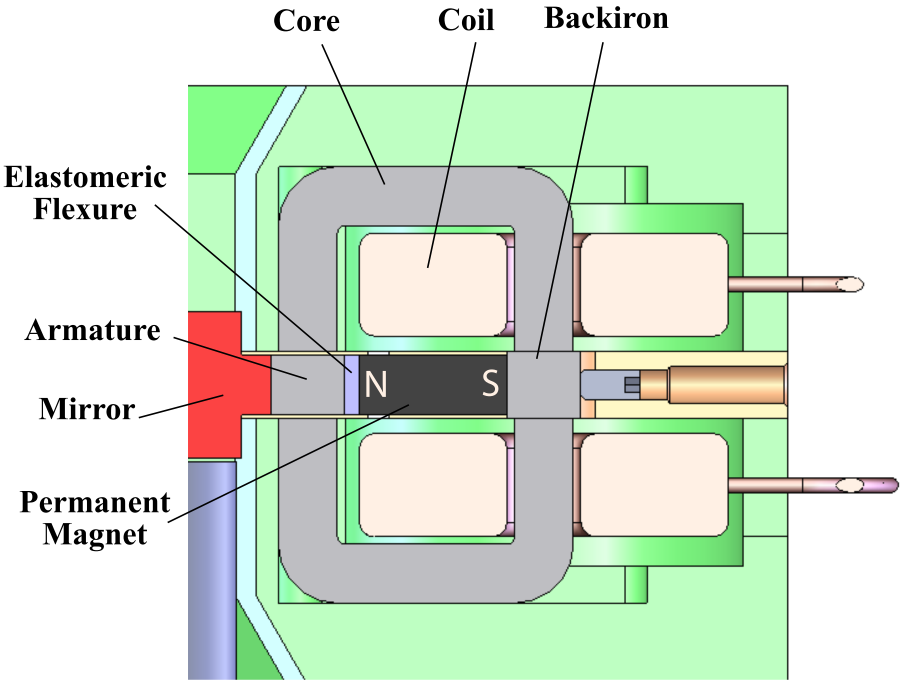

Figure 3-5: Cross-sectional view of the magnetic actuator showing component detail.

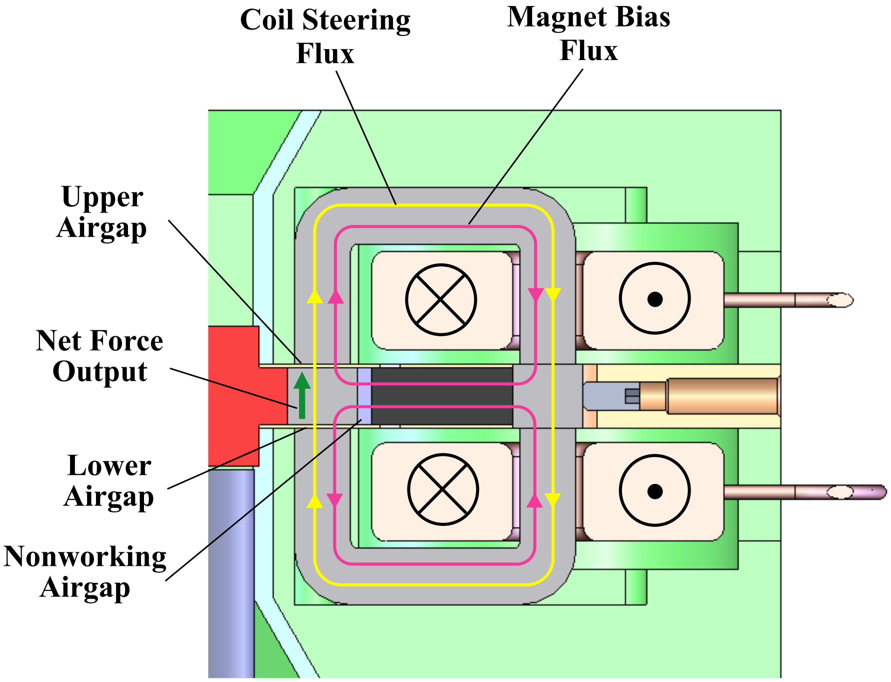

Figure 3-6: Flux-steering actuator principle of operation.

coil. In Figure 3-6, current flowing into the page is denoted symbolically by the circle with an ‘X’ through it, while current out of the page is indicated by a circle with a dot. Per the right-hand rule[8], applying a positive-valued current will produce a magnetic flux pointing downward, as indicated by the yellow arrows in Figure 3-6. Recall that since the permanent magnet has a high reluctance to externally applied magnetic fields, the preferred path for the coil flux is through the cores and armature.

The interaction between the DC permanent magnet flux and the time-varying coil flux produces the actuator output force. Considering the lower core half, it is clear that the magnet flux acts in the opposite direction as the coil flux, and therefore if the two flux magnitudes are equal, they cancel each other. In the upper core half, the opposite is true. Here, the coil flux acts in the same direction as the permanent magnet flux, producing an additive effect. As a result, applying a positive coil current causes the flux passing through the upper airgap to increase, while simultaneously reducing the flux through the lower airgap. Per Equation 3.27, this causes a net upward force on the armature. It is easy to see also that reversing the current direction will reverse the steering flux, and thus produce a downward force on the armature.

3.2.1 Magnetic Circuit Analysis

With the basic operational principle in mind, we must now analyze the magnetic circuit to quantitatively determine the actuator output. Examining the magnetic elements of Figure 3-5 and the flux paths shown in Figure 3-6, we can construct a circuit diagram to represent the behavior of the magnetic network. The diagram is shown in Figure 3-7.

Note that the circuit layout is chosen to roughly correspond to the physical layout of Figure 3-6. Starting in the center with the permanent magnet, we use a Thévenin equivalent representation of an ideal MMF source Ψ_m in series with the magnet internal reluctance $\mathcal{R}_m$. The flux Φ_m through this circuit leg must cross the nonworking

---

[8] The right-hand rule for a coil states that if you curl the fingers of your right hand along the direction of positive current flow, the induced magnetic field vector points in the direction of your thumb.

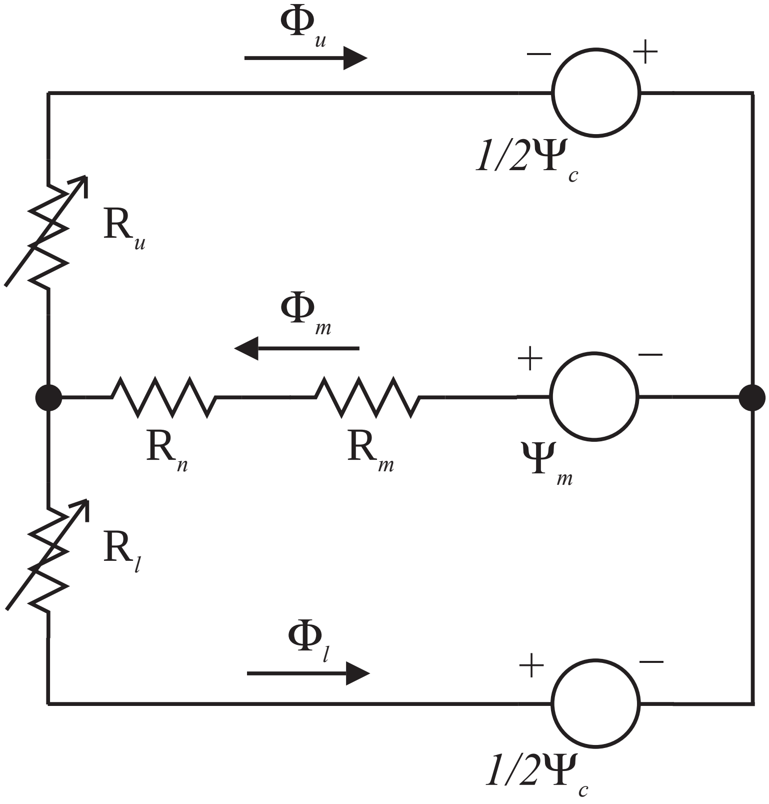

Figure 3-7: Magnetic circuit representation of the flux-steering actuator.

airgap between the magnet's north pole and the armature (note that this 'gap' is physically filled by the elastomer flexure, but from a magnetic standpoint we can simply model it as air). This airgap reluctance is denoted by $\mathcal{R}_n$. Inside the armature, the circuit branches into two symmetrical flux paths. The reluctance $\mathcal{R}_u$ represents the upper armature airgap, which varies depending on the armature position. The flux through this gap is denoted by Φ_u; it links the upper coil, which we represent as $1/2\Psi_c$ because the coils are wired in series.

The lower branch of the circuit diagram is basically identical; however, note that the sign of the coil MMF source is switched in order to be consistent with the right-hand rule established by the upper coil. The result of this is that the flux generated by the lower coil is opposite of the defined direction of Φ_l. It will be important to keep track of these signs in the analysis.

Flux Analysis

To begin, we first make use of Equation 3.14 to write loop equations for the magnetic potential changes around the circuit loops. For the potential drops of each reluctance, we employ the constitutive law given by Equation 3.15. For the upper, lower, and outer loops respectively, the results are:

$$\mathcal{R}_m\Phi_m + \mathcal{R}_n\Phi_m + \mathcal{R}_u\Phi_u = 1/2\Psi_c + \Psi_m \tag{3.28}$$

$$\mathcal{R}_m\Phi_m + \mathcal{R}_n\Phi_m + \mathcal{R}_l\Phi_l = -1/2\Psi_c + \Psi_m \tag{3.29}$$

$$\mathcal{R}_u\Phi_u - \mathcal{R}_l\Phi_l = 1/2\Psi_c + 1/2\Psi_c = \Psi_c \tag{3.30}$$

Notice, however, that Equation 3.30 is a linear combination of Equations 3.28 and 3.29. It is obtained by subtracting 3.29 from 3.28, and therefore contains no new information. To find the third equation required to solve for the unknown fluxes, we use Equation 3.13 to analyze the circuit nodes. If we define the node on the right hand side of Figure 3-7 to be at ground potential, we only need to consider the left-hand node to obtain

$$\Phi_m - \Phi_u - \Phi_l = 0 \tag{3.31}$$

We now have three equations in three unknowns, which we can represent in matrix form as (with $\mathcal{R}_{mag} \equiv \mathcal{R}_m + \mathcal{R}_n$):

$$\begin{bmatrix} \mathcal{R}_{mag} & \mathcal{R}_u & 0 \\ \mathcal{R}_{mag} & 0 & \mathcal{R}_l \\ 1 & -1 & -1 \end{bmatrix} \begin{bmatrix} \Phi_m \\ \Phi_u \\ \Phi_l \end{bmatrix} = \begin{bmatrix} 1 & \frac{1}{2} & 0 \\ 1 & -\frac{1}{2} & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \Psi_m \\ \Psi_c \\ 0 \end{bmatrix} \tag{3.32}$$

Though requiring a bit of algebra, this system is easily solved for the unknowns Φ_m, Φ_u, and Φ_l to obtain

$$\Phi_m = \frac{1/2\Psi_c(\mathcal{R}_l - \mathcal{R}_u) + \Psi_m(\mathcal{R}_l + \mathcal{R}_u)}{\mathcal{R}_{mag}(\mathcal{R}_l + \mathcal{R}_u) + \mathcal{R}_l\mathcal{R}_u} \tag{3.33}$$

$$\Phi_u = \frac{\Psi_c(\mathcal{R}_{mag} + 1/2\mathcal{R}_l) + \Psi_m\mathcal{R}_l}{\mathcal{R}_{mag}(\mathcal{R}_l + \mathcal{R}_u) + \mathcal{R}_l\mathcal{R}_u} \tag{3.34}$$

$$\Phi_l = \frac{-\Psi_c(\mathcal{R}_{mag} + 1/2\mathcal{R}_u) + \Psi_m\mathcal{R}_u}{\mathcal{R}_{mag}(\mathcal{R}_l + \mathcal{R}_u) + \mathcal{R}_l\mathcal{R}_u} \tag{3.35}$$

The magnetic state is now completely known. The remaining tasks are to evaluate the reluctances in terms of the machine parameters, redefine Ψ_m and Ψ_c in terms of the magnet and coil constitutive relations, and finally to apply Equation 3.27 to the working airgap fluxes to obtain the actuator force output.

To specify the airgap reluctances, we need the length and area of each gap. In the AFSM design, I set the nonworking airgap area, magnet pole area, and upper and lower working airgap areas equal, and defined them under a common variable a_p[9]. The fixed length of the combined gap is $l_m + g_n$. The upper and lower working airgap lengths are variable and related to one another, since closing the upper airgap opens the lower by an equal amount, and vice versa. Defining the (equal) gap length when the armature is in the neutral position as x_o, the upper and lower gaps at an upward armature displacement x are, respectively:

$$g_u = x_o - x \tag{3.36}$$

$$g_l = x_o + x \tag{3.37}$$

Using Equation 3.16, the reluctances are then

---

[9] The rationale for setting theses areas equal is that in the null position with no current, the bias flux density will be equal to half the residual magnet flux density. When the coils are energized to drive the one working gap flux to zero, the flux density in the other gap will be about equal to the residual magnet flux density, and will not saturate.

$$\mathcal{R}_{mag} = \frac{l_m + g_n}{\mu_o a_p} \tag{3.38}$$

$$\mathcal{R}_u = \frac{x_o - x}{\mu_o a_p} \tag{3.39}$$

$$\mathcal{R}_l = \frac{x_o + x}{\mu_o a_p} \tag{3.40}$$

Per Equation 3.14, the coil MMF (from both coils together) is

$$\Psi_c = N i_c \tag{3.41}$$

where N is the *total* number of turns for both coils, and i_c is the coil current. Finally, the magnet MMF is, using Equation 3.19

$$\Psi_m = M_o l_m \tag{3.42}$$

Note here that the internal reluctance term in Equation 3.19 has already been captured in the nonworking airgap reluctance.

Terminal Voltage

We now use the flux linkage relation given in Equation 3.24 to determine the voltage across the coils. Noting the flux directions denoted in Figure 3-7, we add the contributions from each coil to obtain

$$v = \frac{N}{2}\frac{d\Phi_u}{dt} + (-\frac{N}{2}\frac{d\Phi_l}{dt}) = \frac{N}{2}\frac{d}{dt}(\Phi_u - \Phi_l) \tag{3.43}$$

Substituting Equations 3.34 and 3.35 into the above gives

$$v = \frac{N}{2}\frac{d}{dt}\left(\frac{\Psi_c(2\mathcal{R}_{mag} + 1/2(\mathcal{R}_l + \mathcal{R}_u)) + \Psi_m(\mathcal{R}_l + \mathcal{R}_u)}{\mathcal{R}_{mag}(\mathcal{R}_l + \mathcal{R}_u) + \mathcal{R}_l\mathcal{R}_u}\right) \tag{3.44}$$

Finally, we make use of the constitutive relations given in Equations 3.38 through 3.42, which after some algebra results in

$$v = \frac{d}{dt}\left(\frac{N^2\mu_o a_p(l_m + g_n + x_o)i_c + N\mu_o a_p M_o l_m x}{2x_o(l_m + g_n) + x_o^2 - x^2}\right) \tag{3.45}$$

This equation gives the coil voltage in terms of the physical parameters of the actuator. Although there are many terms, most of them are constant. The only time-varying quantities are the coil current i_c and the armature displacement x. Defining the leading constant terms as

$$\begin{aligned}
\alpha &\equiv N^2\mu_o a_p(l_m + g_n + x_o) \approx N^2\mu_o a_p l_m \\
\beta &\equiv N\mu_o a_p M_o l_m \\
\gamma &\equiv 2x_o(l_m + g_n) + x_o^2 \approx 2x_o l_m
\end{aligned} \tag{3.46}$$

(with the approximations accounting for the fact that $l_m \gg x_o$ and $l_m \gg g_n$), we write the result very simply as

$$v = \frac{d}{dt}\left(\frac{\alpha i_c + \beta x}{\gamma - x^2}\right) \tag{3.47}$$

The next step is to differentiate Equation 3.47 with respect to time. We use equation 3.25 to separate the current and displacement terms. The partial derivatives with respect to i_c and x are, respectively

$$L = \frac{\partial}{\partial i_c}\left(\frac{\alpha i_c + \beta x}{\gamma - x^2}\right) = \frac{\alpha}{\gamma - x^2} \tag{3.48}$$

$$K_m = \frac{\partial}{\partial x}\left(\frac{\alpha i_c + \beta x}{\gamma - x^2}\right) = \frac{\beta\gamma + 2\alpha i_c x + \beta x^2}{(\gamma - x^2)^2} \tag{3.49}$$

The above equations are both nonlinear in x, due to the x^2 term in the denominator of Equation 3.47. We can simplify these relations, however, by realizing that the armature displacements are small, which makes the x^2 term negligible relative to the γ term. Applying this approximation results in constant values for L and K_m, as follows:

$$L \approx \frac{\alpha}{\gamma} = \frac{N^2\mu_o a_p(l_m+g_n+x_o)}{2x_o(l_m+g_n)+x_o^2} \approx \frac{\mu_o N^2 a_p}{2x_o} \tag{3.50}$$

$$K_m \approx \frac{\beta}{\gamma} = \frac{N\mu_o a_p M_o l_m}{2x_o(l_m+g_n)+x_o^2} \approx \frac{\mu_o M_o N a_p}{2x_o} \tag{3.51}$$

In the AFSM design, the approximations are valid due to the small 250 μm stroke. As a point of reference, the actual γ value in the AFSM was 6.7×10^{-6} m^2, while the value of x^2 was $\leq 6.3 \times 10^{-8}$ m^2, about 100 times smaller. Thus, the variation of L and K_m with stroke was less than 1%. The final combined voltage at the actuator terminals is:

$$v = \frac{N^2\mu_o a_p(l_m+g_n+x_o)}{2x_o(l_m+g_n)+x_o^2}\frac{di_c}{dt} + \frac{N\mu_o a_p M_o l_m}{2x_o(l_m+g_n)+x_o^2}\frac{dx}{dt} \tag{3.52}$$

$$\approx \frac{N^2\mu_o a_p}{2x_o}\frac{di_c}{dt} + \frac{N\mu_o a_p M_o}{2x_o}\frac{dx}{dt} \tag{3.53}$$

Actuator Force

The actuator force output is given by Equation 3.27. Here, we must sum the contributions from the upper gap and lower gap fluxes to obtain the net force output:

$$F_{act} = \frac{\Phi_u^2 - \Phi_l^2}{2\mu_o a_p} \tag{3.54}$$

Also, the force in the nonworking (radial) direction is given by

$$F_{nw} = \frac{\Phi_m^2}{2\mu_o a_p} \tag{3.55}$$

This force is reacted by the elastomer flexure, as well as the opposing actuator.

Although we can substitute the equations for flux, reluctance, and coil and magnet MMF to obtain a closed-form solution for these forces in terms of the actuator parameters, as was done for the voltage, the expression becomes quite messy algebraically due to the need to square the flux and reluctance terms. Instead, I performed the calculations separately during my numerical exploration of the design space.

With the electromechanical coupling relations established, the actuator analysis is complete. In Appendix A, I show that the results obtained above using circuit analysis are identical to a direct analysis using Maxwell's equations. In the remaining sections of this chapter, I discuss the design of the physical actuator components.

3.2.2 Material and Geometry Selection

The final task in the actuator design is to turn the numbers into a physical machine. To do this for the AFSM, I developed a spreadsheet using the equations derived in the previous section, and performed a numerical study of the actuator performance (output force), as well as its effect on the mechanical hardware (mirror acceleration). The magnetic equations reveal the following relationships regarding the actuator performance, which I used as guidelines for my study:

- From Equation 3.54, it is evident that the force output depends on the *magnitude* of the fluxes in the working airgaps, not the direction, due to the fact that they are squared[10]. The sensitivity of the actuator (i.e. the change in force with a given change in current) is directly proportional to the amount of biasing flux from the magnet. This flux determines the (fixed) differential between the upper and lower gap fluxes.

- The limitation on maximum force is determined by the maximum flux through the airgaps, which is in turn limited by pole area, the magnetic materials used and the available coil MMF. High flux is obtained either by increasing the pole area a_p or the flux density B. Since increasing a_p negatively affects the armature inertia, the best choice is to maximize B. The limitations in doing this are the saturation characteristics of the magnetic conductor, and the current limit of the coils. With B established, a_p is then selected to provide the required force.

---

[10] This is also consistent from the perspective of energy minimization. The airgaps represent the highest energy storage gradients (potentials) in the system, so the magnetic forces act to close the gaps, thereby minimizing the total energy in the system.

- The permanent magnet and the coil both generate MMF to establish the airgap flux; however the magnet is the largest contributor by far. This can be seen by setting Equation 3.34 or 3.35 to zero and solving for the coil MMF. Using Equation 3.35, the result is

  $$\Psi_c = \Psi_m \frac{\mathcal{R}_u}{\mathcal{R}_{mag} + 1/2\mathcal{R}_u} \tag{3.56}$$

  Since $\mathcal{R}_u$ is much smaller than $\mathcal{R}_{mag}$, it is easy to see that Ψ_c is much smaller than Ψ_m (in the case of the AFSM, the ratio was about 3%). This information is useful to size both the magnet and the coil.

With the above guidelines in mind, I used my spreadsheet to set the numerical parameters for each component of the actuator.

Armatures

As mentioned in Chapter 2, the design of the armatures must consider mechanical as well as magnetic factors. Unfortunately, these factors directly compete. Even though a low inertia is mechanically desirable, the magnetic requirements limit the armature density and volume to minimum values. The former limit is due to the fact that materials with magnetic permeability high enough to be treated as magnetic conductors all have densities similar to iron (8 g/cm^3). The latter is because all cross sections of the armature normal to the actuator flux path must maintain a minimum cross-sectional area to avoid magnetic saturation.

In the case of the chosen SM-2 alloy, the saturation flux density is about 1.4 tesla. By comparison, the flux density output of the permanent magnet (considering the drops due to the airgaps) is about 1.1 tesla. Thus, to avoid saturation, the nonworking airgap armature pole area can only be made slightly smaller, at minimum, than the magnet pole area. However, making the armature area smaller than the magnet incurs excessive fringing and leakage losses, so I chose them to be equal.

Since the actuator is designed to drive all of this flux to one working airgap or the other, the same constraints are present for the working airgap armature pole areas.

In addition, a sufficient area is required to produce the desired force output. Since increasing the radial dimension of the armature has the consequence of increasing the inertia quadratically, I chose instead to increase the tangential dimension. This decision resulted in the long thin armature geometry used in the final design. It represented the best compromise between structural integrity, flux-carrying capacity, force output, and inertia. The final armature design characteristics are shown in Table 3.1

Table 3.1: Armature Design Characteristics

Property	**Value**
Dimensions	20x5x4 mm
Material	SM-2 HB
Saturation Flux Density	1.4 tesla
Flux Density at Maximum Force	1.1 tesla
Active Pole Area	80 mm^2
Maximum Force Output	36.9 N

Permanent Magnet

The force and inertia requirements that set the length and height of the armature in turn place constraints on the height and width of the permanent magnet. Since in this design there is no intermediary piece of soft magnetic material to focus the magnet flux into the armature, the magnet pole face itself must match up directly with the armature nonworking airgap area. If it did not, the flux from the magnet would certainly leak to the stationary actuator components as mentioned above, and thus reduce the efficiency of the output.

Given that the area is set, two variables remained to obtain the required amount of energy storage in the magnet: the magnetization density M_o (or equivalently, the residual induction B_o, a quantity preferred by most manufacturers) and the magnet length. The range and performance envelope of available permanent magnet materials limit the former, while machine envelope constraints limit the latter.

The obvious choice for the magnet material is neodymium-iron-boron (NdFeB). It has by far the highest residual induction of any modern material, is magnetically linear over a wide range of magnetic field environments (important, since this is assumed by Equation 3.17), and also has a high coercive force (resistance to demagnetization). The one drawback to NdFeB is its relatively low range of allowable operating temperature. Certain grades of NdFeB can be charged to B_o values as high as 1.45 tesla; but these are the most sensitive to temperature increases. Since I was concerned that actuator heating might be a problem, I chose a material grade with a slightly lower value of B_o, 1.3 tesla, but with a comfortable temperature rating of 150°C. As it turns out, heating was not an issue during testing of the AFSM prototype, and therefore I could have taken advantage of a higher grade material for increased energy density.

With the material chosen, I selected a magnet length to give the required amount of MMF bias to the actuator magnetic circuit. Based on my numerical evaluations, an adequate length was 10 mm, making the final magnet dimensions 20x10x4 mm. The direction of magnetization was along the 10 mm dimension.

Cores

The configuration of the magnetic magnetic cores, both in material and geometry, involved considerable deliberation during the actuator design. The cores are the longest part in the magnetic circuit, and thus represent the largest potential for losses in the conduction path. They are also important because they heavily influence the form factor and cost of both the actuator and the prototype as a whole.

In order to mitigate eddy current losses, magnetic conductors that operate at high frequencies (such as those used in electrical power transformers) are generally constructed using a series of thin laminations with electrically insulating material in between. Because of Faraday's law, electrical currents are induced in the magnetic conductor due to the time-varying magnetic field. These currents circulate within the conductor, and are dissipated as heat due to the resistivity of the material. As such, a percentage of the power supplied to the magnetic circuit is lost, resulting in reduced mechanical output. As the operating frequency increases, the losses due to eddy current effects become greater.

Although it does not exist in reality, an ideal magnetic conductor has infinite permeability to magnetic flux, but zero conductivity (that is, infinite resistivity) to electrical currents. Such a conductor cannot support the generation of eddy currents, because the property of zero conductivity means that current flow is impossible. Laminated core designs approach this ideal behavior by forcing the eddy current to flow inside their thin metallic strips. The thin laminates are oriented such that they provide a high permeability path for magnetic flux, yet impose a high resistivity path to the eddy currents that are are induced normal to the flux[11].

Since the AFSM is required to operate at frequencies up to 10 kHz, using such a laminated core design is essential. As several methods for creating laminated structures exist, I had to decide which one would be best suited for the AFSM. For maximum optical functionality, I first considered several design iterations using chemically-etched laminations with a low profile of varying cross section, as shown in Figure 3-8. I also considered core designs with an asymmetric profile, which would place the bulk of the actuator below the mirror. However, I discovered that the cost and lead time associated with this type of design was prohibitive.

A low-cost alternative to a laminated stack was a tape-wound core. These can be made cheaply and efficiently by winding a long thin ribbon of magnetic material (the “tape”) around a shaped mandrel. As it is wound, an adhesive is placed on the ribbon to both bind the layers together and provide an electrically insulating barrier. After winding, the core is cut in half and the cut surfaces are ground and etched to prevent shorting of the laminations and provide a uniform pole face.

The downside of using this type of design is that only simple geometries of uniform cross section are allowable, and symmetric designs are the most cost-effective. In order to accommodate this, I decided to increase the profile of the actuator at the expense of the optical window and made it symmetric, resulting in the profile shown in Figure 3-5 and in Chapter 1. Since for this project I was more interested in the mechanical

---

[11] By Lenz’s law, the induced currents orient themselves such that they in turn induce a magnetic field in the opposite direction as the applied field; that is, the current direction is normal to both the applied and induced magnetic field.

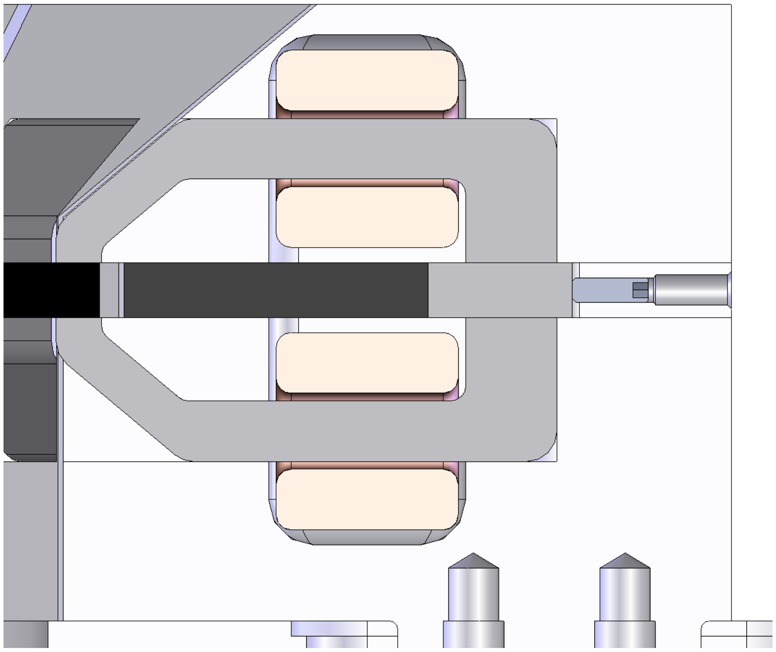

Figure 3-8: Early design iteration showing low-profile core for increased optical access to the mirror.

performance of the mirror itself than its impact on a fictitious optical platform, I deemed this an acceptable compromise.

I chose MK Magnetics in Adelanto, CA to wind the cores. The final characteristics of the core design are shown in Table 3.2, below.

Table 3.2: Actuator Magnetic Core Design Characteristics

Property	**Value**
Lamination Width	20 mm
Lamination Thickness	0.001" (0.0254 mm)
Winding Buildup	4 mm
Space Factor	0.83
Lamination Material	50% Ni-Fe ('Deltamax')
Saturation Flux Density	1.5 tesla
Maximum Relative Permeability μ_r	5×10^4 to 1.5×10^5
Initial Relative Permeability	1.7x10^4

Coils

The last step in the actuator design was to settle on a configuration for the coil. Coil design involves a number of factors in the magnetic, electrical, and thermal domains. Respectively, the coil must produce the correct magnetic field, at the specified current and voltage, without overheating due to its own resistivity.

With the permanent magnet output known, the maximum coil MMF is given by Equation 3.56. The resulting value is used to specify the applied current and the number of coil turns. Here some iteration is involved vis-a-vis the performance of the driving power amplifier in terms of current, voltage, and RMS power. Generally, the current i_c is chosen as some percentage of the amplifier's rated output, and then the number of turns N is set relative to this value. One must then consider the resulting terminal voltage at the operating frequency due the coil inductance, which is given by Equation 3.50. If the voltage is higher than the amplifier's maximum output, a more powerful amplifier must be specified.

Note that, aside from losses that vary slightly as the coil parameters are changed, the power throughput of the coil and actuator remains unchanged during these design iterations. Only the relationship between current and voltage is affected. Late in the AFSM design, I was able to make use of this fact to redesign the coil in response to the availability of low-voltage, high-current power amplifiers. The details of this change are described in Chapter 4.

With the coil current and number of turns established, the next step in the coil design is to determine the coil geometry as a function of thermal limits. The thermal properties of the coil wire determine the wire diameter, and this combined with the number of turns sets the required coil volume. Since a coil consisting of stacks of circular wire contains a large percentage of empty volume in addition to the metal conductor volume, a 'packing factor' must be applied to determine the physical coil dimensions. Both the thermal limit and the packing factor are design 'rules of thumb' that are established from experience. I consulted with Professor Trumper and my coil builder, Fred Sommerhalter, and found that an air-cooled copper coil wound with medium-gauge wire is conservatively limited to a current density[12] of 10^6 to 10^7 A/m^2 and a packing factor of 0.6 to 0.7.

Using these numbers, along with the magnetic MMF requirements and the amplifier output ratings, I designed the AFSM actuator coils with the following characteristics. Note that the tabulated values are for a single actuator; that is, identical upper and lower coils wired in series. Also, in order to increase the packing factor for each coil I used ten wires wrapped in parallel instead of a lower gauge, thicker wire (the equivalent gauge size is given in the table).

Table 3.3: Coil Design Characteristics

Property	**Value**
Required Coil MMF Ψ_c	257 A-turns
Number of Turns N	20
Wire Gauge (Diameter)	10 x #22 AWG (10 x 0.645 mm)
Equivalent Wire Gauge	#12 AWG
Packing Factor	0.7
Cross Section Dimensions	10 x 7 mm
Maximum Current	14 A
Maximum Current Density	4.28 x 10^6 A/m^2
Terminal Voltage at 5 kHz, 25 μrad	30.45 V
Coil Resistance	8.7 milliohms
Coil Inductance	67.9 microhenries
Peak Power Consumption	426 watts
RMS (Average) Power Consumption	213 watts

The final step in the analysis was to predict the output characteristics of the actuator based on my spreadsheet runs. These are given in Table 3.4, below.

With the actuator and mechanical components specified, I was ready to fabricate and assemble the AFSM hardware. This activity is the subject of the following chapter.

---

[12] From my industry experience, I also knew that a good rule of thumb, albeit in rather archaic units, is to design for greater than 300 "circular mils per amp" for continuous operation. A "circular mil" is the area of a 0.001" diameter circle, and is equal to 5.066×10^{-10} square meters.

Table 3.4: Actuator Force Output Characteristics

Property	**Value**
Peak Force at maximum current	36 N
Average Force	28 N
Actuator Magnetic Stiffness	-0.1246 N/μm

Chapter 4

AFSM Assembly and Electrical Integration

In this chapter I describe the fabrication and assembly of the AFSM hardware, as well as the integration of the power electronics to drive the system and the capacitance probes and optical feedback sensors that I used to measure the system performance. I also provide some insight into my philosophy for achieving precision in the design, especially with respect to airgaps and sensor placement. While not heavy in technical content, proper execution of these steps were nevertheless vital to obtaining successful test results later in the project, and so warrant the discussion that follows.

4.1 Mechanical Fabrication and Assembly

Although I had not originally intended to do so, I performed most of the AFSM component fabrication myself in the MIT Laboratory for Manufacturing and Productivity (LMP) machine shop. I did so mostly for economic reasons, since the fabrication and assembly budget for the AFSM was fixed and I anticipated needing additional funds for the later stages of the program. However, working in the machine shop has always been an enjoyable activity for me, so I was happy to have the opportunity do do so for this project.

On the other hand, my machining skills were fairly meager in comparison to what was needed to produce adequate parts for the AFSM prototype. Fortunately, I had plenty of help from the man running the LMP shop, Gerry Wentworth. Gerry is an expert machinist with decades of experience at MIT and Honeywell, and even more significantly, he has plenty of experience preventing wayward graduate students from destroying their thesis projects. Gerry provided lots of help in all areas of my fabrication activity, including the MasterCAM software for generating CNC programs, setting up and running the Bridgeport CNC mill, use of the surface grinder, and cutter grinding for some of the more delicate parts. I owe many thanks to Gerry for helping the fabrication go smoothly.

Figure 4-1: Gerry Wentworth and the Bridgeport Torq-Cut TC3 CNC mill, on which many of the AFSM parts were machined.

4.1.1 Machined Parts and Subassemblies

Below, I describe briefly some the fabrication of some of the important AFSM parts and subassemblies, and how their features influenced the AFSM performance.

Actuator Housings

The actuator housings perform several functions in the AFSM design. First, they provide structural support and positioning for the coil and cores. Second, they provide via their high thermal conductivity a means of cooling the actuators. Third and most importantly, through the precise thickness and parallelism of their upper and lower surfaces they provide a reference datum surface off which the armature airgaps are set.

The completed housings are shown in Figure 4-2. The pockets on the interior of the housings serve to position the cores and coils prior to potting them in place. By using the housing features this way, I could position the core poles relative to the backiron and armatures simply by referencing the external features of the housing, which made assembly much simpler. This is described further in Section 4.1.2.

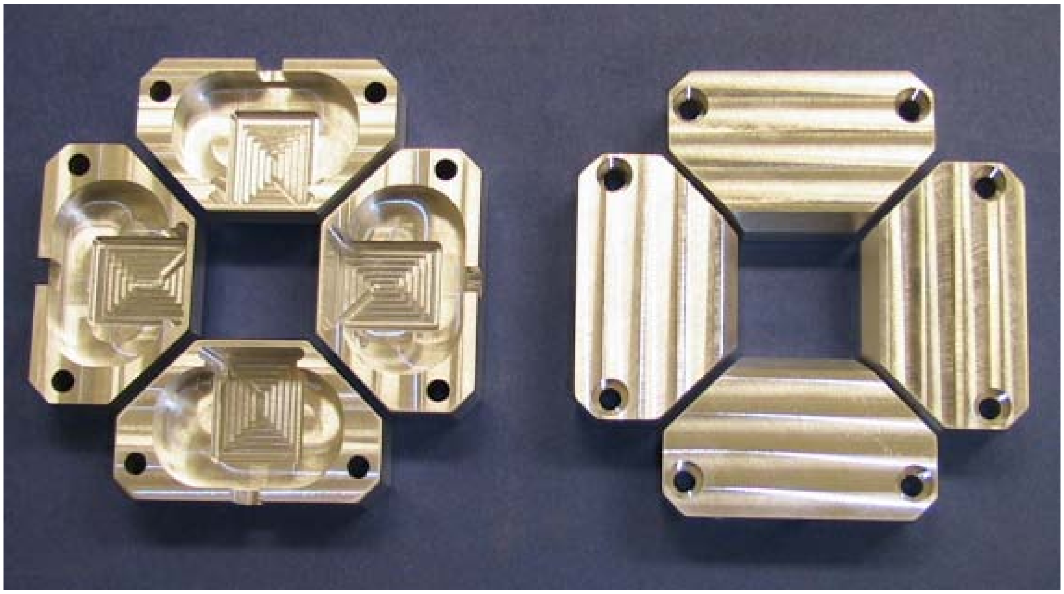

Figure 4-2: Fabricated ASFM actuator housings prior to core and coil installation.

AFSM Base Plate

The base plate for the AFSM provides structural support for all the active AFSM components, and therefore it contains several important design features. Most prominently, the flat upper surface is the primary datum feature relative to which the remaining critical features are positioned. Since my plan was to achieve precision in the design primarily through the easily manufacturable features of flatness, parallelism, and thickness, properly establishing this basis feature was essential. As such, I specified that it have a flatness of less than 10 microns. After machining I hand-lapped the surface to ensure that it met specification.

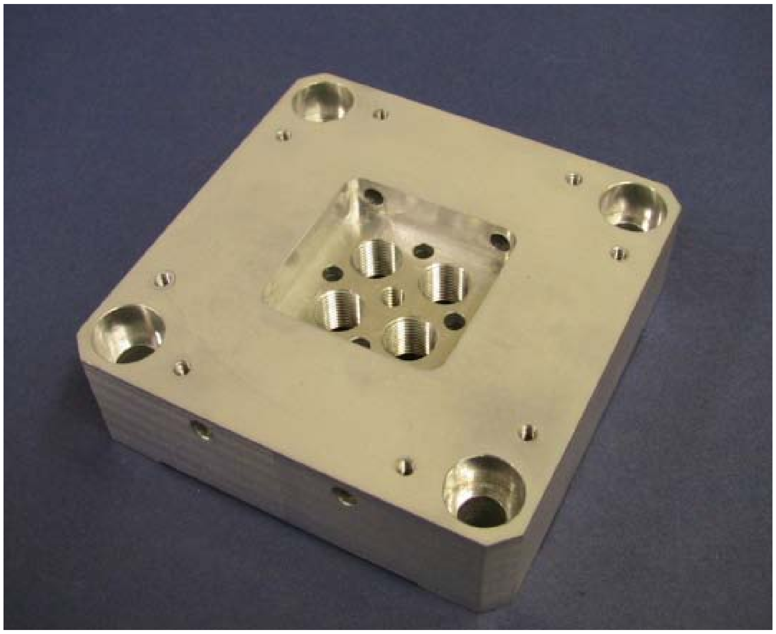

Figure 4-3: The fabricated AFSM base plate.

A few other features on the baseplate are noteworthy. First, it has a ratio of thickness to length of 30%, which ensures high bending stiffness. Given the high frequency range of the AFSM, exciting unwanted modes in the structure was a distinct possibility, and as such I took steps to prevent this. The high baseplate stiffness was one such measure. Also, to preserve precision, I milled relief features into the underside of the base plate. This ensured small, well-defined contact surfaces at the four bolt hole mounting locations, and reduced the possibility of the plate flexing when tightened down due to the underside surface being out-of-flat.

Backiron Housing and Magnet Assembly

The AFSM actuators were assembled as a ‘sandwich’ type structure, with actuator halves on either side and the backiron, magnet, and backiron housing in the middle. This middle assembly was arguably the most critical of the AFSM features, because its thickness and parallelism directly determined the accuracy of the armature airgaps. Because of this, I deliberately machined the backiron and housing oversized, and then assembled them together with the magnets and ground all four assemblies at the same time using a surface grinder[1]. By doing this I assured that the critical features of all four assemblies were matched exactly. This approach paid dividends later on, as I measured the output of the elevation and azimuth axes to be nearly identical and decoupled–which would not have occurred had the airgaps not been equally spaced.

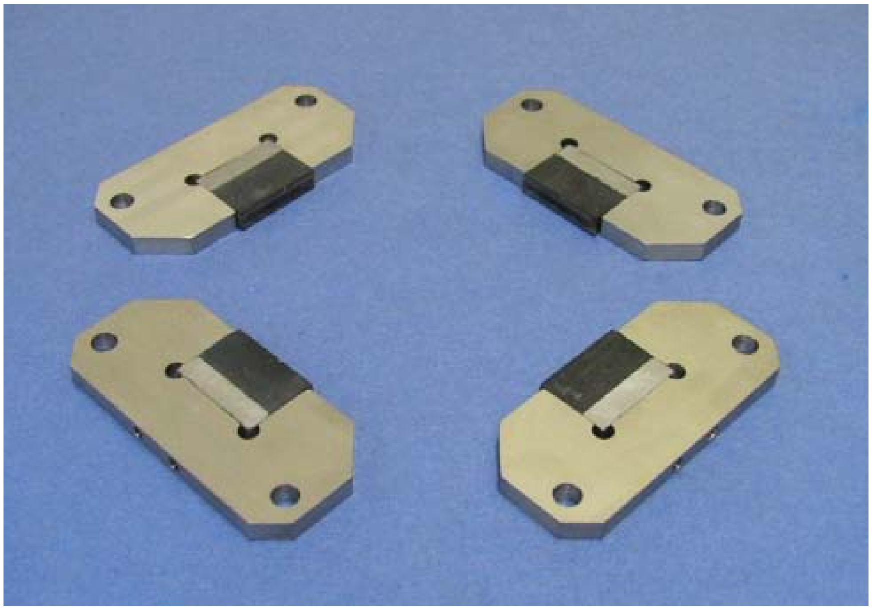

Figure 4-4: The AFSM backiron assemblies prior to final grinding.

Mirror and Armature Assembly

I intentionally designed the mirror and armatures to be as simple as possible in order to achieve symmetry in both geometry and mass properties. Thus, the individual

---

[1] Note that the thickness of the middle assembly was ground to a final dimension of 4.600 mm, but the magnet thickness as-received was only 4 mm. To build up the magnet thickness, I applied a layer of potting compound to each side prior to the grinding operation.

fabrication of these parts was straightforward. One issue I faced, however, was how to assemble the armatures to the mirror. Ironically, the high force density of the normal force actuator means that transmitting the force to the outside world is somewhat of a challenge. This is because the actuator is working magnetically on three sides of the armature, and therefore these sides must be left open and accessible. Because of the long narrow armature design in the AFSM, two of the remaining sides (the ends) were of little use for mounting purposes. This left only one long side available for assembly.

Since I was trying to minimize inertia, I did not want to use additional mechanical fasteners to attach the armature. Therefore, the only option was to bond the parts in some way. Welding is not viable, since the metals are dissimilar (i.e. aluminum and iron cannot be welded). Brazing is another possibility, but I could not find a standard brazing process that would accommodate the SM-2 alloy. Such a process would require a dedicated development project, and I did not have money or time to do so.

The remaining recourse was to use a structural adhesive to bond the parts together. I had never used an adhesive for such a critical application before, and therefore I researched the issue quite extensively. I settled on a two-part epoxy product from 3M, Scotch-Weld$^{\text{TM}}$DP460. Its shear and peel strengths are among the highest of any commercially available epoxy, and it is particularly suited to bonding aluminum and other metals. Even so, at a maximum shear strength of only 4500 psi, the bond strength is significantly inferior to a metal-to-metal joint (for example, the shear strength of aluminum is about 25,000 psi).

In addition, the strength of epoxy adhesives is highly dependent on the surface finish of the mating parts, as well as their cleanliness. Strength is also maximized when the thickness of the adhesive bond is established at a specific value (for DP460, this was 0.005 to 0.008 inches). I took pains in assembling the parts to clean them thoroughly and abraded them slightly to achieve the desired finish. To achieve the proper bond line thickness, I used fine wires of appropriate diameter to separate the parts during the bond cure. I also made a fixture to align the armature to the mirror during bonding, as shown in Figure 4-5.

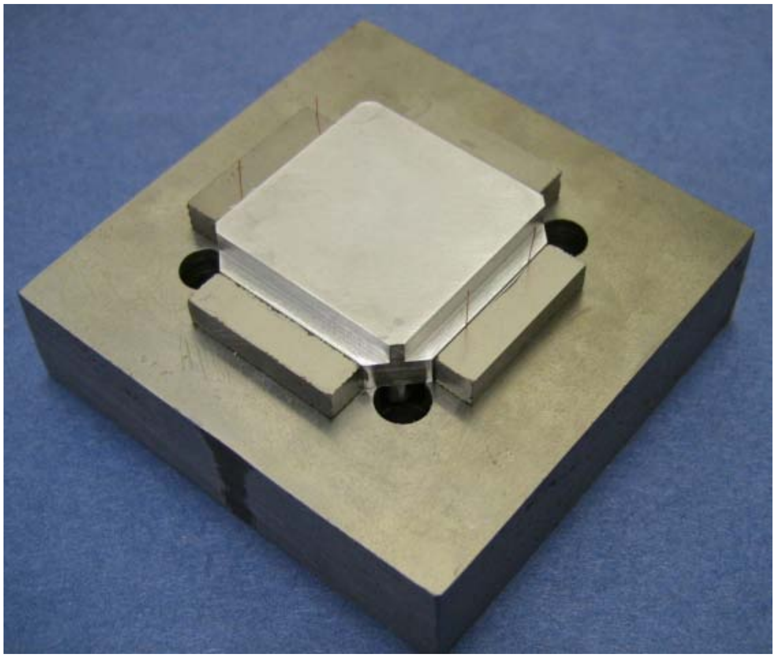

Figure 4-5: The mirror and armature assembly during epoxy bonding of the armatures to the mirror. Note the fine wires used to establish the optimal bond thickness.

After allowing the bond to fully cure, I lapped the upper and lower surfaces of each armature flush with the mirror mating surface, and applied several coats of polyurethane to absorb the energy of impact in the event of instability. The final step was to bond the elastomer bearings to the armature edges. For this, I used Loctite Black Max®380, a cyanoacrylate-based adhesive specially formulated for rubber bonding.

Axial Flexure

I discussed the axial flexure in Chapter 2, but I return to it here from a fabrication and assembly standpoint. One objective for designing the part was that I did not want it to dictate the precise axial positioning of the mirror. Instead, I intended to use the armature airgaps, referenced to the lower actuator halves (which were in turn referenced to the base plate), to set the axial position, and then bring the axial flexure *to* this position and lock it in place. To accomplish this, I created a threaded feature which allowed the flexure to move axially relative to the base plate, and made the ‘head’ of the flexure pilot into a mating feature in the mirror, which was then secured with an epoxy bond. Once positioned, I used thread locking compound and a jam nut to lock the lower portion of the flexure into the base. This methodology allowed the axial flexure to effectively perform its kinematic constraint function without influencing the assembled position of the mirror. See Figure 4-6.

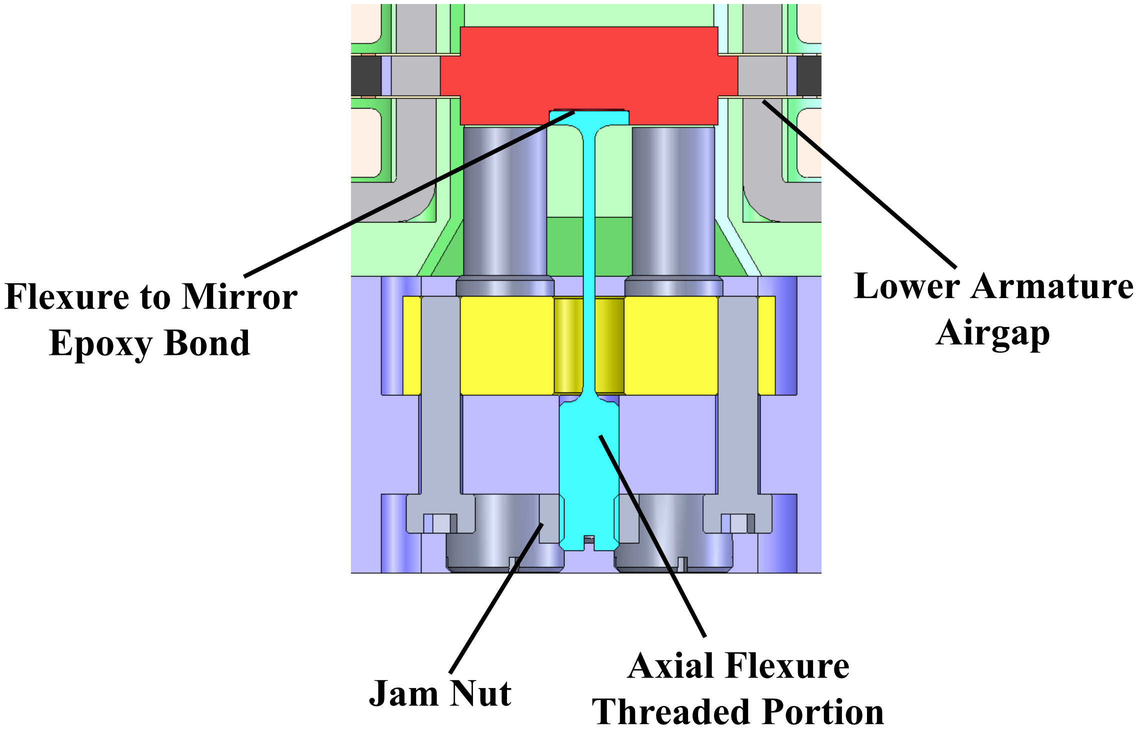

Figure 4-6: Illustration showing the axial flexure assembly technique. The lower armature airgap sets the mirror position, and the flexure is adjusted axially to accommodate it.

As mentioned in Chapter 2, fabrication of the axial flexure was challenging due to the small size of the part and the requirement for precise dimensions and surface finishes. I made the part manually on a lathe, turning down the large features first and then gradually reducing the diameter of the actual flexure portion. Turning the fillet radii on each end was a particularly delicate operation. Gerry Wentworth ground a special single-point tool for me to use for this purpose, and I used it to finish the entire flexure portion, thereby ensuring smooth transition between the straight features and the fillets. The last step was to polish the flexure portion to ensure that the surface was free of minute imperfections that would induce stress concentrations.

Since the flexure is made of a precipitation-hardening steel, heat treatment of the part is necessary before putting it into service. For the axial flexure, I chose the condition H900 heat treatment. My father, Bob Kluk, works with many small machine shops and manufacturers, and was able to have one of them perform the heat treatment for me. This proved to be a considerable time saver, as most heat treatment facilities are reluctant to provide services for only one (tiny) part!

Sensor Housings and Clamp

One of the problems I encountered during the layout of the AFSM design was how to position and clamp the capacitance probes. Since they were my primary means of obtaining mechanical feedback of the mirror angular position, it was critical to first position them precisely, and then to lock them rigidly in place to eliminate measurement error during testing. What was needed was a way to decouple the positioning and locking features, such that the act of locking the sensor would not affect its position. In addition, I wished to align the sensors laterally relative to the mirror position rather than the threaded holes in the base. Doing this would ensure proper placement of the sensors relative to the mirror center, thereby improving the resulting probe angular measurements.

To accomplish these goals, I designed a floating sensor clamp based on the familiar split-hoop method for clamping a cylindrical object. I reamed the clamp bores to achieve a tight slip fit with the sensor body, and used the central hole in the clamp to align it precisely to the axial flexure hole in the base[2]. Once aligned, I bolted the clamp in place with four screws, and then installed the cap probes.

Unfortunately for this purpose, the capacitance probes are cylindrical, and have no features useful for axial positioning. To remedy this problem, I machined hollow externally-threaded housings, which were designed to be slipped over the lower portion of the probes and secured in place with epoxy. The thread on the housings was a

---

[2] In the actual assembly, I used the axial flexure to position the mirror laterally before clamping the mirror in place. Therefore, piloting the clamp on the flexure mounting hole aligned the clamp with the center of the mirror.

special extra fine thread with a 1 mm pitch in order to allow small axial adjustments. Since I wanted to ensure that the housings were precisely concentric with the probe bodies to prevent binding, I made a fixture to hold the probes square, and the housings concentric to them during the epoxy cure time. The bonding fixture is shown in Figure 4-7.

Figure 4-7: The capacitance probes installed in the alignment fixture during bonding of the probe housings (threaded brass parts over the probes).

With the clamp aligned and the probe housings bonded, I installed the probes in the body, as shown in Figure 4-8. During testing, I found the adjustment process to be very easy. I simply loosened the clamp screw, rotated the probe body as needed with a screwdriver, and re-tightened the clamp. Although the probes did move slightly when tightened, the displacement was only a few microns, an error which was easily offset using the zeroing potentiometer on the probe conditioning electronics.

4.1.2 Actuator Coil and Core Assembly

The assembly of the magnetic cores, coils, and actuator housings into the actuator half assemblies was another key to attaining proper output force symmetry over the AFSM dynamic range. Such assembly work is something of an art form, especially for custom designs such as the AFSM.

Figure 4-8: Probe clamp with probes and axial flexure installed.

Fortunately, Professor Trumper was able to put me in touch with an experienced designer and builder of these assemblies, Fred Sommerhalter. In addition to performing this work professionally at Rockwell Automation's Anorad division, Fred also builds magnetic assemblies in his shop at home, both for his own personal projects and on a consultation basis for small custom work. It was in this latter capacity that I secured Fred's services.

I consulted with Fred during the early stages of the design, and we worked out several key details that made the actual assembly go smoothly. One such detail was the coil design itself. As mentioned in Chapter 3, I had to make some last-minute changes to the coil to accommodate the low-voltage, high-current power amplifiers that were available for the AFSM testing. Because I had designed for the opposite case (high-voltage, low-current), my original coil consisted of 100 turns of #22 AWG wire. To adjust for high current values, the most direct solution is to simply use fewer turns of a smaller gauge (i.e. larger diameter) wire. However, as the wire gauge decreases the packing factor also drops, which in turn results in a larger coil. Given that the actuator housings were already designed and constructed, I did not have room to accommodate such a coil.

Instead, Fred and I settled on a somewhat unusual design change. To achieve the equivalent cross-sectional area of a larger gauge size yet keep the packing factor small, we decided to wind ten individual strands of #22 AWG wire in *parallel*, each with ten turns. In doing so, the coil geometry, number of amp-turns, and power consumption remained the same, yet the current capacity was increased tenfold while the terminal voltage dropped by a factor of ten. As an added benefit, the small wire diameter had the advantage of mitigating resistive losses at high frequencies due to the skin effect, as discussed in Chapter 3.

Despite solving the coil geometry problem, the new ten-parallel-strand design added an additional challenge to the winding and forming of the coils. Fred developed special tooling for this purpose, as shown in Figure 4-9.

To keep the coil intact after winding and forming, Fred decided to make use of bondable magnet wire. This is a special type of wire which has a thin coating of a

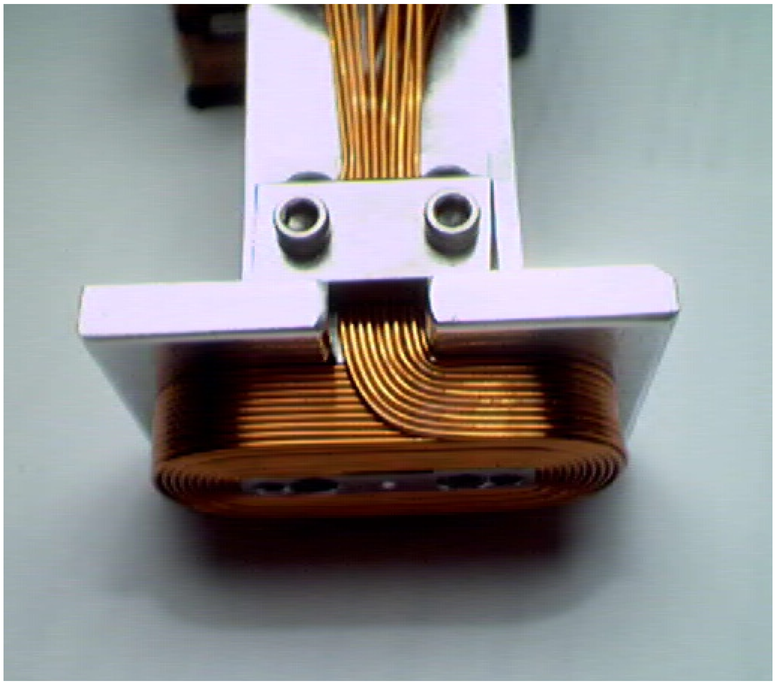

Figure 4-9: Stainless steel mandrel and forming tools used to wind the AFSM coils, shown with a coil installed. Photo courtesy of Fred Sommerhalter.

thermally-activated adhesive over the normal wire insulation. When the formed coil is heated (either by external means such as an oven, or simply by passing a large current through the coil itself), the adhesive flows between the wire strands and bonds them together. The coil then retains its shape after removal from the mandrel, as shown in Figure 4-10.

The next step was to install the core assemblies into the actuator housings, place the finished coils around the cores, and pot the components in place. I shipped the cores from MK Magnetics and the actuator housings made at MIT to Fred, and he performed the assembly and potting. See Figure 4-11.

After curing of the potting compound, the final and most critical activity was to grind the upper and lower surfaces of the finished assemblies to precise flatness, thickness, and parallelism. As mentioned, the actuator halves established the datum plane off of which the armature airgaps were set. In addition, careful grinding with proper tooling was necessary to avoid smearing the core laminations, which would defeat their design purpose by allowing eddy currents to flow. Fred subcontracted this

Figure 4-10: A finished coil assembly after forming and bonding. Photo courtesy of Fred Sommerhalter.

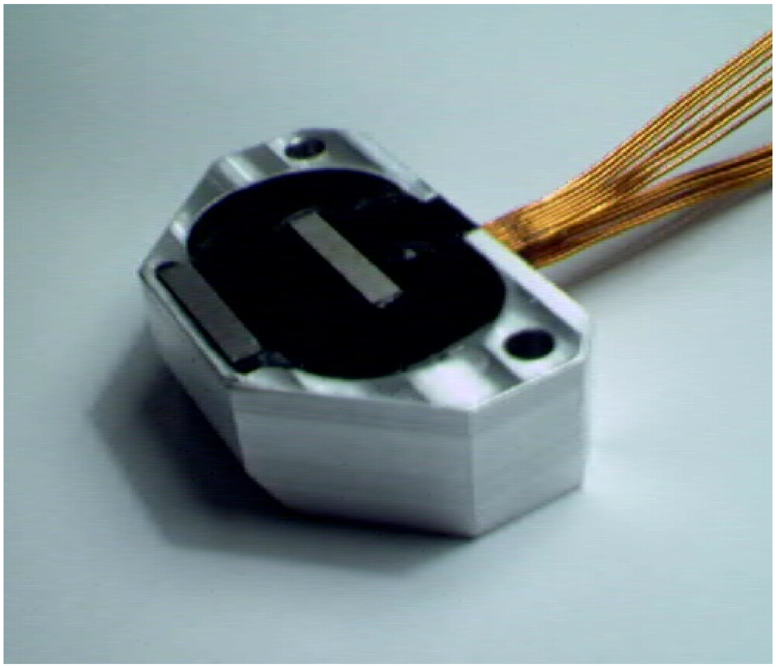

Figure 4-11: An assembled actuator half filled with potting compound. Photo courtesy of Fred Sommerhalter.

work to a trusted machine shop, and their machinist performed the grinding operation flawlessly. The complete set of actuator halves installed on a special grinding fixture is shown in Figure 4-12.

With the grinding complete, Fred packaged and sent the completed assemblies back to MIT. They are shown as-received in Figure 4-13. Note that I ordered twelve assemblies, but only used eight in the actual AFSM hardware. Of the remaining four, I used two as a test inductor load for the power amplifier diagnostics, and two as spares.

4.1.3 AFSM Final Assembly

With a full set of complete parts and subassemblies (see Figure 4-14), I was ready to begin the final assembly of the AFSM. Since most of the machine precision was built into the components themselves, I was able to complete the process in a timely manner, with minimal use of fixtures and external alignment techniques.

The first step was to align the capacitance probe clamp to the base and lock it in place. To attain the best possible alignment, I custom machined a brass alignment pin to match the diameters of the axial flexure hole in the base and the clamp center through-hole. I matched the clamp and base using the alignment pin (see Figure 4-15 and secured the clamp in place with four bolts.

Next, I measured the thickness of the actuator halves to ensure consistency from part to part. I selected the four most closely-matched assemblies and used them as the lower actuator halves to establish the datum reference surface for the mirror armature airgaps. After bolting them in place, I made four matched shims, each with a thickness of 300 micrometers, and placed them over the lower core poles. This is shown in Figure 4-16.

I installed the mirror over the shims and test fit the axial flexure to ensure that it threaded cleanly into the mirror counterbore with zero lateral displacement. Once I was satisfied with the alignment, I coated the mirror counterbore with a small amount of structural epoxy (here I used the same 3M product used for the armature-to-mirror bond) and secured it in place with a simple hold-down plate, as shown in Figure 4-17.

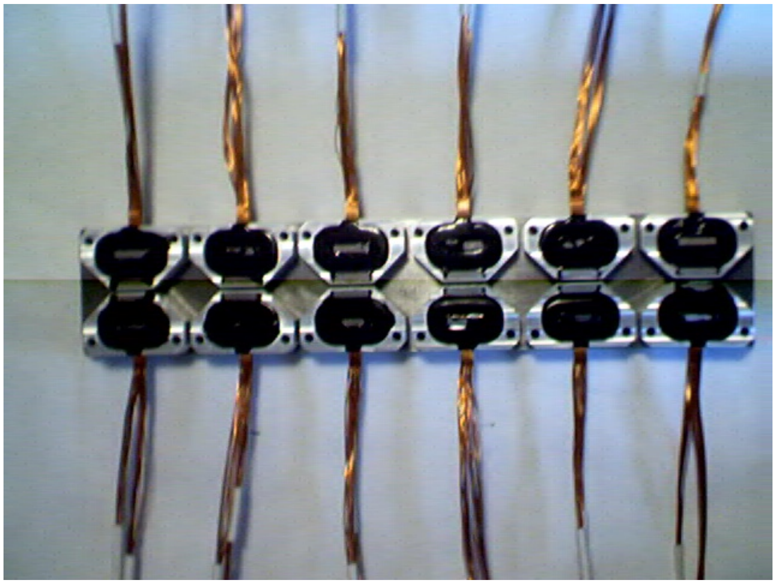

Figure 4-12: The potted actuator halves installed on a fixture plate ready for grinding. Photo courtesy of Fred Sommerhalter.

Figure 4-13: Finished actuator halves after final grind.

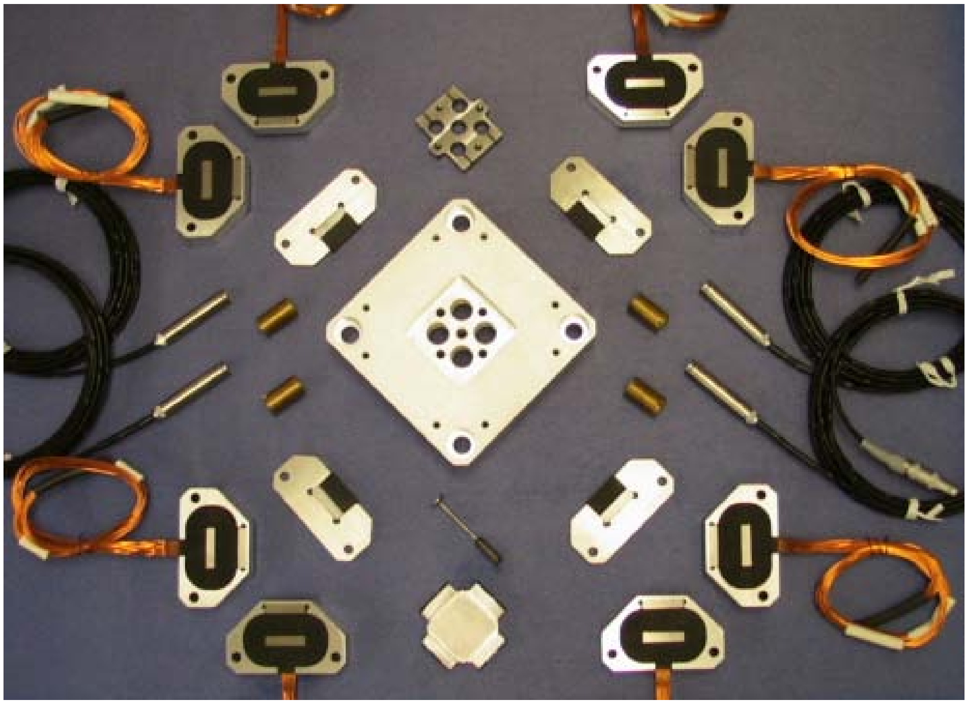

Figure 4-14: The complete set of fabricated AFSM components, prior to final assembly.

Figure 4-15: The probe sensor clamp installed in the base plate using an alignment pin (brass part in the center). Note the access holes in the base plate for the clamp screws at the bottom of the figure.

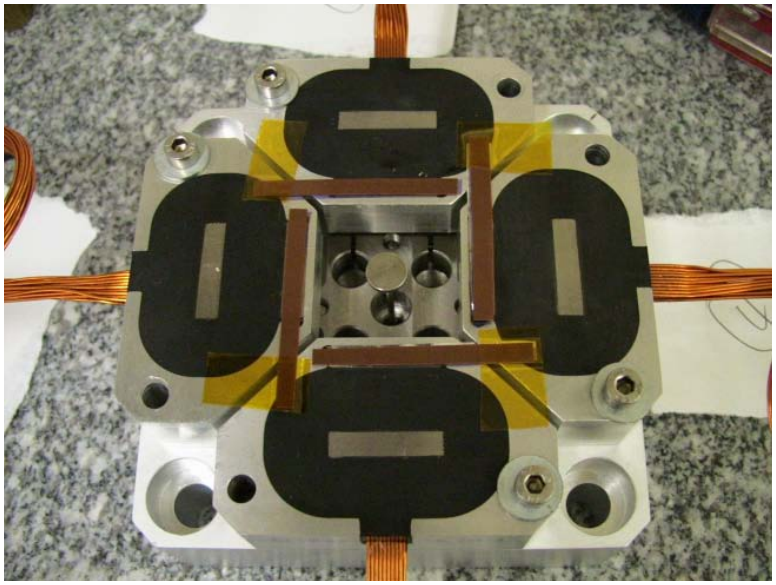

Figure 4-16: Thin shims in place over the lower actuator core poles in preparation for mirror assembly. The lower actuator halves are temporarily bolted to the AFSM base.

I then threaded the axial flexure into the mirror from the underside, and secured it into the base using thread locking compound and a jam nut (see Figure 4-6). I left the assembly in place overnight while the adhesives cured.

With the armature neutral location established, I bonded the magnet faces to the elastomeric bearings using the same cyanoacrylate adhesive I had used for the bearing-to-armature bond. The nice feature here was the the permanent magnets were more or less self-aligning, as they were attracted to the both the armature and the rear core poles.

With the magnets installed, the simple remaining steps were to install the preground backiron housings over the magnets, remove the airgap shims, and install the upper actuator halves. This completed the base assembly (Figure 4-19). To orient the assembly for optical bench testing, I installed it on a pre-machined angle plate, and then threaded the capacitance probes into the rear of the assembly, as shown in Figures 4-20 and 4-21.

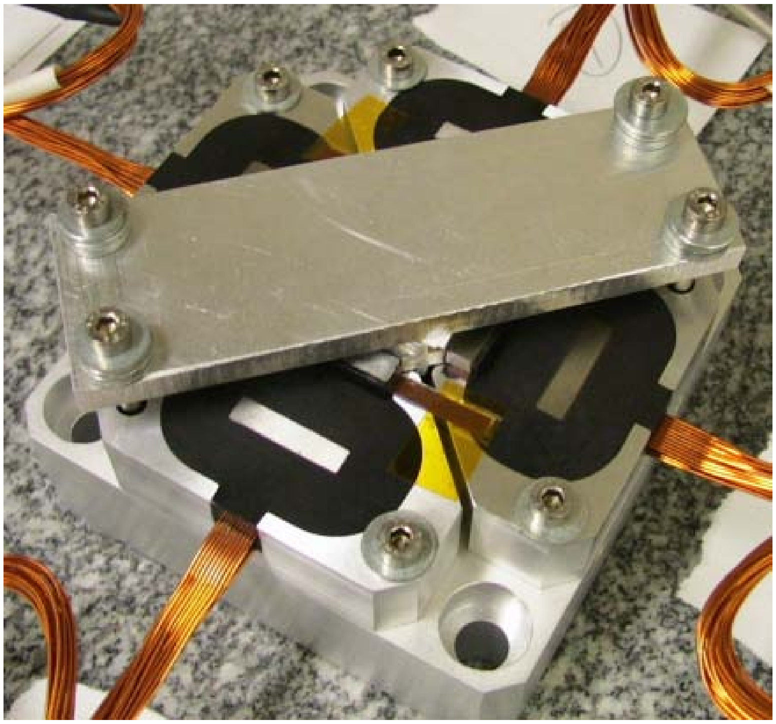

Figure 4-17: The mirror assembly clamped in place with a fixture plate prior to bonding of the axial flexure.

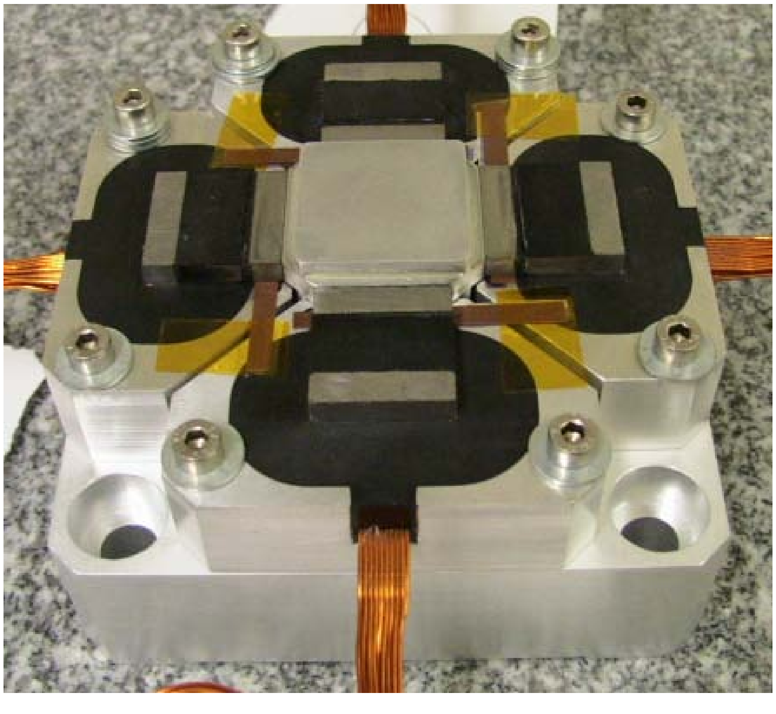

Figure 4-18: The magnet assemblies bonded to the elastomeric bearings and armatures.

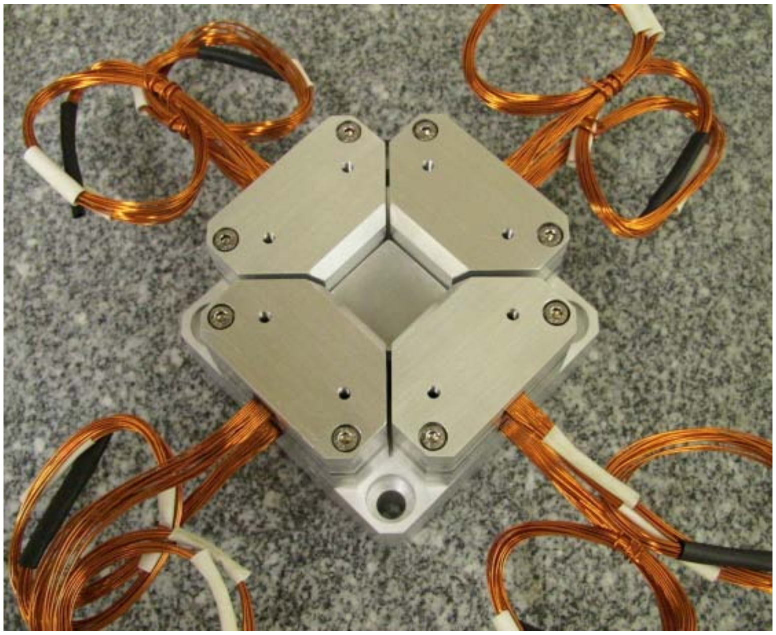

Figure 4-19: The finished AFSM base assembly.

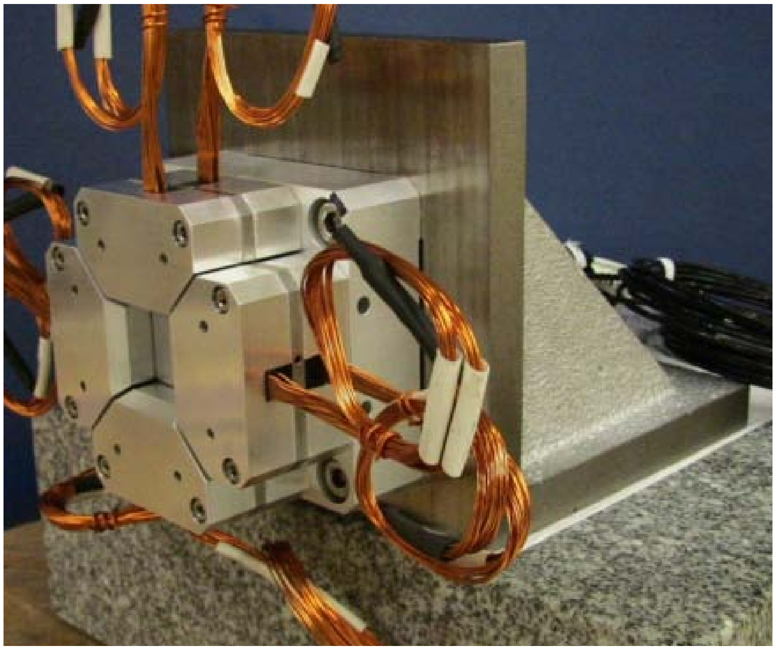

Figure 4-20: The base assembly installed on its angle plate.

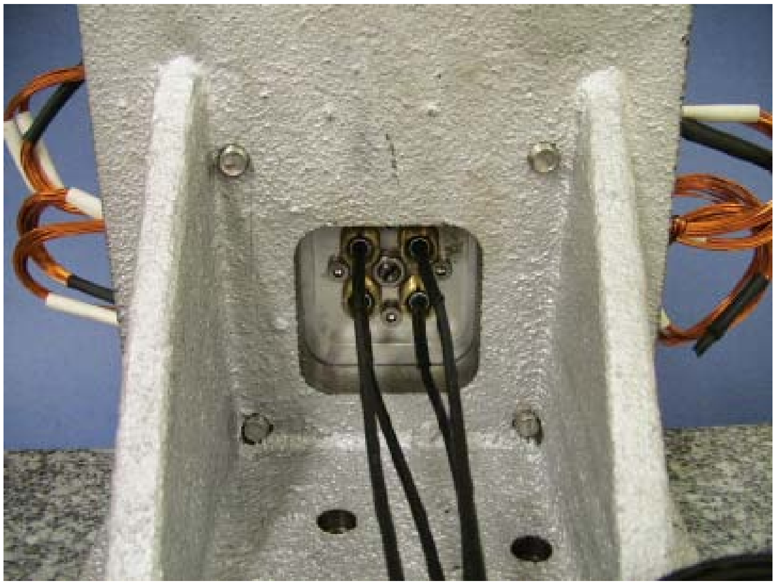

Figure 4-21: Rear view of the AFSM assembly showing the installed capacitance probes.

4.2 Power Electronics and Sensors

I now turn to discussion of the support systems which drive and sense the AFSM hardware. Since the controller for this system is to be implemented electronically, it is necessary to transduce the control output signals into high power current inputs to the AFSM electromagnetic actuators. This is the job of the power amplifiers and current compensation electronics. Similarly, the mechanical output from the AFSM mirror motions must be converted into an electrical signal in order to be used as feedback. For the AFSM, this is accomplished through the use of two different types of sensors: the mechanical capacitance probes, and the optical quad cell detector. Each sensor has a suite of electronics to condition, scale and rotate their signals into a form that the controller can process.

4.2.1 Power Amplifier Selection

With the AFSM coil design from Chapter 3 complete, I was able to calculate the electrical power requirements for driving it, and used these to select an appropriate set of power amplifiers, one for each actuator. Because of the inductive impedance load at high frequencies, I required an amplifier with high output power (at least 450 watts) and a high fidelity, linear frequency response from DC to at least 20 kHz.

A switching-type amplifier of the type commonly used for driving servomotors and other low-bandwidth hardware would not be appropriate for the AFSM. Most commercially available versions of this device have switching frequencies of 20-30 kHz. Since I wished to drive the AFSM to 5 kHz or greater, the phase loss incurred by the switching action would have been prohibitive. In addition, the mirror mechanical time constant is small enough that it can respond to the current ripple produced by switching at these frequencies. Moreover, switching amplifiers generate large levels of electrical noise at the switching frequency, which would cause unwelcome interference with the low-power control electronics. Therefore, using a linear amplifier was essential for this project.

It is important to keep in mind that although the 426 W peak power requirement for the actuators is large, the vast majority of this power is *reactive*; that is, power returned to the source during each electrical cycle. The actual *dissipative* power consumed by the actuators through electrical, magnetic, and mechanical losses is comparatively small at less than 10 watts[3]. Nevertheless, the power requirements still required the use of a heavy-duty amplifier.

As luck would have it, Jamie Burnside at Lincoln Laboratory knew of a set of power amplifiers that had been procured for a prior program and were no longer in use. The set consisted of three Techron model 7541 linear amplifiers delivering a maximum of 559 watts RMS power each, which was more than enough to suit my needs for the AFSM. The specifications for the 7541 are given below.

Only three 7541s were available, however, which left me in need of a fourth amplifier. Since the 7541s were no longer being manufactured by Techron, I procured a newer close relative, the Techron model 7560. This amplifier had more than twice the output power (and weighed about twice as much) as the 7541. Its specifications are given in the following table.

---

[3] Dissipative and reactive power are most conveniently represented mathematically in complex phasor notation, and therefore are commonly referred to as ‘real’ and ‘imaginary’ power, respectively.

Table 4.1: Techron 7541 Power Amplifier Characteristics

Property	**Value**
Maximum Output Power (RMS)	559 watts
RMS Output Voltage	33 volts
RMS Output Current	17 amps
Slew Rate	21 V/μsec
Voltage Bandwidth (2 ohm load)	25 kHz

Table 4.2: Techron 7560 Power Amplifier Characteristics

Property	**Value**
Maximum Output Power (RMS)	1350 watts
RMS Output Voltage	64 volts
RMS Output Current	16 amps
Slew Rate	36 V/μsec
Voltage Bandwidth (4 ohm load)	45 kHz

The two power amplifier models are shown in Figure 4-22. Only one 7541 can be seen in this figure; the other two were mounted at the rear of the electronics rack.

4.2.2 Power Amplifier Current Compensators and Signal Inverter

The default operational mode of the power amplifiers is voltage signal amplification. Since the force output of the AFSM actuators is proportional to current, I wished to control this variable instead. Although the model 7541 has a selectable current control mode, the 7560 does not. Furthermore, the 7541 documentation states that the internal current control mode decreases the available amplifier bandwidth, depending on load. For these reasons, I decided to design my own current compensation for the amplifiers.

I sensed the AFSM coil current by placing a 0.5-ohm power resistor (actually, two 1-ohm resistors wired in parallel) in series with the actuator coil. I then measured the voltage across the resistor using a differential amplifier, which I used as feedback

Figure 4-22: Model 7541 (top) and 7560 (bottom) power amplifiers mounted in their racks during AFSM testing. Note also the shield terminations on the four cables at the output terminal block at middle-right.

for the current compensator circuitry. The sense resistor configuration is shown in Figure 4-23.

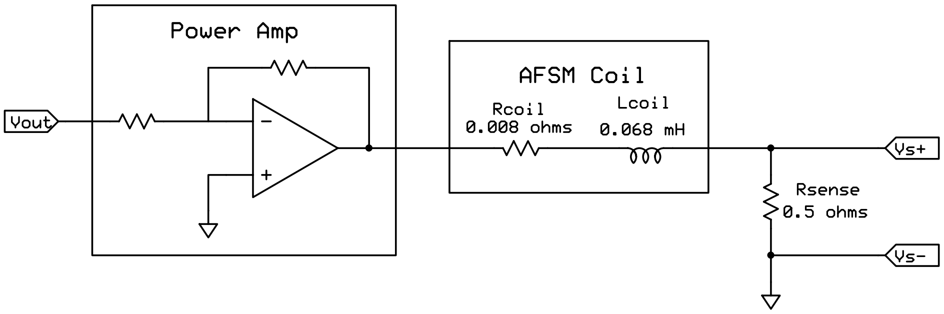

Figure 4-23: Power portion of the control loop showing the power amplifier and sense resistor in series with the AFSM actuator (modeled as a resistor in series with an inductor).

To determine the appropriate compensation, I first measured the open loop frequency responses of the amplifiers under the load of the AFSM coils and power resistors. Using a pair of spare actuator halves set to appropriate airgaps, I was able to simulate the impedance load of the AFSM. The resulting frequency response is shown in Figure 4-24. With this plot I was also able to obtain the inductance of the actuator by using the first-order breakpoint frequency. I found that the experimental time constant result was within 1% of that predicted by the coil design values of 67.9 μH and 8.7 milliohms.

Current Compensator Circuit

With the frequency responses known, I designed the controller. The first order inductor breakpoint in the frequency mid-band allowed me to use a simple proportional controller and cross over in the 20-40 kHz range. I designed the analog circuit shown in Figure 4-25 to perform the current compensation[4]. It consists of an input preamplifier for command scaling, a differential amplifier (INA117) for the current sense feedback, and a simple offset circuit using two voltage regulators as input[5]. These

---

[4] In this and subsequent schematics shown in this thesis, I have omitted the operational amplifier power supply connections for clarity. In each case, the op-amps are powered with ±15 Vdc and decoupled through 0.1 μF capacitors.

[5] Mike Boulet added the zero offset portion of this circuit in a later revision.

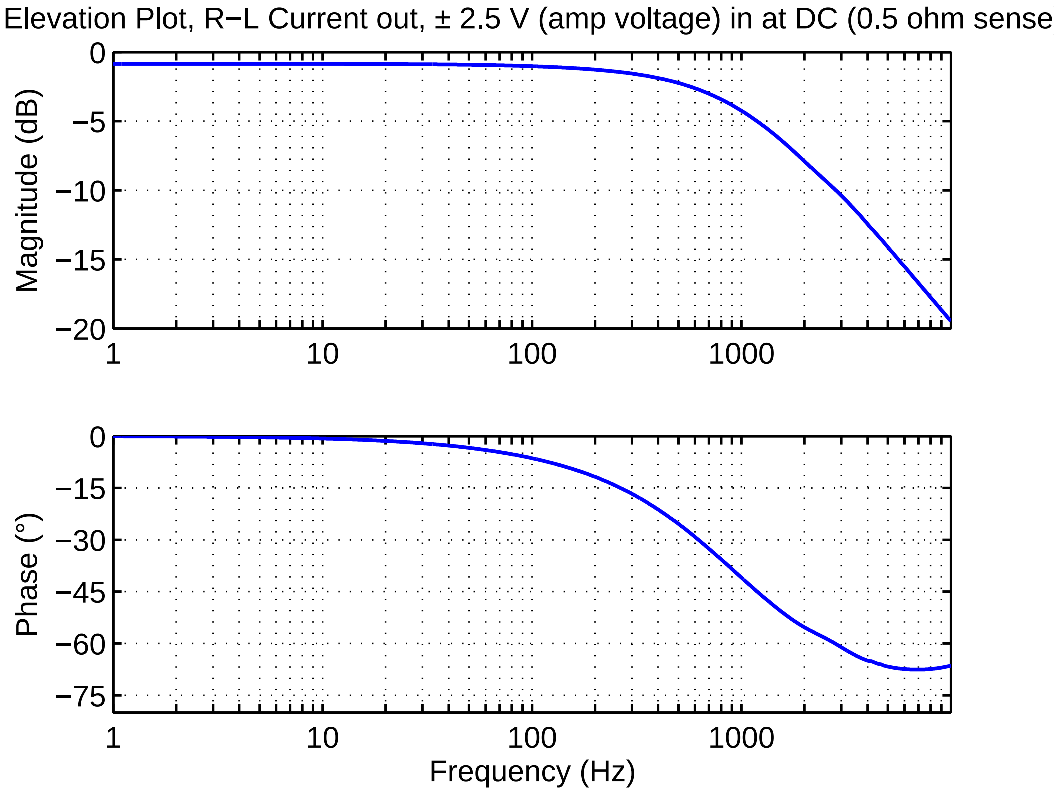

Figure 4-24: Frequency response of the AFSM actuator and series sense resistor.

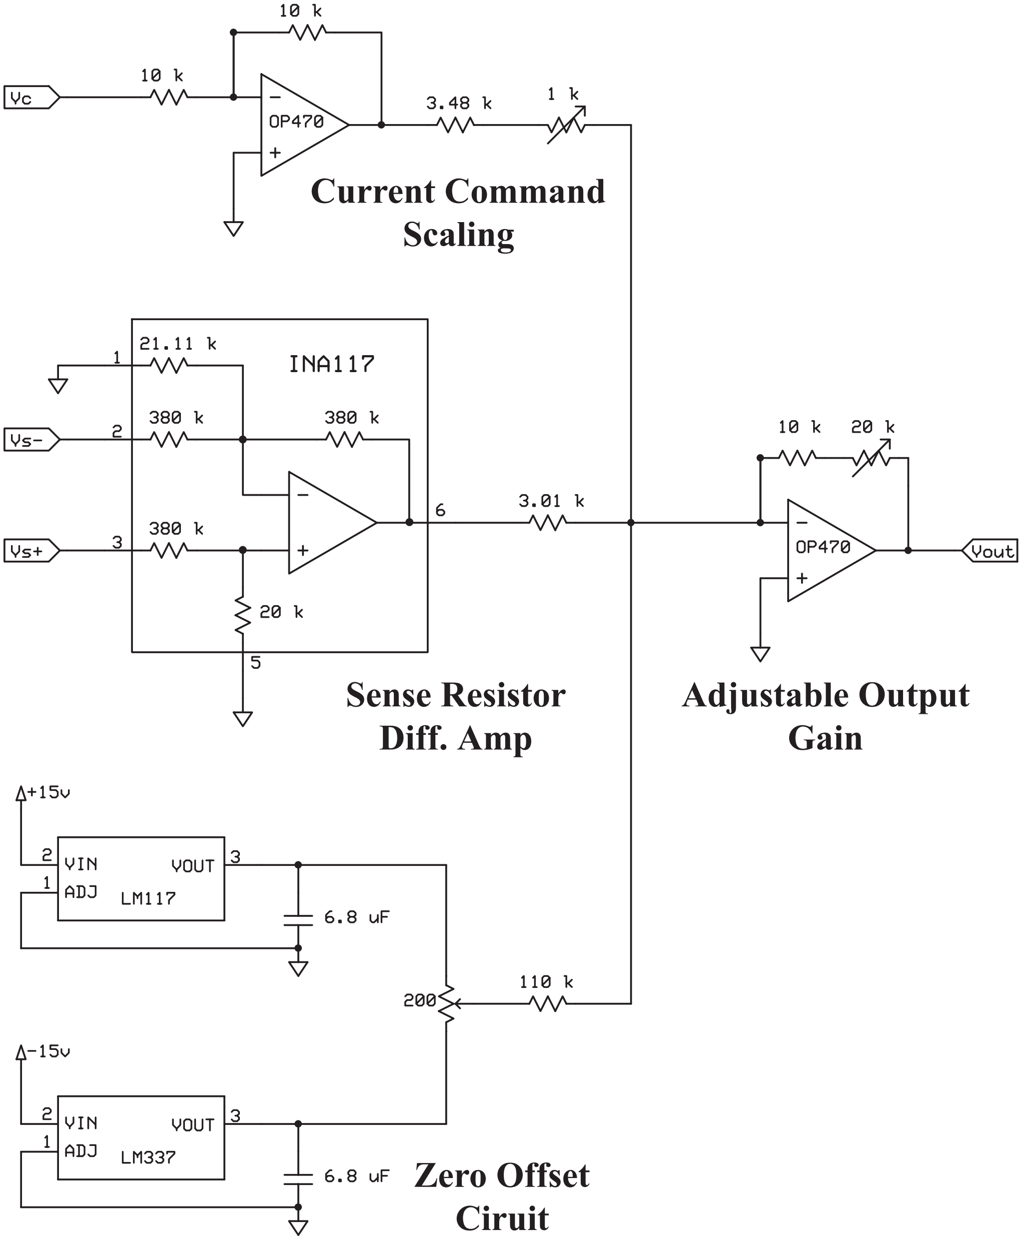

Figure 4-25: Current compensator circuit for the power amplifiers (one circuit per amplifier).

three modules sum into the right-hand op-amp circuit, which implements the proportional gain. The circuit output is fed to the power amplifier inputs through the inverter circuit described in the next section.

In my initial test runs at low frequency the current compensation loop worked well. However, I discovered while performing tests at higher frequency that when combined with the phase roll-off of the amplifier, the current loop displayed a large resonant peak at about 40 kHz, as shown in Figure 4-26. This in turn caused a peak to occur in the AFSM position loop. I later reduced the current loop gain to eliminate this problem, at the expense of some current loop bandwidth.

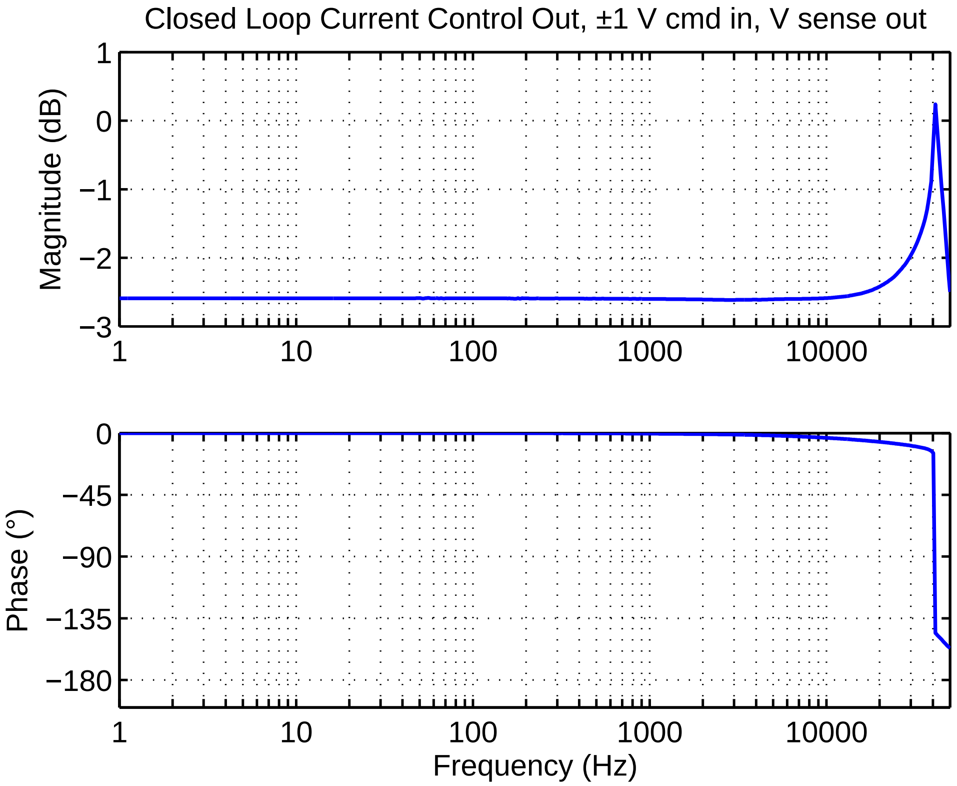

Figure 4-26: Closed-loop current response showing resonant peak at 40 kHz.

Inverter Circuit

Since each AFSM coil operated from its own independent amplifier, I needed to develop coordinated 'push / pull' actions between pairs of actuators using command signals. Since I anticipated some coupling effects between the azimuth and elevation axes, I wanted to retain the ability to independently control the four actuators electronically. Doing so would allow for MIMO control in three degrees of freedom (the two original mirror rotations plus axial motion) if necessary. However, for my initial testing this was not important, so I designed a simple circuit to split the command signal for a single axis into inverted and unmodified outputs. When sent to opposing pairs of actuators, the desired push / pull behavior was obtained. The circuit is shown in Figure 4-27.

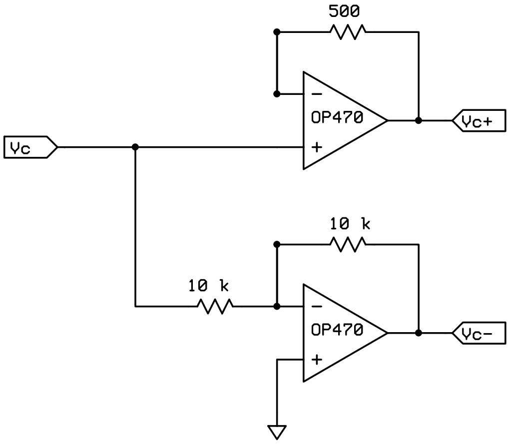

Figure 4-27: Command signal inversion circuit (one circuit per actuator pair).

In later stages of testing, I found that there were no significant cross-axis coupling effects, and as such there was no need for independent control over each actuator. I therefore retained the inversion circuitry in the final electronics board.

4.2.3 Capacitance Probes

To directly sense the mechanical output of the AFSM, I needed a high bandwidth, non-contact device with enough range to cover the entire mirror motion, but with high resolution in order to hold position to micro-radian level precision. Several sensing technologies are available that meet these criteria, including eddy current sensors, optical encoders, capacitance probes, and laser interferometers. However neither optical encoders nor laser interferometers package well into the AFSM volume, and the latter has the added disadvantage of being very expensive. Eddy current sensors are relatively inexpensive and package well, but they rely on the same fundamental physics (i.e. magnetic fields) as the actuators, and therefore may be subject to false measurements due to interference.

Although moderately expensive, capacitance probes are the best choice for the AFSM. The cylindrical probe housings are slim enough to be installed behind the mirror, and the electric fields they use to sense displacement are not subject to interference from the actuators. In order to attain both the range and resolution requirements, I needed to select a probe with a relatively large sensing diameter of 5 mm. With guard and housing, the final probe diameter was 8 mm–just small enough to fit within one quadrant of the square mirror.

Although I could have used only three probes to obtain full angular displacement information in the two axes of interest, I decided to use four so that I could average the readings and thereby obtain a more accurate angular measurement. The averaging technique is also useful for rejecting the effects of modal excitation at one of the sensor locations, as described in [28]. Also, having four probes was useful for detecting any static or dynamic asymmetries in the mirror.

I purchased a set of four probes from ADE Corporation, together with signal conditioners and a gauging console. The probes prior to installation are shown in Figure 4-28; the electronics console is in Figure 4-29. The probe performance specifications are given in Table 4.3, below.

Table 4.3: ADE 2805 Capacitance Probe Characteristics

Property	**Value**
Range	±250 μm
Standoff	500 μm
Resolution	112 nm RMS
Bandwidth	10 kHz
Linearity	0.04%

4.2.4 Capacitance Probe Rotation Electronics

Of course, the capacitance probes produce a local linear measurement of the mirror, not the desired angular output. To manipulate the signal into a form useful for the controller, I designed an intermediate electronics module to rotate and scale the output of the four probes.

Before discussing the design, an understanding of the geometry involved is helpful.

Figure 4-28: The set of four capacitance probes prior to installation in the AFSM.

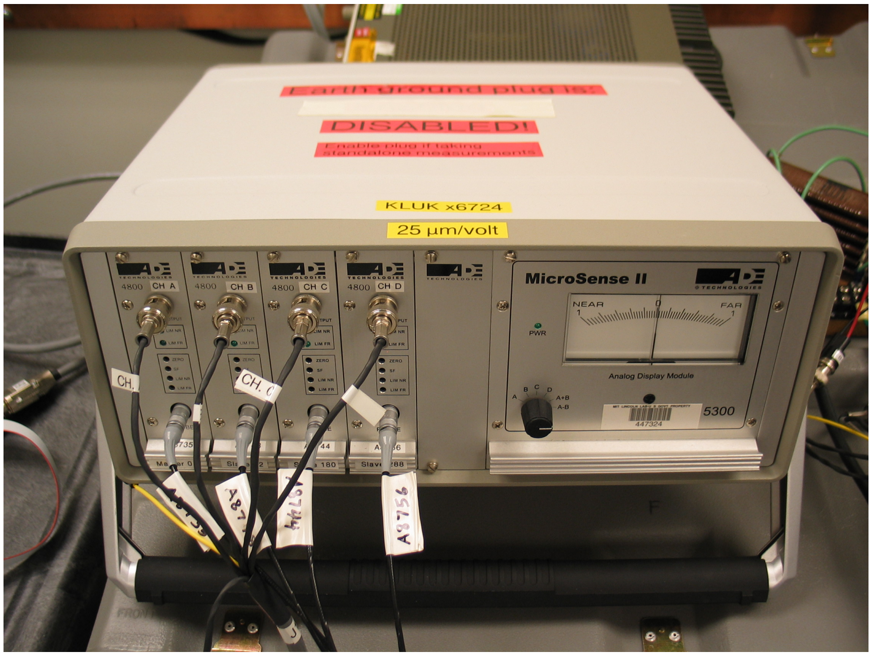

Figure 4-29: Capacitance probe conditioning electronics.

I labelled the four probes as A, B, C, and D, as shown in Figure 4-30. The mirror is depicted by the square in the foreground, and its front face is normal to the Y axis. The local displacement of the mirror is measured by the capacitance probes as the distance from the probe face to the rear mirror surface. The direction of positive displacement is that of increasing distance from the probe face.

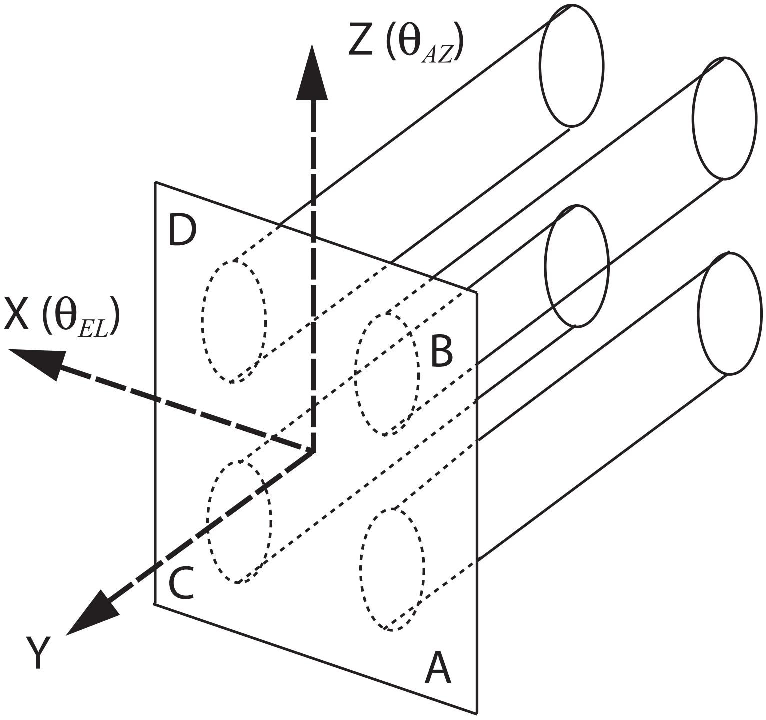

Figure 4-30: Capacitance probe labels and mirror coordinate definitions.

From the figure, it is clear that under ideal conditions (i.e. no probe error or misalignment), a pure mirror rotation about the positive Z axis (θ_{AZ}) will produce negative displacements in probes A and B, and positive displacements in probes C and D. Similarly, rotation about X (θ_{EL}) results in positive displacements in probes A and C, and negative displacements in probes B and D. This is illustrated in Figure 4-31.

Using Figure 4-31, we can easily derive the mirror rotations as a function of the probe output voltages. Making use of the small angle approximation, it is evident that assuming the mirror rotates about its center, $\theta \approx x/d$. However, we have four sensor measurements available, and we can take advantage of this by averaging the

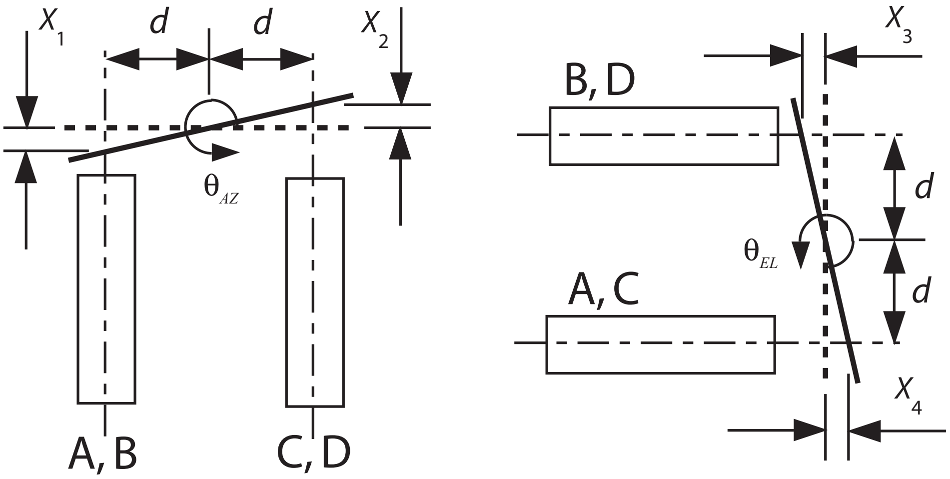

Figure 4-31: Probe displacements for θ_{AZ} and θ_{EL} mirror rotations.

measurements. This removes the constraint that the mirror rotate about its center, and also makes the measurement impervious to axial motions of the mirror. For the case of azimuthal rotation, the average displacement x_{AZ} is

$$x_{AZ} = \frac{x_2 - x_1}{2} = \frac{\frac{C+D}{2} - \frac{A+B}{2}}{2} = \frac{1}{4}(C + D - A - B) \tag{4.1}$$

where A, B, C, and D are the displacements of each probe. For rotation in elevation, we have

$$x_{EL} = \frac{x_4 - x_3}{2} = \frac{\frac{A+C}{2} - \frac{B+D}{2}}{2} = \frac{1}{4}(A + C - B - D) \tag{4.2}$$

All that remains is to divide the results by the distance d, and multiply by the sensitivity constant of the probes and a scaling factor for the control signal. In the AFSM design, the probe centers are located at a radial distance of 12 mm from the mirror center, 45 degrees off the rotation axes. Hence, $d = 12/\sqrt{2} = 8.485$ mm. Also, the probe sensitivity is 25 μm/volt. In addition, since I expected a range of ±10 mrad and wished to use full-scale voltage signals of ±10 volts, I applied a signal scaling factor of 1000 volts/rad to the result. With these constants, the final conversion relations for elevation and azimuth in terms of signal voltages were, respectively

$$V_{AZ} = 0.7366(V_C + V_D - V_A - V_B) \quad (4.3)$$

$$V_{EL} = 0.7366(V_A + V_C - V_B - V_D) \quad (4.4)$$

where V_A, V_B, V_C, and V_D are the voltage signal outputs from each probe.

I implemented the rotation math using the analog circuit shown in Figures 4-32 and 4-33. The first group of INA105 differential amplifiers at the left of Figure 4-32 provided signal input buffering, while the second group on the right performed initial subtraction of the signals. These two outputs are simply summed at the inverting output stage (see Figure 4-33) to produce the azimuth signal, but the $A - D$ portion must be inverted prior to the final summing operation. The scaling factor of 0.7366 is implemented using the 10k potentiometer at the output stages.

4.2.5 Optical Sensors and Electronics

Just before I stared work on the AFSM project, Larry Hawe in our lab was just finishing his thesis work [11], which explored advanced control methods for the Lincoln Laboratory heritage FSM, the HBSM-D. Upon completion of his work, Hawe's experimental setup became available for me to use. I therefore designed the AFSM dimensions to be compatible with the existing optical test setup, which included a helium-neon laser, attenuator, focusing lens, a quad-cell detector, and custom rotation electronics, as shown in Figure 4-34. Hawe describes each of these components in detail in his thesis [11], so I cover them only briefly here.

The quad-cell detector consists of four very small photodiodes arranged in a diamond pattern. Each photodiode produces a signal proportional to the light level impinging on its surface.

Electrically combining the four signals gives the position of a focused spot of light relative to the quad cell position. To obtain properly referenced and scaled azimuth and elevation position information, a separate electronics board is used to perform the conversion. This board is shown in Figure 4-35.

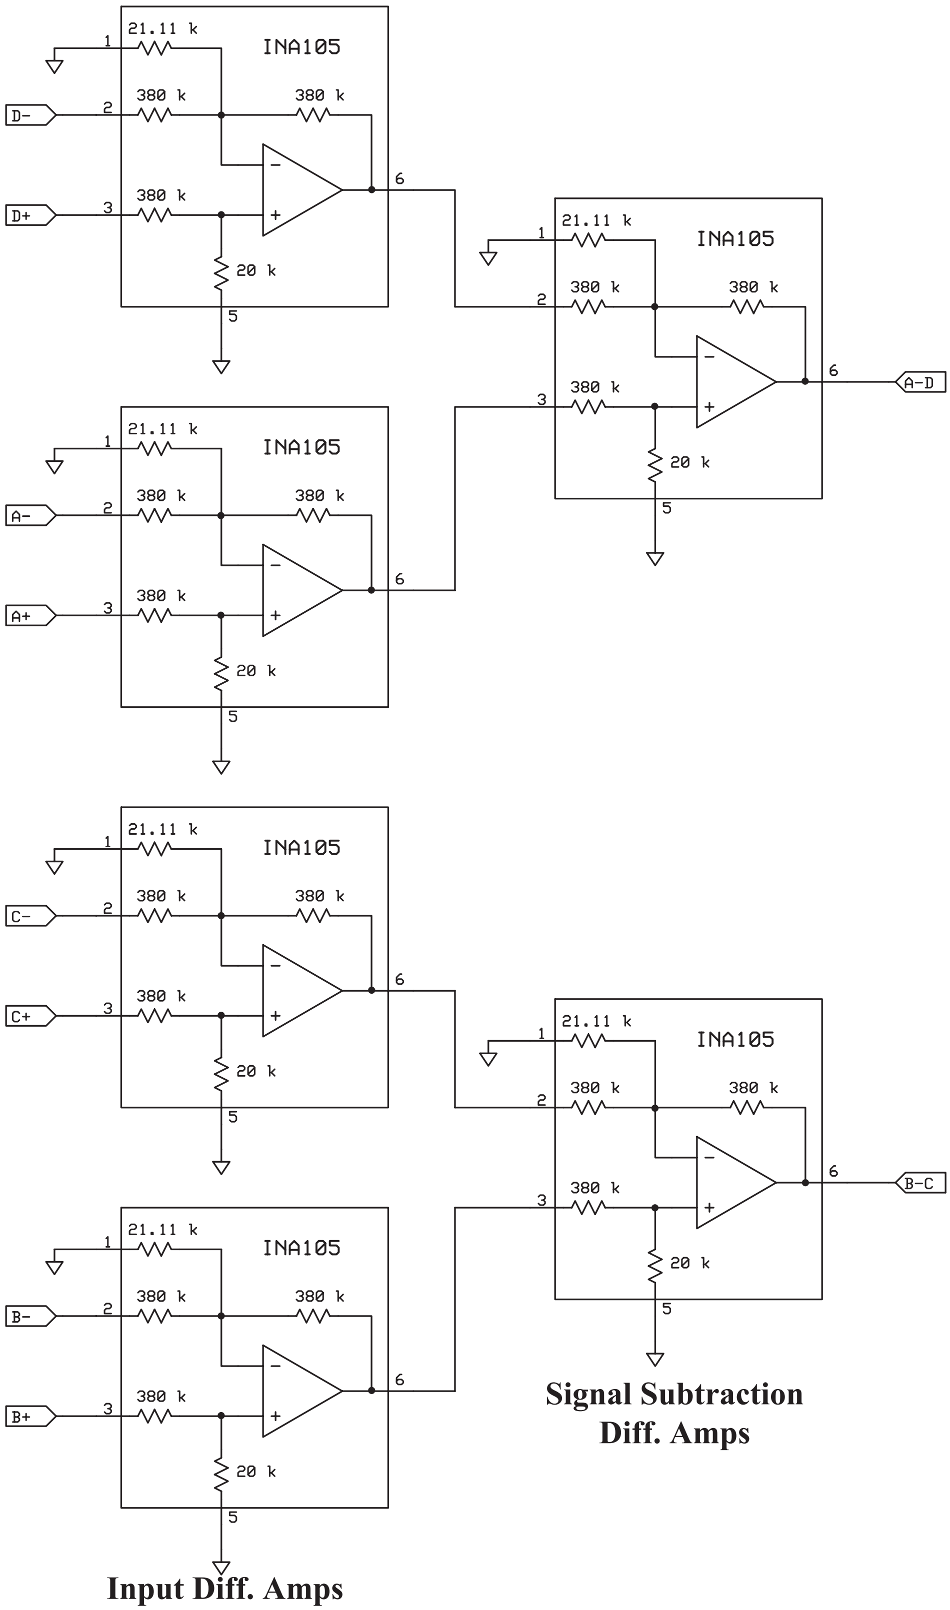

Figure 4-32: Input stage of the capacitance probe rotation electronics (signal buffering and subtraction).

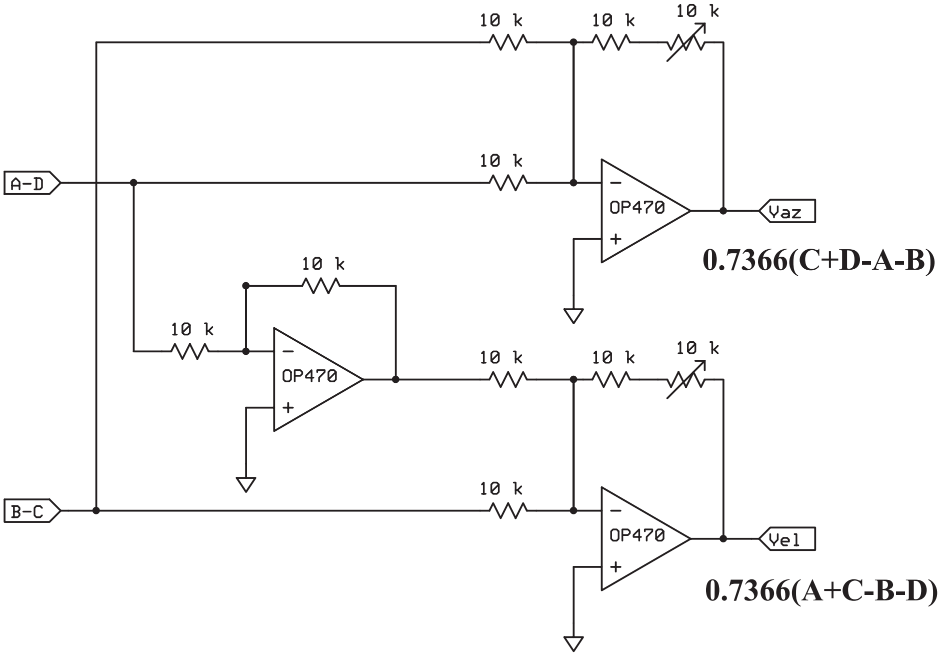

Figure 4-33: Output stage of the capacitance probe rotation electronics (final subtraction and scaling).

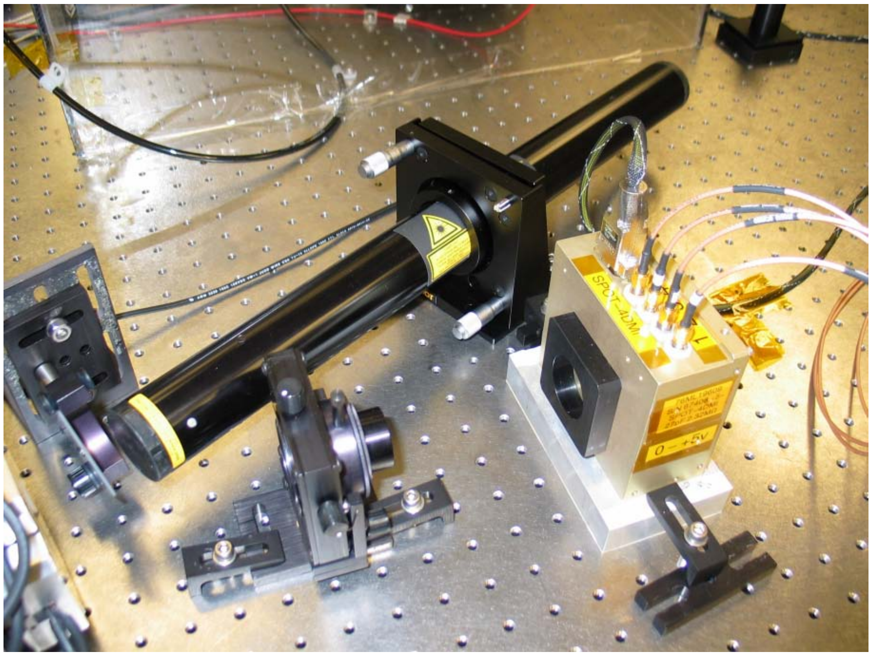

Figure 4-34: The hardware used for collecting optical feedback measurements for the AFSM.

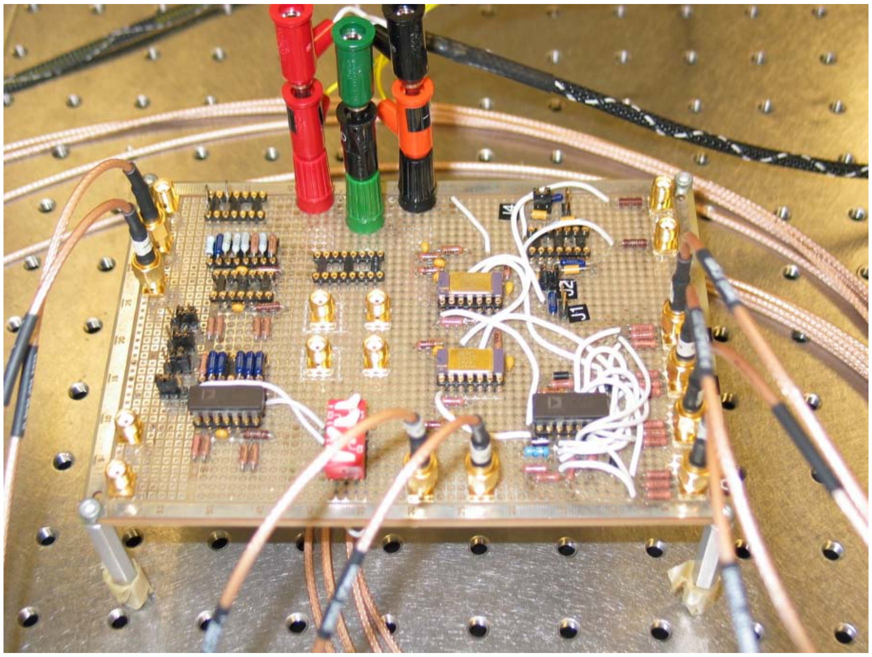

Figure 4-35: The quad cell detector rotation and scaling electronics.

One advantage of using the optical setup is that the angular resolution of the system is essentially infinitely variable. This is because the quad cell may be placed at any distance from the mirror (closer for coarse resolution; farther for fine). In addition the conditioning electronics have an exceptionally low signal-to-noise ratio and high sensitivity (the full range across the approximately 1 mm^2 active area was ±10 volts). With the distances used for my setup, I calibrated the sensitivity, using the capacitance probes as my baseline angle measurement, at approximately 25 microradians per volt.

With the AFSM assembled and all of the elements for driving and measuring it in place, I connected them all together to form the complete system, as drawn schematically in Figure 4-36. Of all the blocks in the figure, only the capacitance probe and optical controllers had yet to be designed and built. The controllers, as well as the system identification testing I performed in order to design them, are covered in the following chapter.

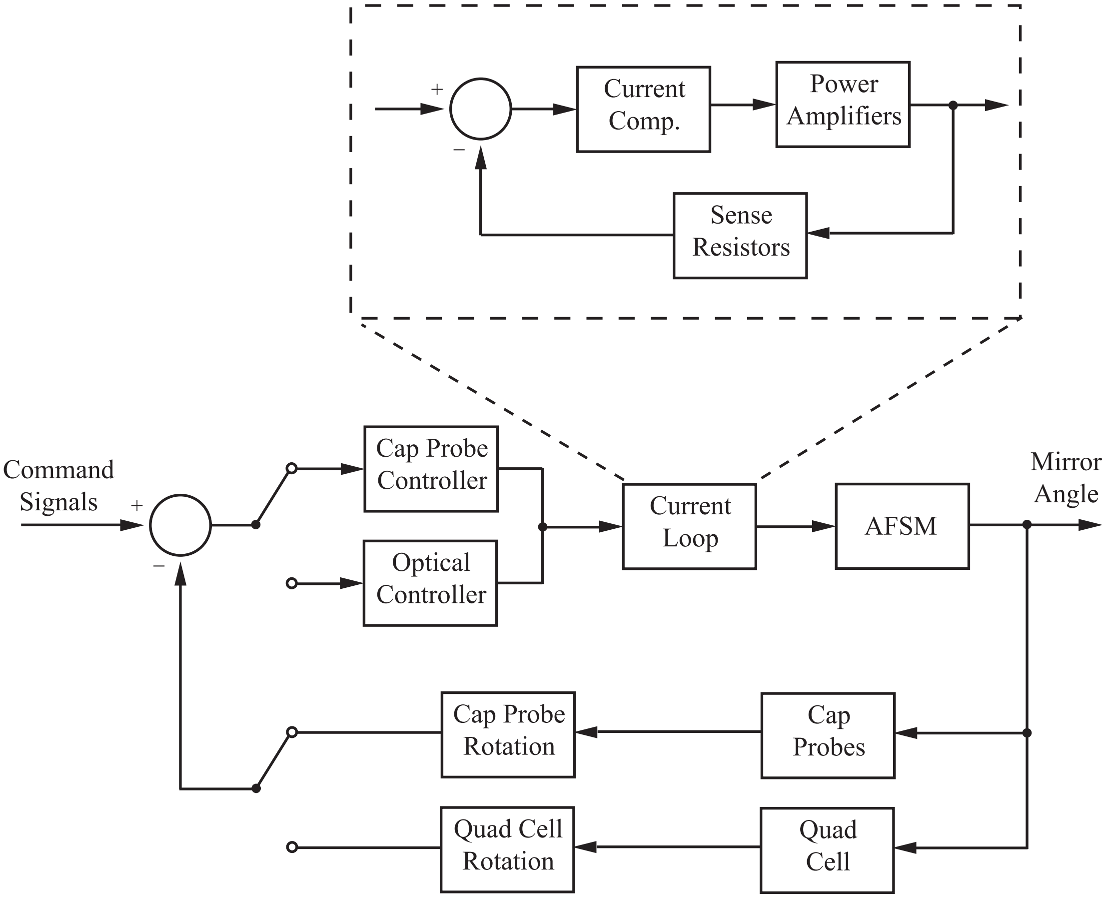

Figure 4-36: Block diagram of the AFSM system components.

Chapter 5

System Identification, Controller Design, and Experimental Results

In this chapter I describe the AFSM system identification and controller design. Since I used two different feedback measurements (capacitance probes and the optical quad cell), it was necessary to perform separate identification studies for both measurement types. This also made it necessary to design and build separate controllers for each type of feedback.

All dynamic measurements herein were taken using the HP 3562A dynamic signal analyzer (DSA)–an old but very powerful electronic measurement and analysis tool. Through many swept sine frequency response measurements with the DSA, I was able to completely characterize the critical details of the AFSM dynamics[1].

Interestingly, one of the major challenges in using the DSA is getting the data off the machine. Due to its age, the DSA is not set up with modern data transfer tools (such as a USB port) that we now take for granted. Fortunately, my job was made easy since I was able to download the analyzer data into Matlab using a GPIB-to-Ethernet converter and a series of scripts written by Mike Boulet. I owe many thanks to Mike for his long hours working on this and other tasks during the final stages of this project.

---

[1] Sometimes, I uncovered some peculiarities in the dynamics that I wish I hadn't!

Figure 5-1: The AFSM electronic support and test hardware at Lincoln Laboratory. The HP 3562A DSA is on the left.

5.1 System Design Using Capacitance Probe Feedback

I began the system identification of the AFSM by using the capacitance probes for feedback. I did so because the the probes' measurement range (±250 microns) includes the full expected mirror travel, which allowed me to characterize the full mirror range and drive it at many different current amplitudes. The optical sensor, on the other hand, works only in a limited portion of the travel, and as such is only suitable for small-signal measurement and feedback.

5.1.1 Open-Loop Dynamics

Before the capacitance probe rotation and scaling electronics were fully operational, I decided to take a preliminary frequency response using a single probe output in order to get an initial glimpse into the system dynamics. This data set would of

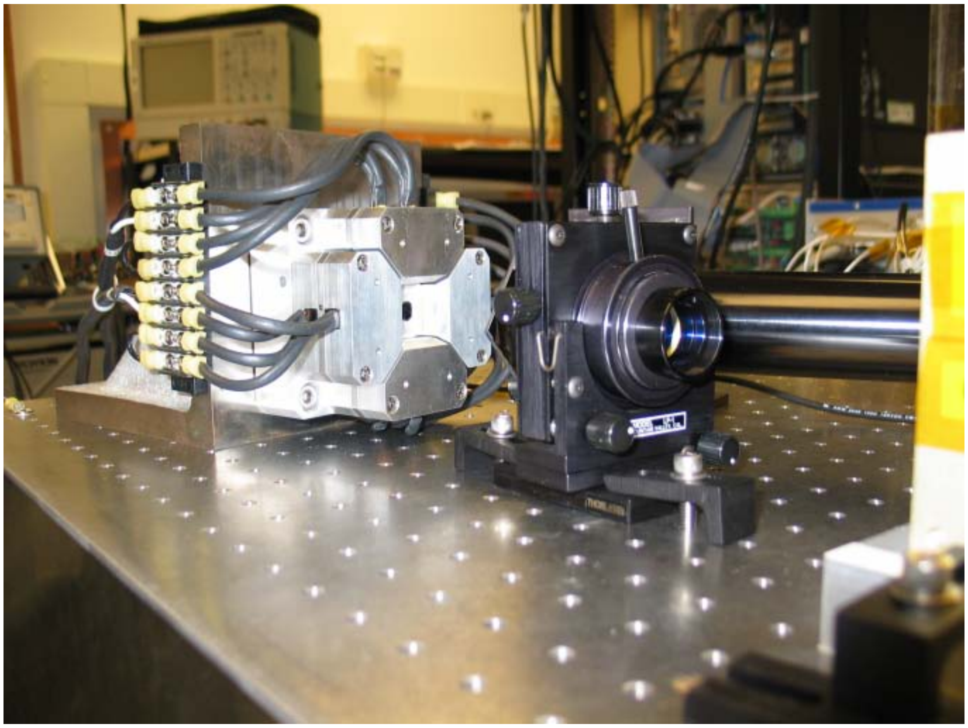

Figure 5-2: The AFSM installed on the optical table during testing, along with optical test hardware.

course be unacceptable for a controller design due to the gain mismatch and imprecise representation of true mirror angle, but nevertheless the overall shape of the response, particularly with regard to the resonant frequencies and phase profiles, was valuable information that would be immediately useful.

The initial system dynamics are shown in Figure 5-3. As expected, several details are immediately apparent. The first and most surprising thing I noticed is that the location of the first suspension mode resonant peak was far higher than the design frequency, at about 2.1 kHz. Contrasting with the design value of 356 Hz, the measured data shows a stiffness increase of about 25 times the design value. Note that I did not expect to see a resonant peak *at* the frequency suggested by the DC-based design, since it was known based on Barton's work [17] that stiffening of the elastomer bearings with frequency does occur. However, the actual location of the peak indicated a much higher stiffening effect than I expected. Later in the program I performed an extensive investigation into the high stiffness issue, which I discuss in Appendix B.

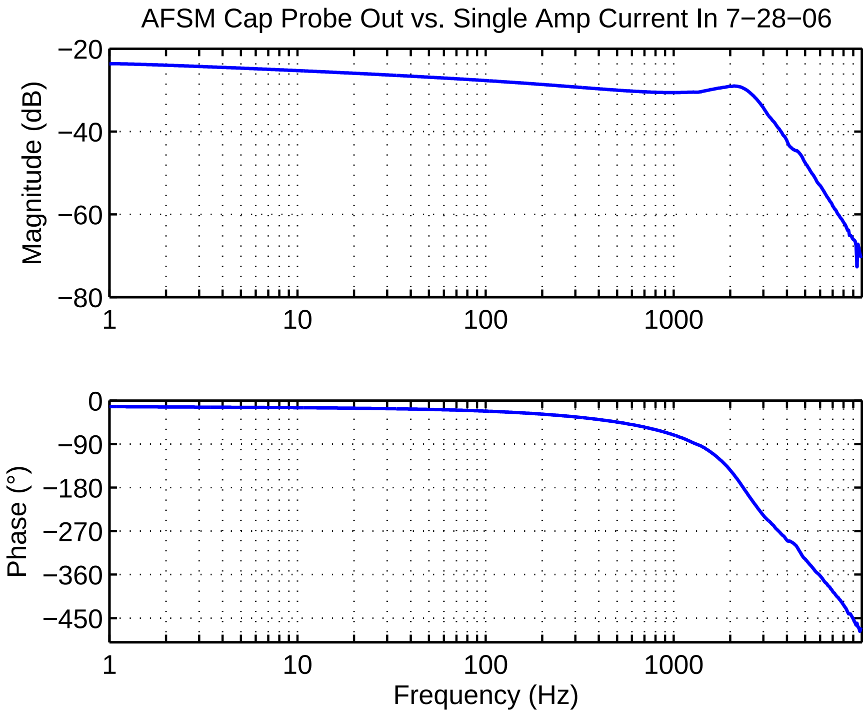

Figure 5-3: Initial frequency response of elevation axis.

In addition to the resonant peak location, several other characteristics in the data are notable. First, there are no spurious structural modes evident through 10 kHz. One very weak complex zero-pole pair is noticeable between 4 and 5 kHz, but the mode is co-located and thus has no effect on the phase profile. The frequency suggests that it may be a very weakly-coupled axial mode. Second, there is a nonzero phase of about -15 degrees present at DC, which slowly begins to decrease after about 100 Hz. This shift is modeled as part of the viscoelastic properties of the elastomer bearings. Third, the resonant peak is very well damped due to the elastomer bearing dissipation. The peak magnitude is only about 2 to 3 dB, indicating a damping ratio of about 0.4. Finally, at frequencies above the resonance the magnitude slope is about -65 dB/decade, which is much steeper than the -40 dB/decade which would be expected for a pure second-order system. This behavior is confirmed by the phase plot, which shows a steadily-decaying phase roll-off instead of the flat response that would be expected following the 180-degree transition at resonance.

The latter characteristic is the most troublesome for purposes of control design.

Since the phase response at crossover must be greater than -180 degrees for stability, attempting to cross over at frequencies above resonance will require large amounts of phase lead, which cannot be provided without a corresponding increase in magnitude at higher frequencies. As the order of the controller lead component is increased (in this case, to above third order), the magnitude slope becomes too shallow to allow for a well-defined crossover–and may even become positive–at which point further compensation is not possible.

Realizing that I could not meet the bandwidth performance goals with these dynamics, I set out to determine the exact origins of the phase roll-off characteristic. After researching all possible contributions, I came to the conclusion that the measured behavior was the composite response of several individual phenomena. They are, from most significant to least:

- **AFSM suspension mode.** The suspended resonance of the mirror on its flexures contributes phase lag that starts at -90 degrees at the resonant frequency and levels off at -180 degrees at higher frequencies.

- **Capacitance probe electronic filter phase lag.** After several conversations with Roy Mallory, the chief engineer at ADE, I discovered that the sensors themselves were the second-largest contributor to the phase response, and the largest factor in the roll-off behavior after resonance. Although the advertised bandwidth of the sensors is 10 kHz, the sensor package is designed with ‘static’ metrology as the primary application, rather than real-time feedback. As such, the sensors and conditioning electronics are designed to have a very accurate (flat) magnitude response at the expense of phase. In the case of the AFSM system, the conditioning electronics incorporate a seven-pole Butterworth electronic filter, whose frequency response is shown in Figure 5-4. Notice that while the magnitude response does indeed display a -3 dB bandwidth of 10 kHz, the phase response rapidly decays past 1 kHz due to the filter design. At 10 kHz, the phase is already below -200 degrees, and continues to decay to nearly -600 degrees at 100 kHz.

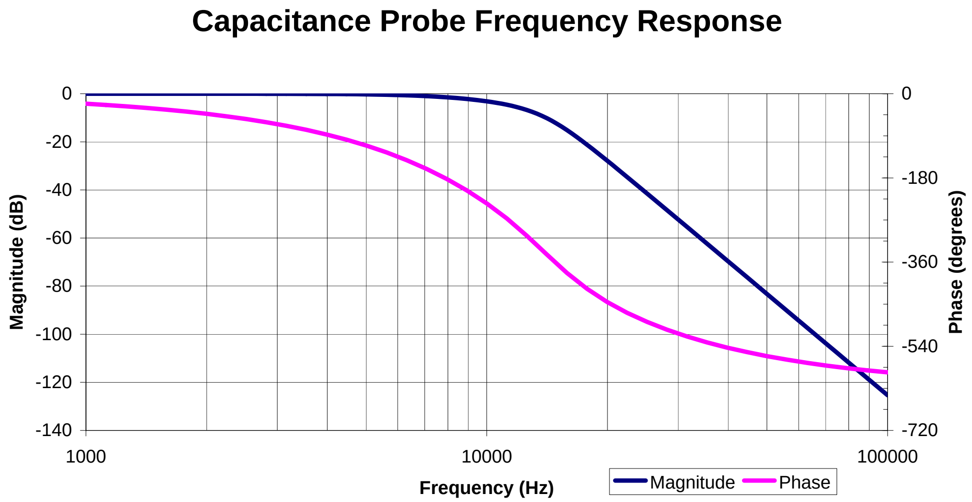

Figure 5-4: Frequency response of the capacitance probe and conditioner, courtesy of Roy Mallory.

- **Elastomer bearing dynamics.** As mentioned above, the stiffening of the rubber bearings with frequency results in dynamic behavior similar to a lag compensator when viewed from the perspective of compliance (i.e. the ratio of position output to force input). Since an accurate dynamic model of the elastomer is not available, determining its precise phase contribution to the overall dynamics is difficult. However, subtracting the known contributions from the measured dynamics suggests that the phase loss is about -30 to -90 degrees at 10 kHz.

- **Magnetic hysteresis and eddy current losses.** Because the input measurement in Figure 5-3 is of current rather than a direct force measurement, the dynamic losses in the magnetic system are present in the data. Although I performed a measurement of the tape-wound core efficiency by bolting two spare actuator halves together (see Figure 5-5), I did not characterize the losses of the SM-2 armature and backiron. The wound cores show very low losses above 100 Hz (the low-frequency roll-off is due to the finite reluctance of the core and air gap), but the SM-2 losses are inevitably much higher. Thus, the magnetic dynamics also contribute to the overall phase decay.

As an additional preliminary task, I explored the change in frequency response

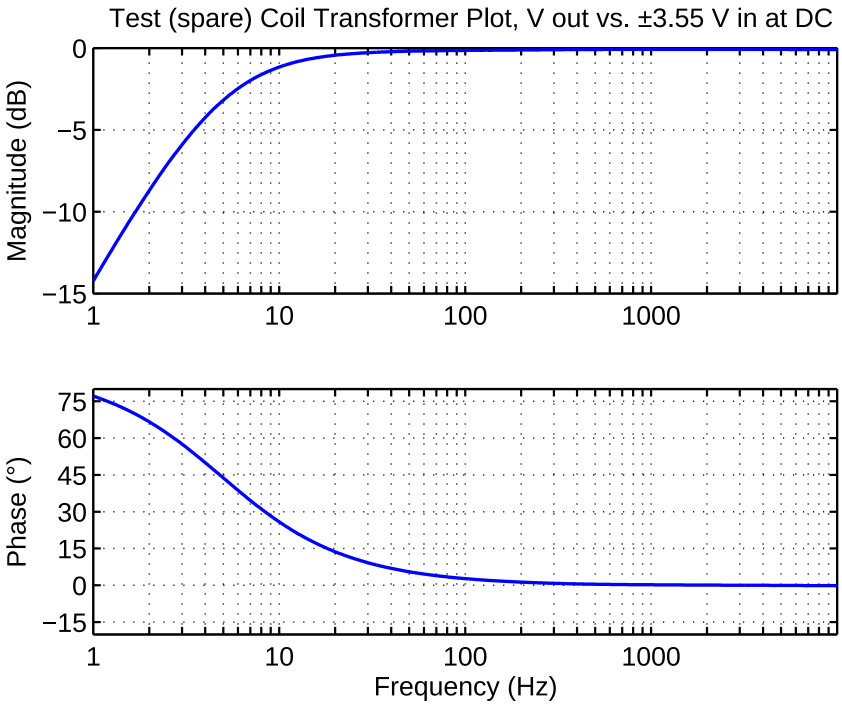

Figure 5-5: Frequency response of the AFSM actuator halves when configured as a transformer (ratio of upper coil input voltage to lower coil output voltage).

with increasing current amplitude. These plots are shown in Figures 5-6 and 5-7. Although not markedly different, it can be observed that the DC gain increases with increasing current amplitude[2], while the natural frequency and damping decrease. This suggests a decrease in the elastomer bearing stiffness and damping with increasing amplitude. Notice also in Figure 5-7 that the weakly-coupled axial resonance is much more pronounced in the azimuth direction than in elevation, indicating a slight degree of asymmetry in the azimuth geometry (presumably due to manufacturing and/or assembly error).

It was also during this sequence of testing that I confirmed the high stiffness behavior of the AFSM. I found that the angular response amplitude relative to the derated force output of the actuators[3] was too low for the rubber to be at the designed stiffness value, even at near DC frequencies. As a result, the maximum measured range

---

[2] The values listed in the plot key are half-wave amplitudes, not peak-to-peak values.

[3] Since the original magnetic analysis does not account for magnetic leakage and fringing effects, the actuator output must be derated relative to the theoretical value to obtain a reasonable estimate of the real force. Here I used the results from a magnetic FEA by Xiaodong Lu [16] to obtain a derating multiplier (on the force output) of 0.7.

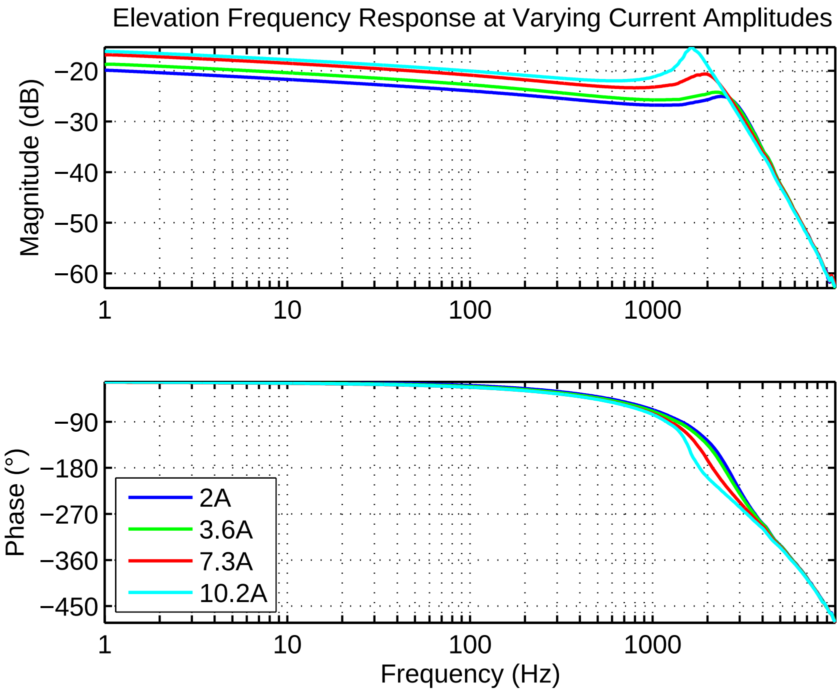

Figure 5-6: AFSM elevation axis frequency responses at various current amplitudes using capacitance probe feedback (ratio of capacitance probe voltage out to amplifier current command in).

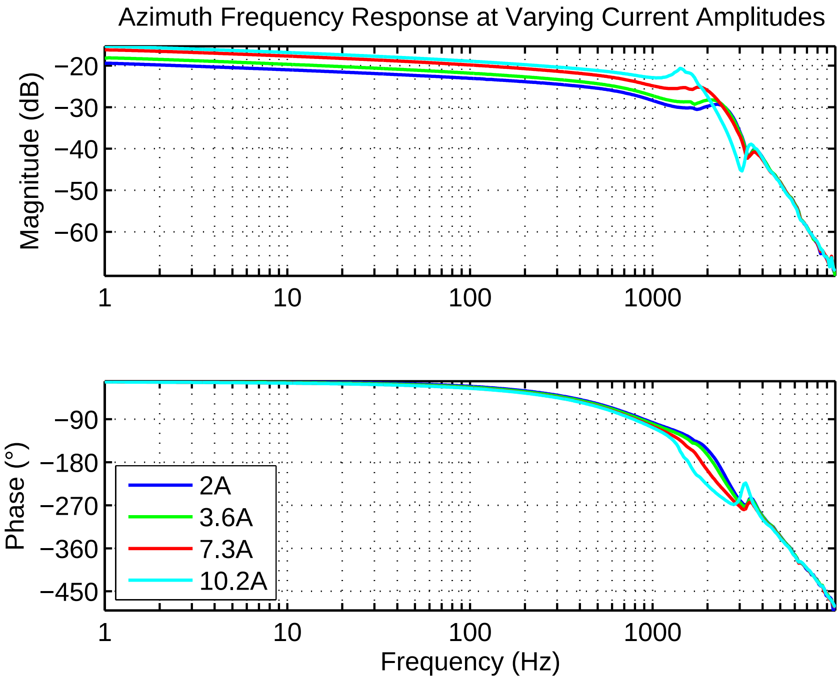

Figure 5-7: AFSM azimuth axis frequency responses at various current amplitudes using capacitance probe feedback (ratio of capacitance probe voltage out to amplifier current command in).

of motion was ±3.5 milliradians, rather than the design value of ±10 milliradians.

Once the probe conversion and rotation electronics were complete and operational, I was able to measure a response using the true angle as the output, such that the closed-loop system would respond correctly to an angular input signal. I used the representative data set shown in Figure 5-8 as the baseline for my controller design, which is discussed in the following section. Notice that the inclusion of the rotation electronics changes mainly the gain of the plot. The actual dynamics change only slightly due to the averaging effect of the combined output from the four sensors.

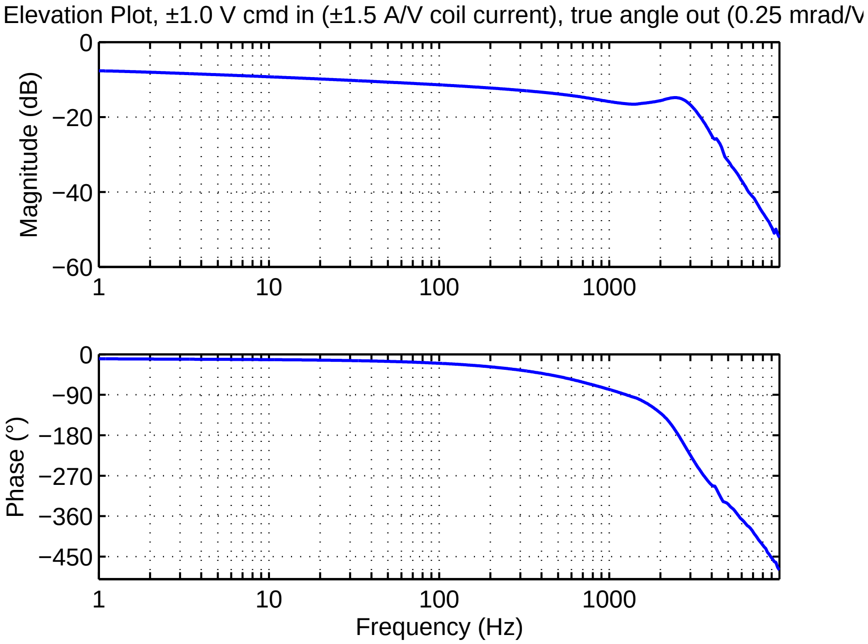

Figure 5-8: AFSM elevation axis frequency response used to design the controller for capacitance probe feedback (ratio of rotated capacitance probe voltage out to amplifier current command in).

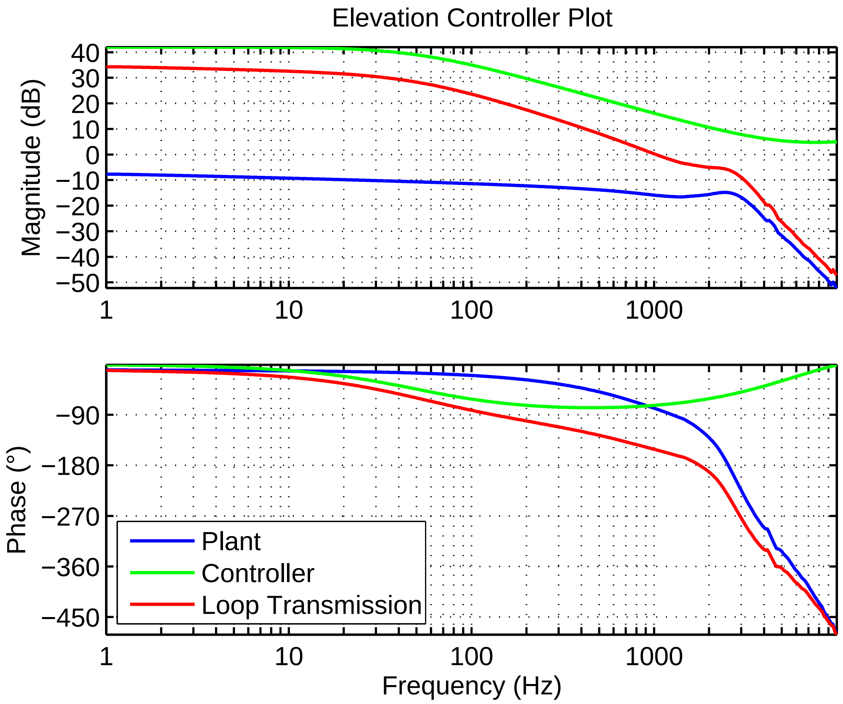

Figure 5-9: Capacitance probe controller design (green) with measured plant dynamics (blue) and resulting loop transmission (red).

5.1.2 Analog Compensator

With the open-loop capacitance response in hand, I proceeded with the controller design. The phase characteristics of the sensors combined with the higher-than-expected suspended mode natural frequency of the plant forced a conservative controller design, but I still tried to extract as much performance as possible.

The features of the design are shown in Figure 5-9. The original open-loop dynamics are plotted in blue, and the designed compensator in green. The red line plots the expected loop transmission when the controller and plant are combined[4].

The controller itself is a simple lead-lag configuration, with a lag pole at 50 Hz, the lag and lead zeros both at 7 kHz, and the lead pole at 25 kHz. The DC gain of the controller is 42 dB, and the transfer function of the controller is

---

[4] I wrote a script in Matlab that allowed me to create these plots using the actual experimental data as input, with the theoretical controller transfer function overlaid. By using this tool, I circumvented the task of trying to fit a high-order transfer function to the data, which resulted in a more accurate controller design. It was also very time-efficient, allowing me to perform many design iterations quickly.

$$G_c(s) = 126\frac{\left(\frac{s}{2\pi(7000)}+1\right)^2}{\left(\frac{s}{2\pi(50)}+1\right)\left(\frac{s}{2\pi(25000)}+1\right)} \tag{5.1}$$

The objective here is to get the gain as high as possible within the controlled bandwidth using the lag stage, and then recover as much phase as possible using the lead stage in the vicinity of crossover. As can be seen from Figure 5-9, the crossover frequency is 1 kHz, with a phase margin of 30 degrees and a gain margin of 5 dB. Notice that I had to place the crossover at a frequency lower than the resonant peak due to the excessive phase loss after resonance. Examining the figure, it is evident that attempting to cross over at the desired bandwidth of 5 kHz would require over 150 degrees of phase lead, which is a formidable challenge even for a triple or quadruple lead design. Also, such a design would not be tractable due to the magnitude issues discussed above. However, since the main contributor to the loss at these frequencies is the capacitance probes, I expected to achieve the required bandwidth when using optical feedback due to the improved sensor dynamics.

Originally my plan was to implement the compensator digitally using a real-time computer system. However, after characterizing the available hardware in the lab facility (a PC with a commercial data acquisition card running Matlab's xPC real-time operating system), I concluded that the maximum sampling frequency of the digital system, about 40 kHz, was too low to provide robust control at the bandwidth and performance level demanded by the AFSM, as it would have added even more phase loss to the system. Therefore, I designed the analog circuit shown in Figure 5-10 to perform the compensator functions.

The details of the analog circuit analysis are presented in Appendix C, so I describe the circuit functionality in broader terms here. It consists of a unity-gain differential amplifier which is used as a summing junction. The output of the junction is fed into an operational amplifier configured as an inverting lag stage. The stage also provides a DC gain which is adjustable via the 2k potentiometer[5]. The output of the lag stage

---

[5] Note that in the iteration of the circuit shown, I decided to reduce the loop gain to obtain better robustness in the loop, so the adjustment range is 55 to 95 (35 to 40 dB) rather than 126 (42 dB) as originally designed.

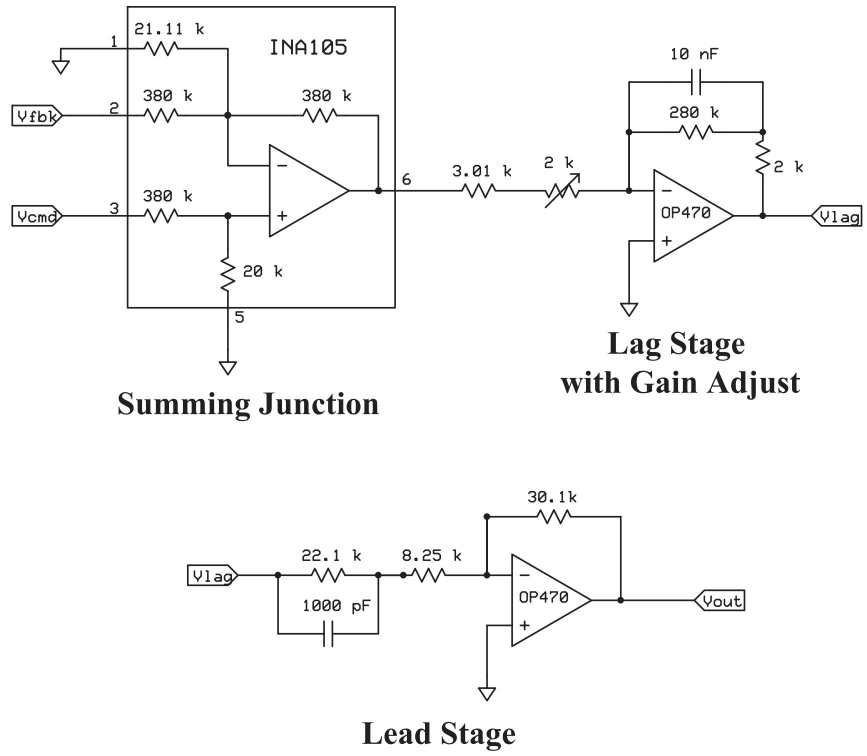

Figure 5-10: The AFSM compensator circuit for use with capacitance probe feedback (one circuit per axis).

feeds into the inverting lead stage, which is configured to have unity gain at DC. Note that since there are two inversions, the resulting signal has positive polarity relative to the command signal.

I built a single-channel breadboard version of the circuit in Figure 5-10 to test the actual compensator response and loop performance. The frequency response of the compensator alone is shown in Figure 5-11, along with the theoretical as-designed response. I was very pleased with the accuracy of the response, which suffered only a slight departure in phase relative to the ideal after 10 kHz (presumably due to non-ideal dynamic behavior of the op-amps and/or passive components)[6].

After verifying the operation of the controller, I connected it to the AFSM amplifier inputs to measure the open-loop dynamics. The result of this test is shown in

---

[6] The high degree of precision obtained by the controllers I designed for the AFSM was due in large part to the extensive selection of premium quality electronic components in the Lincoln Laboratory stockroom. For example, the resistors I used were manufactured to resistance values of three significant figures and tolerances of less than 1%

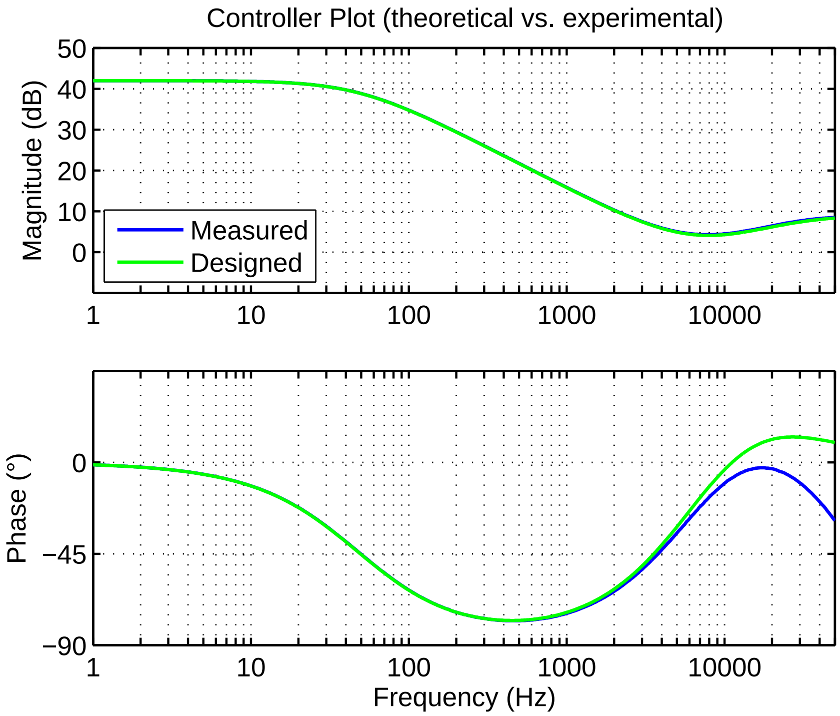

Figure 5-11: Analog capacitance probe compensator measured dynamics overlaid with design values.

Figure 5-12. Aside from a small amount of noise at high frequencies, the measured loop transmission is virtually identical to the expected dynamics in Figure 5-9.

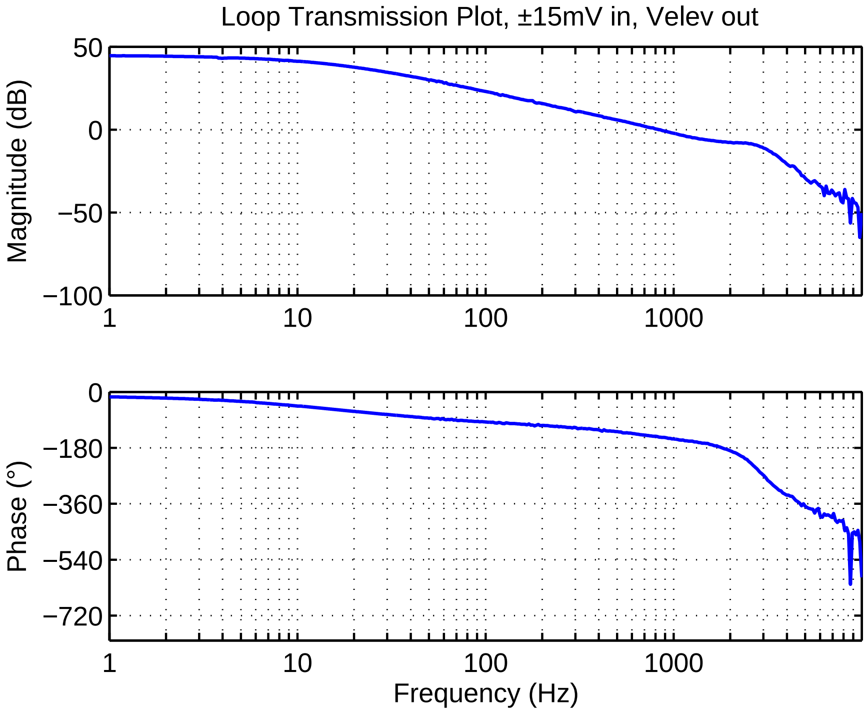

Figure 5-12: Measured loop transmission of the AFSM plus capacitance probe analog controller.

5.1.3 Closed-Loop Performance Using Capacitance Probe Feedback

With the controller designed and operational, the next step was to evaluate the closed-loop performance. With the controller running at a loop gain of 126, the closed-loop plots in elevation and azimuth are shown in Figures 5-13 and 5-14, respectively.

The plots are very similar, with the exception of one subtle difference. While the elevation dynamics appear nominal, the azimuth plot displays a strange ‘bump’ in the transition through the closed-loop peak, with a corresponding small perturbation in the phase response. The source of the bump is troublesome, because no signs of it appeared in the azimuth loop transmission.

In general, I found the azimuth axis of the AFSM to be harder to control than the elevation axis, as it would exhibit sporadic instability in certain situations. Clues to the source of these instabilities may be found by examining Figure 5-7. The major difference between it and the elevation plot of Figure 5-6 is the presence of the co-located mode at 3.5 kHz. Notice that as the current amplitude is increased, the peak magnitude and phase also increase. My hypothesis is that since the closed-loop controller increases the amplifier current output to high levels with increasing frequency, the 3.5 kHz mode may grow to the point of being destabilizing. I encountered no such problems in the elevation axis.

As a verification of this effect, later in my testing I decided to reduce the loop gain of the controller to 80. The closed-loop responses resulting from the reduction are shown in Figures 5-15 and 5-16. Notice that they are now virtually identical, which is consistent with the loop transmission measurements.

The maximum bandwidth performance I could achieve with the capacitance probes was just over 2 kHz, as shown in Figures 5-13 and 5-14. Reducing the gain to 80 resulted in a much flatter magnitude response at the expense of lowering the bandwidth to 1 kHz.

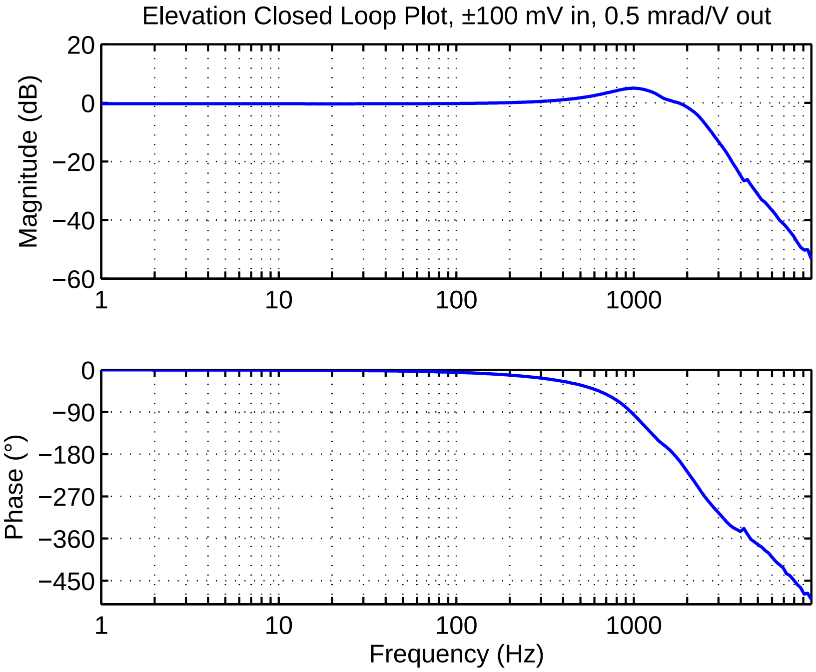

Figure 5-13: Elevation closed-loop frequency response using capacitance probe feedback.

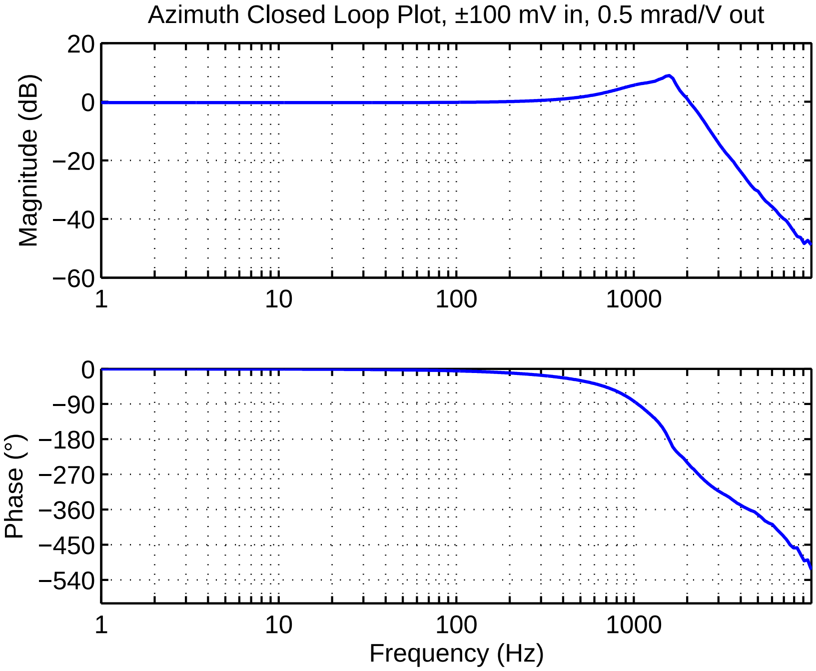

Figure 5-14: Azimuth closed-loop frequency response using capacitance probe feedback.

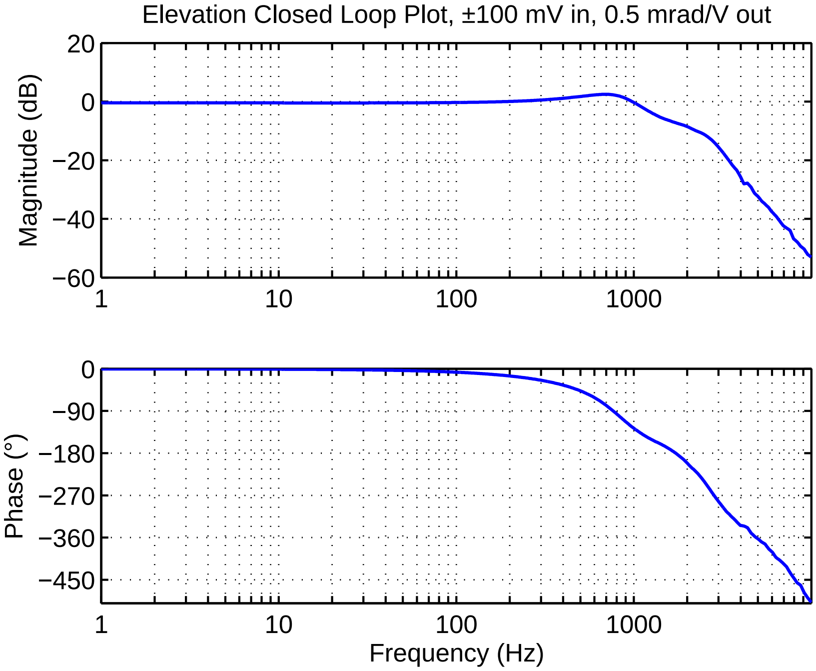

Figure 5-15: Elevation closed-loop frequency response with reduced loop gain.

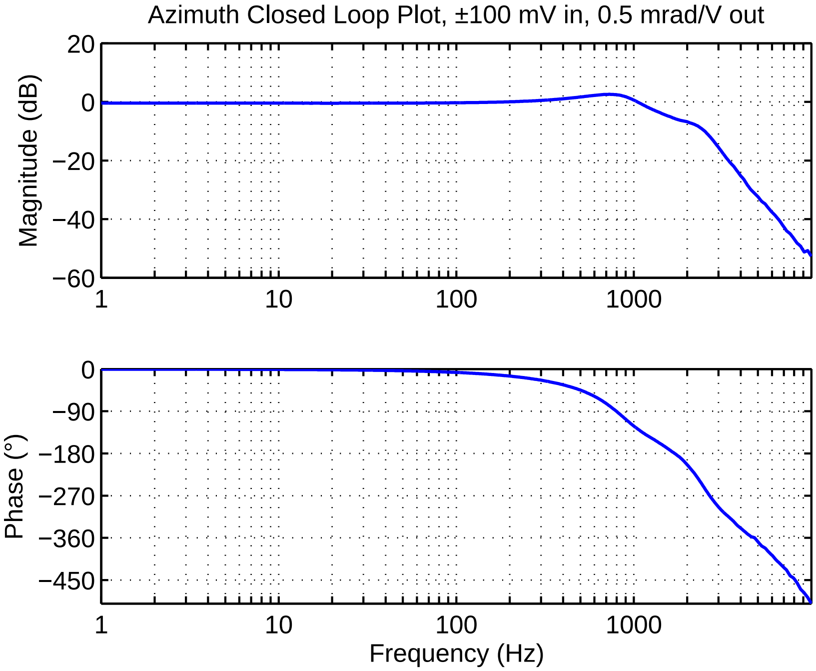

Figure 5-16: Azimuth closed-loop frequency response with reduced loop gain.

The step response of the system under capacitance probe feedback is shown in Figure 5-17. In this case, the DC gain is 126. The rise time of approximately 180 microseconds is consistent with the natural frequency in the frequency response plot, using the approximation $\omega_n \approx 1.8/t_r$. The overshoot is just over 50%, and the settling time is about 1.2 milliseconds.

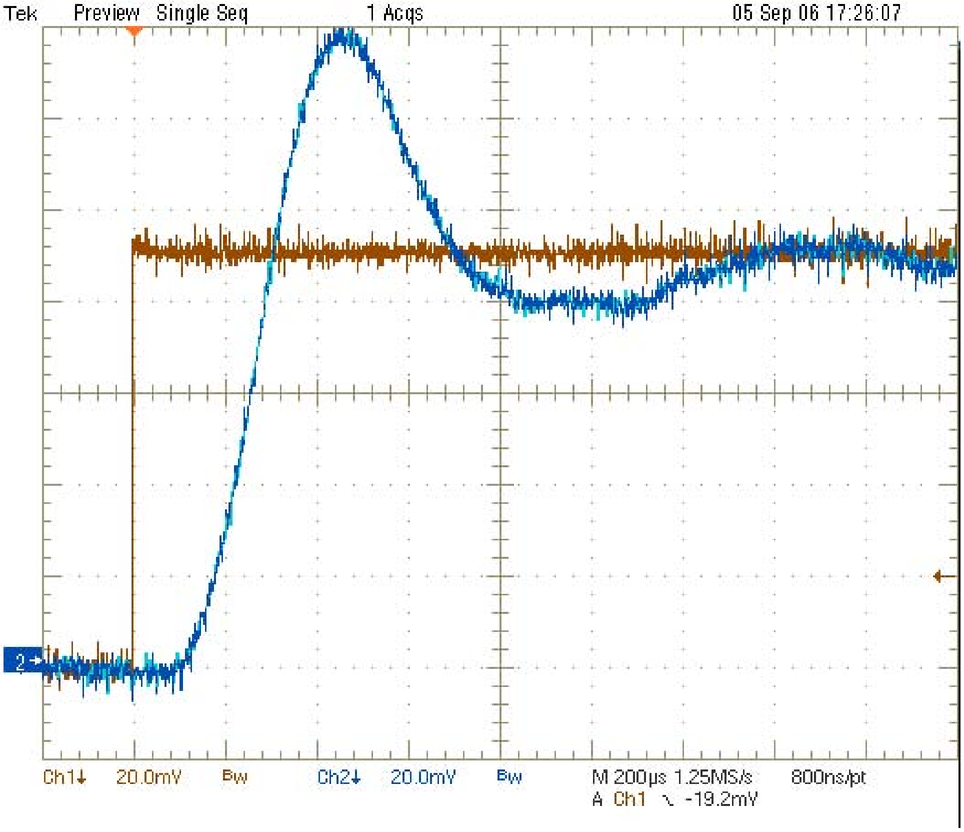

Figure 5-17: Step response of the elevation axis under capacitance probe feedback.

5.2 System Design Using Optical Feedback

Because of the disappointing phase characteristics of the system using the capacitance probes, I was eager to take measurements using the optical feedback setup to determine if better performance could be achieved. After bonding a small glass mirror to the face of the AFSM (since I never actually polished the AFSM "mirror" to an optical finish) and aligning the laser to the quad cell, I calibrated the optical angular output at 25 μrad/volt and proceeded to identify the AFSM system using the new measurement, as described in the following sections.

I gained an appreciation for the precision of the optical setup when I first tried to take quad cell measurements in open air. As it turns out, the air currents in the room coming from the ventilation ducts were disrupting the laser beam enough to introduce a significant (about ± 1 volt) low-frequency random disturbance into the measurement signal! I eliminated this error source by using a plastic enclosure to provide a still-air environment for the test setup, as shown in Figure 5-18.

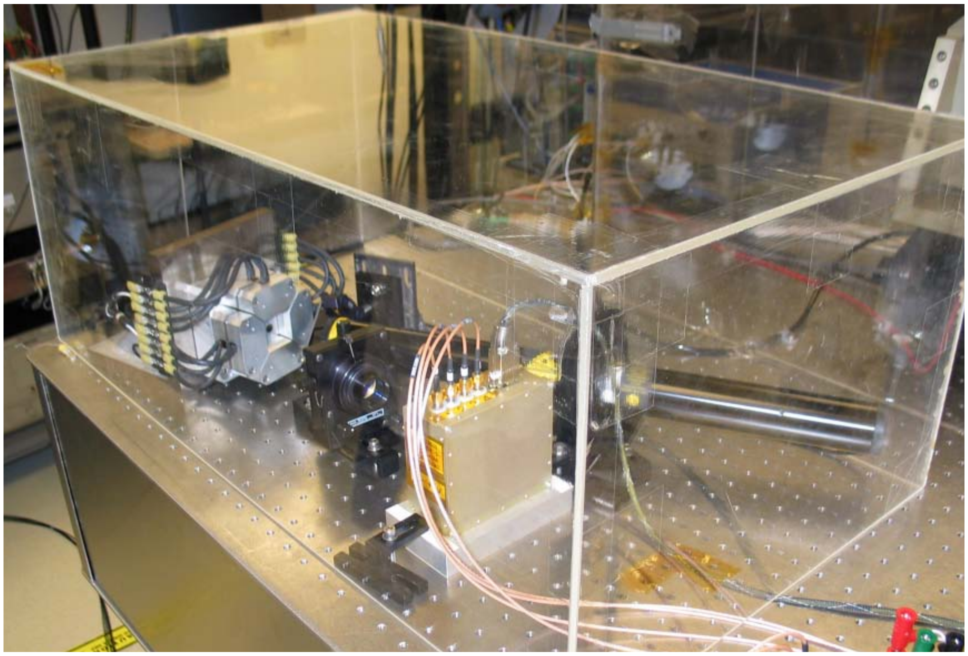

Figure 5-18: The plastic enclosure used to protect the test setup from room air currents.

5.2.1 Open-Loop Dynamics

The elevation and azimuth plant dynamics using optical feedback are shown in Figures 5-19 and 5-20. Here, the higher bandwidth of the optical sensor allowed me to capture meaningful data up to 50 kHz. With the wider frequency range, the first few structural modes of the AFSM mirror are revealed. The first one appears at about 23 kHz, which agrees very well with the first two modes at 22 kHz predicted by the finite element analysis in Chapter 2. An additional group of modes is visible at 35 kHz and the 40 to 50 kHz band, which is also consistent with the FEA predictions of the higher modes.

More importantly, from the figures it is also evident that the high frequency phase characteristics are greatly improved relative to the capacitance probe dynamics. At the target bandwidth in both plots the phase is only about -220 degrees, making the optical system much easier to stabilize than the system using capacitance probe feedback.

To gain a better understanding of the improved phase provided by the optical sensor, I created the overlay plot shown in Figure 5-21. For comparison purposes I adjusted the gain of the capacitance probe magnitude data to match the DC gain of the optical measurement. Here, the difference in phase between the two measurements is clear: At 1 kHz, the capacitance probes add more than 45 degrees of phase lag to the system, and at 10 kHz the difference is greater than 200 degrees! The difference in phase proved to be critical for obtaining maximum performance, as discussed in the following sections.

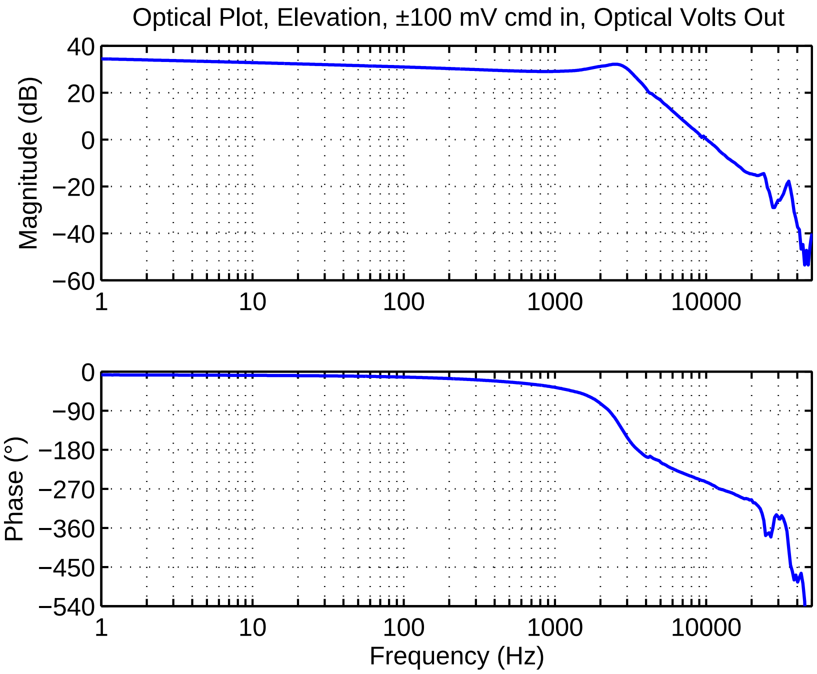

Figure 5-19: AFSM elevation axis frequency response using optical (quad cell) feedback.

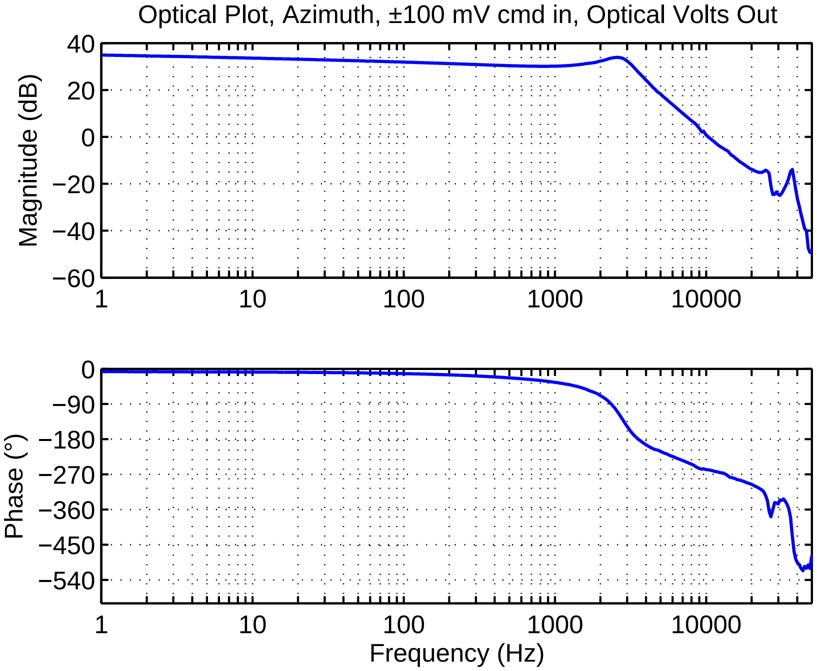

Figure 5-20: AFSM azimuth axis frequency response using optical (quad cell) feedback.

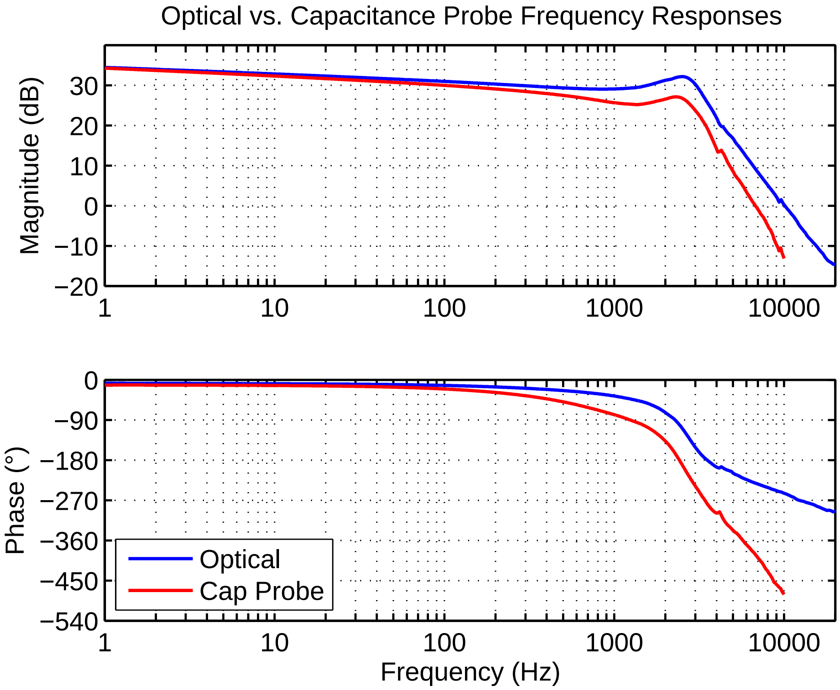

Figure 5-21: Comparison of frequency responses obtained using the optical quad cell versus the capacitance probes.

5.2.2 Analog Compensator

First Iteration

Repeating the general procedure in Section 5.1.2, I designed a controller around the optical AFSM data. The frequency response of my first attempt is shown in Figure 5-22. Its transfer function is

$$G_c(s) = 0.63 \frac{\left(\frac{s}{2\pi(1000)} + 1\right)\left(\frac{s}{2\pi(4000)} + 1\right)^2}{\left(\frac{s}{2\pi(100)} + 1\right)\left(\frac{s}{2\pi(20000)} + 1\right)^2} \tag{5.2}$$

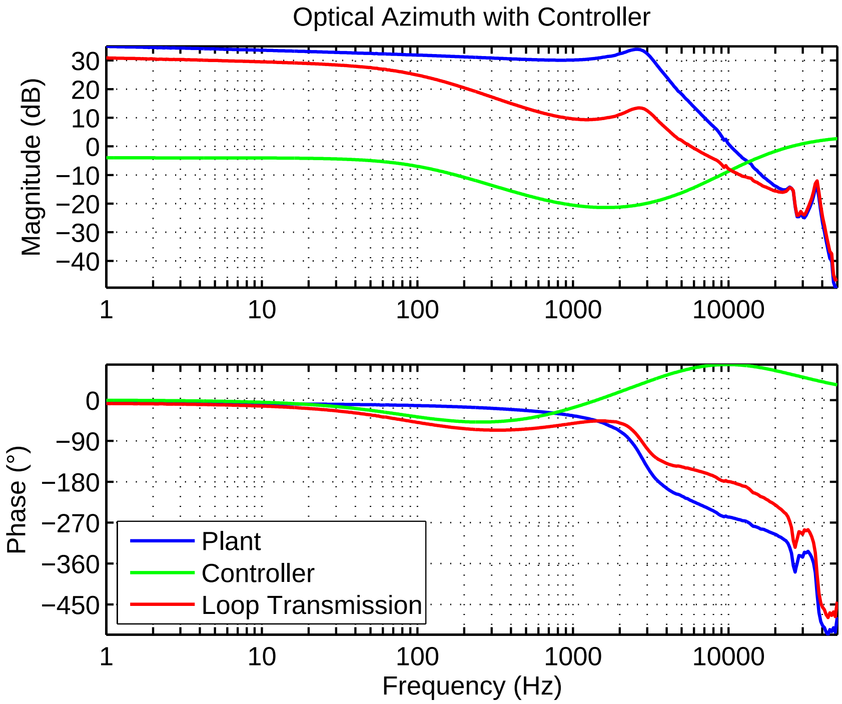

Figure 5-22: Optical controller design (green) with measured plant dynamics (blue) and resulting loop transmission (red).

The new controller design is a lag compensator for low frequency gain, along with a double lead compensator for extra phase boost (about 70 degrees) at crossover. Here the improved phase characteristics allowed me to cross over at 5.5 kHz, which is beyond the suspended resonant frequency per my original intentions in the mechanical design. Note that because of the very high sensitivity of the quad cell electronics I actually had to lower the controller loop gain below unity to achieve reasonable design

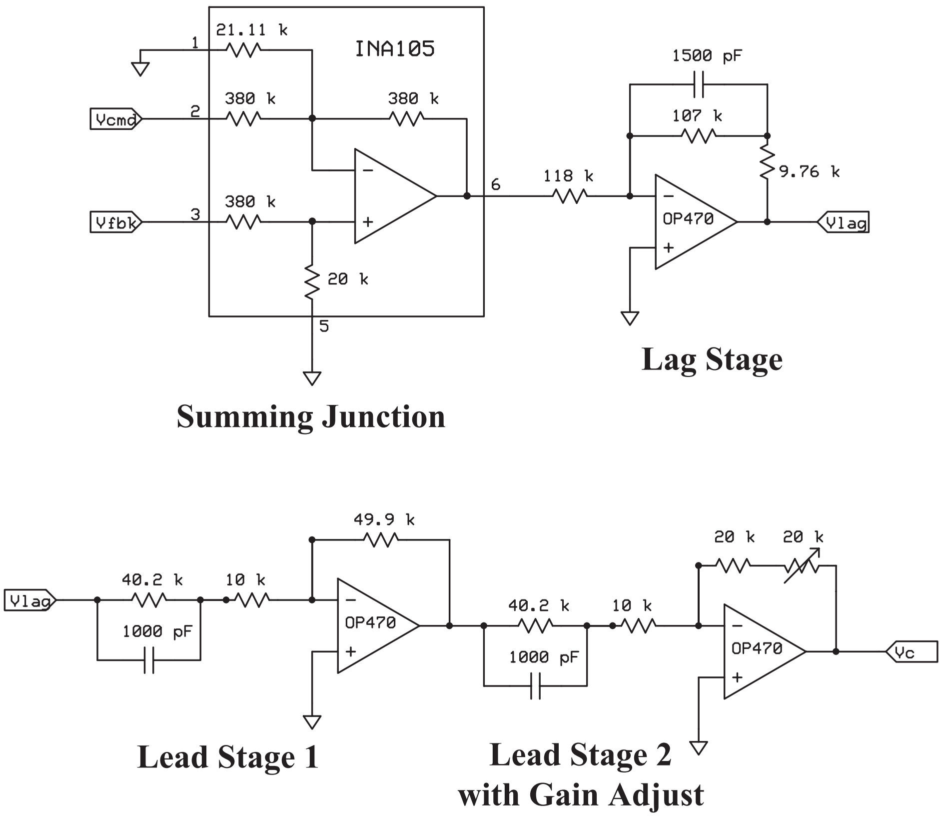

Figure 5-23: The AFSM compensator circuit for use with optical feedback (one circuit per axis).

margins. The phase and gain margins are about 30 degrees and 7 dB, respectively. In addition, the potentially dangerous resonant peaks at 23 and 35 kHz both have comfortable gain margins of greater than 10 dB.

With the optical controller operating at even higher frequencies than the capacitance probe version, I again ruled out using a digital controller, and designed the new analog circuit shown in Figure 5-23 to perform the hardware compensation. Besides the exact component values, it differs from the capacitance probe compensator in only a few ways: The additional lead stage requires another op-amp, which causes the circuit to become inverting. Therefore, the command and signal inputs at the summing junction are swapped to restore proper signal polarity. Also, because the loop gain is attenuating, the gain adjustment is moved to the output stage of the compensator. The idea here is to maintain large signals in as much of the circuit as possible.

In a similar manner to the previous compensator, I built a single-axis breadboard version of the circuit design and ran a stand-alone frequency response before connecting it to the hardware. The measurement is shown in Figure 5-24, again with very accurate results. The measured loop transmission is in Figure 5-25. Note that because of the slight amount of extra phase loss in the real hardware after about 4 kHz, the measured phase margin (20 to 25 degrees) is not as high as predicted in the design. I later lowered the loop gain to provide extra margin.

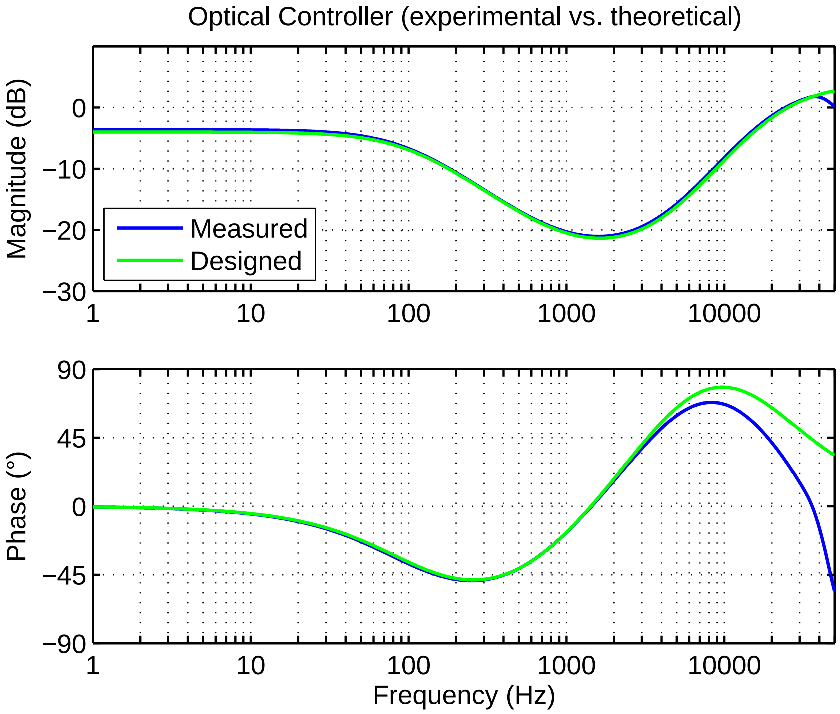

Figure 5-24: Analog optical compensator measured dynamics overlaid with design values.

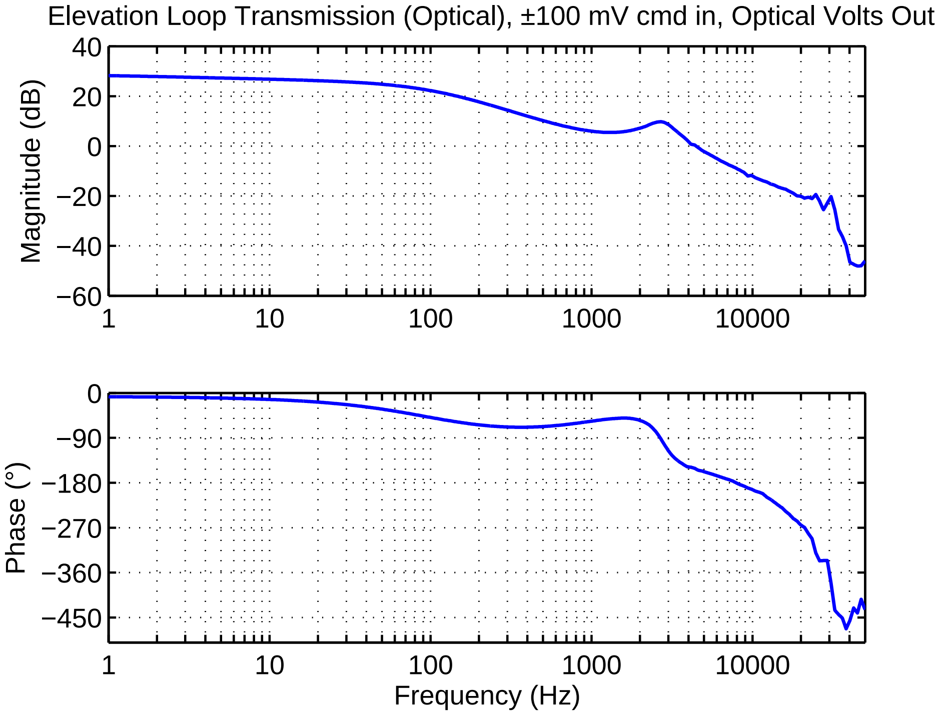

Figure 5-25: Measured AFSM loop transmission plot with optical compensator.

Second Iteration

Near the end of the project, after implementing two-axis control, I decided to revisit the optical controller, both to obtain better performance and to provide better stability margins for the azimuth axis. I observed that I could obtain better gain in the bandpass by moving the lag pole from 100 to 50 Hz, and the zero from 1 to 3 kHz. Similarly, I obtained extra phase from the lead compensators by moving the poles to 30 kHz and the zeros to 3 kHz. The nice thing was that neither of these changes moved the crossover frequency, which allowed me to maintain bandwidth. The new controller transfer function is

$$G_c(s) = 2\frac{\left(\frac{s}{2\pi(3000)}+1\right)^3}{\left(\frac{s}{2\pi(50)}+1\right)\left(\frac{s}{2\pi(30000)}+1\right)^2} \tag{5.3}$$

and the design frequency response is as shown in Figure 5-26. Both the gain and phase margins have improved, to 11 dB and 40 degrees, respectively.

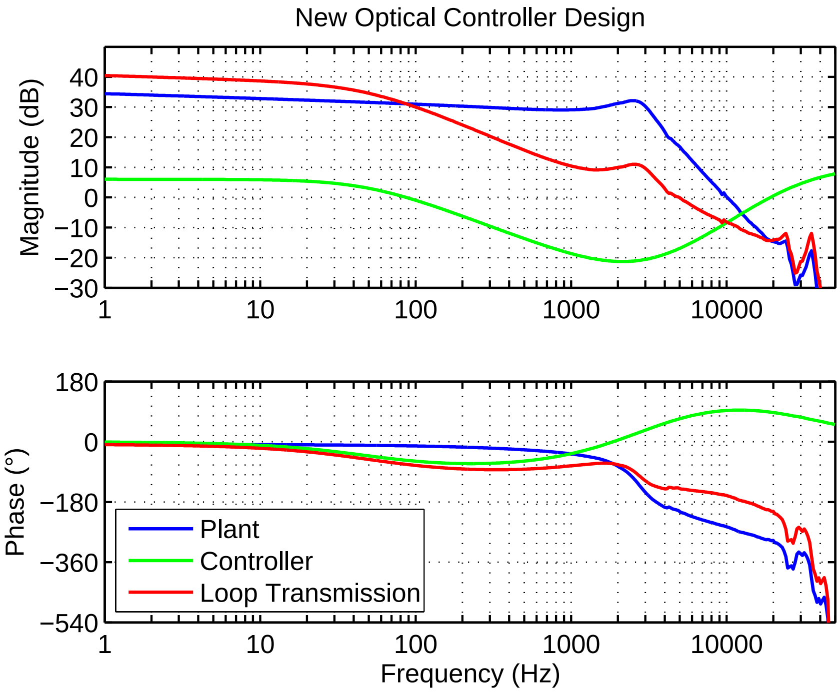

Figure 5-26: Second optical controller design (green) with measured plant dynamics (blue) and resulting loop transmission (red).

After modifying the controller via appropriate passive component changes, I measured the new controller frequency response. Again it proved satisfactory, as displayed in Figure 5-27. Notice that, in addition to the increased phase by design, the experimental plot displays less high-frequency phase loss relative to the theoretical than the original design in Figure 5-24. This of course improved the stability margins in the physical system.

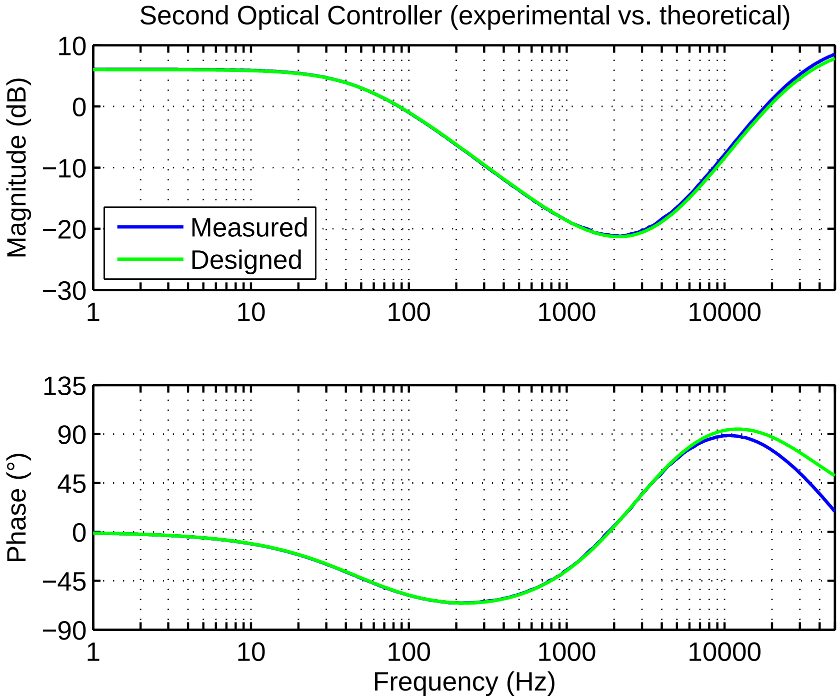

Figure 5-27: Second optical compensator measured dynamics overlaid with design values.

5.2.3 Closed-Loop Performance Using Optical Feedback

First Iteration

The breadboard closed-loop performance of the elevation axis with the original optical compensator is shown in Figure 5-28. This plot represents the maximum bandwidth I was able to achieve with the AFSM hardware: About 10 kHz at -3 dB. Of course, there is also about +8 dB of peaking at resonance due to the low phase margin. In addition, the magnitude displays a small droop after about 350 Hz as the loop gain

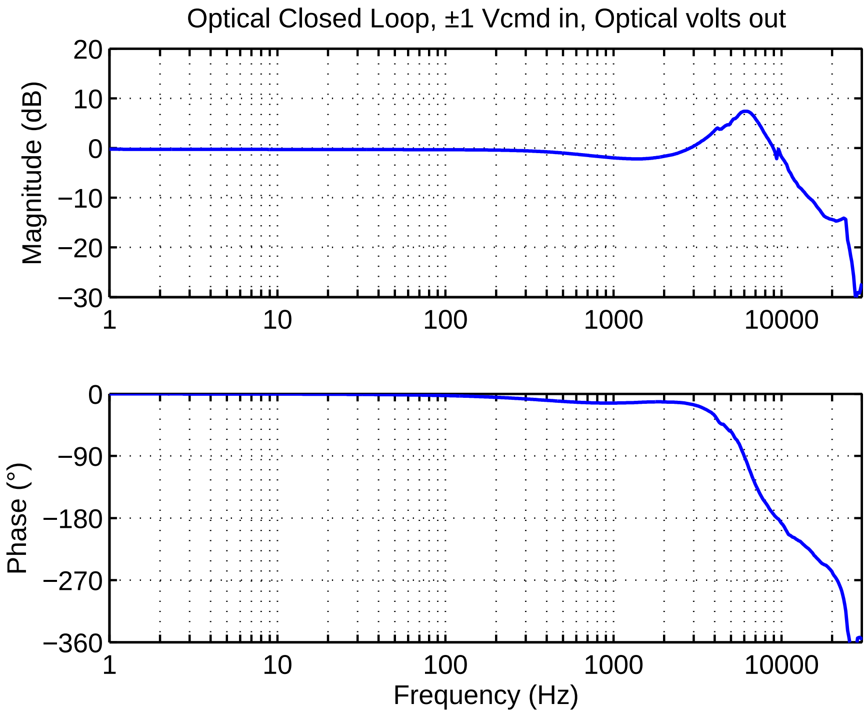

Figure 5-28: Elevation closed-loop frequency response using optical feedback, showing 10 kHz bandwidth.

begins to drop approaching the crossover frequency.

Unfortunately, with the loop gains as-designed I was not able to keep the azimuth loop stable to replicate the same performance in this axis. The loop transmission plot for the azimuth axis appeared normal, but when I ran the closed-loop it would go unstable at about 1.3 kHz, as shown in Figure 5-29.

The puzzling aspect of the instability is that no evidence of unstable behavior is apparent in the loop transmission plot. However, as mentioned in the previous sections, a clue may be found by examining the varying-amplitude plot of Figure 5-7. The increasing amplitude of the resonance as the controller increases the amplifier current magnitude may eventually cause the loop to go unstable. Moreover, since the optical controller has a crossover (5.5 kHz) that is very near the frequency of the resonance (3.5 kHz), the effect that merely produced an unusual dynamic behavior in the capacitance probe azimuth response may have proved fatal for the optical response.

My first attempted solution was to reduce the loop gain to produce additional phase margin. This turned out to be adequate: a 3 dB reduction in gain resulted in the plot in figure 5-30. Despite the gain decrease, the bandwidth is still respectable at approximately 7 kHz. In addition, the resonance peak magnitude is greatly reduced relative to Figure 5-28, at 4 dB.

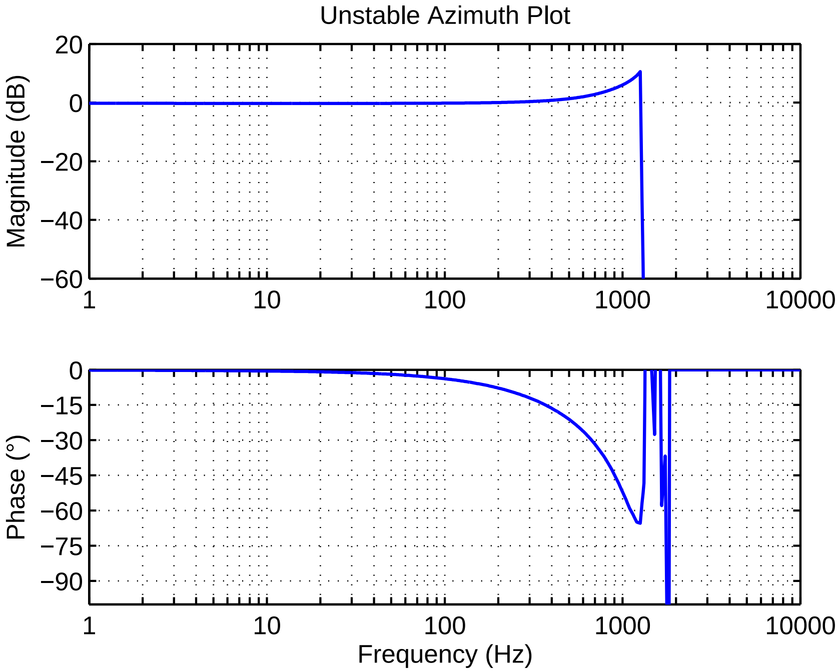

Figure 5-29: Azimuth closed-loop frequency response using optical feedback. The loop went unstable during the test at 1.3 kHz.

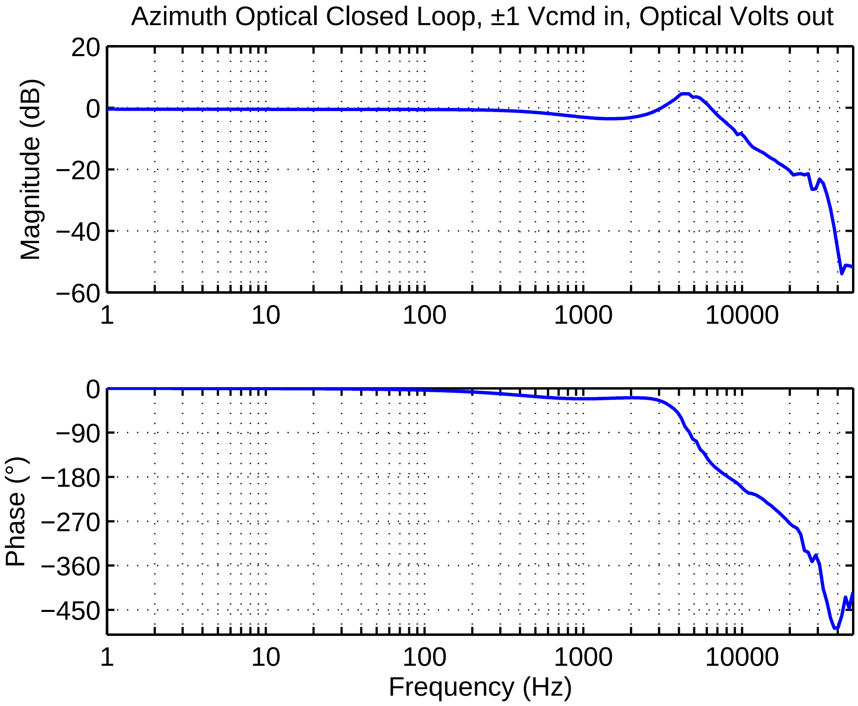

Figure 5-30: Azimuth closed-loop frequency response with reduced gain.

Second Iteration

I was able to achieve slightly better performance from the second iteration of the optical controller. The closed-loop frequency response is shown below in Figure 5-31. It is very similar to Figure 5-28, especially in that the bandwidth is the same–roughly 10 kHz. However, the magnitude performance of the new controller within the bandpass is flatter relative to the old design, and the increased loop gain helps the system track the command reference at 0 dB, especially below 500 Hz. The resonant peak has been reduced to about 6 dB, and the dip in magnitude from 600 Hz to 3 kHz is only about -2 dB at its lowest point.

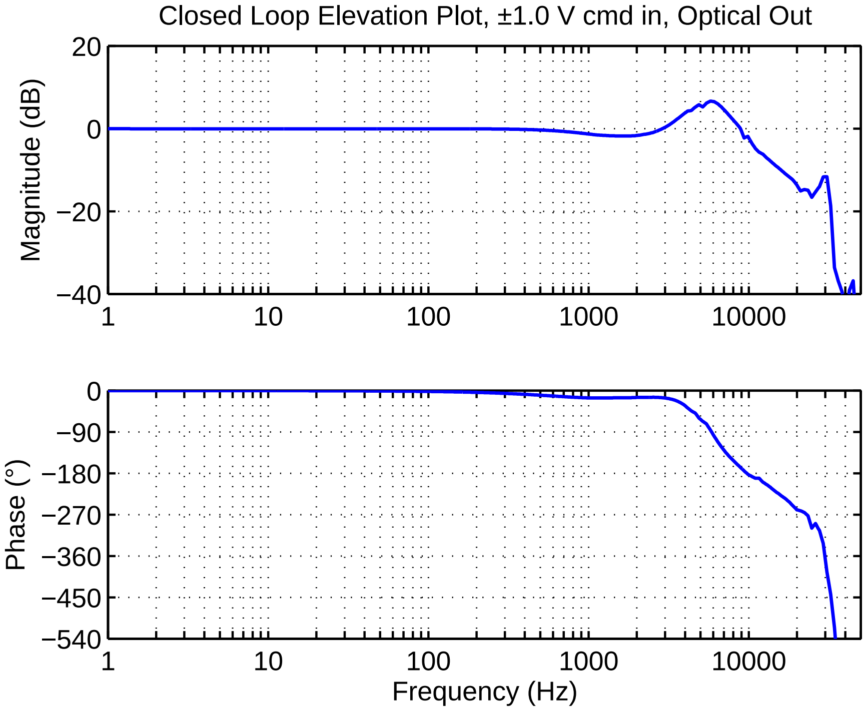

Figure 5-31: Elevation closed-loop frequency response using the new optical controller, showing flatter magnitude but retaining 10 kHz bandwidth.

Using this controller I was able to obtain pleasing step responses. The responses for various step sizes are plotted in Figures 5-32 and 5-33. The small-signal step, Figure 5-32 displays excellent dynamics with a rise time of about 40 microseconds, a settling time of 400 microseconds, and overshoot of only 25%. The performance degrades somewhat with the large step size in Figure 5-33, but is still respectable at

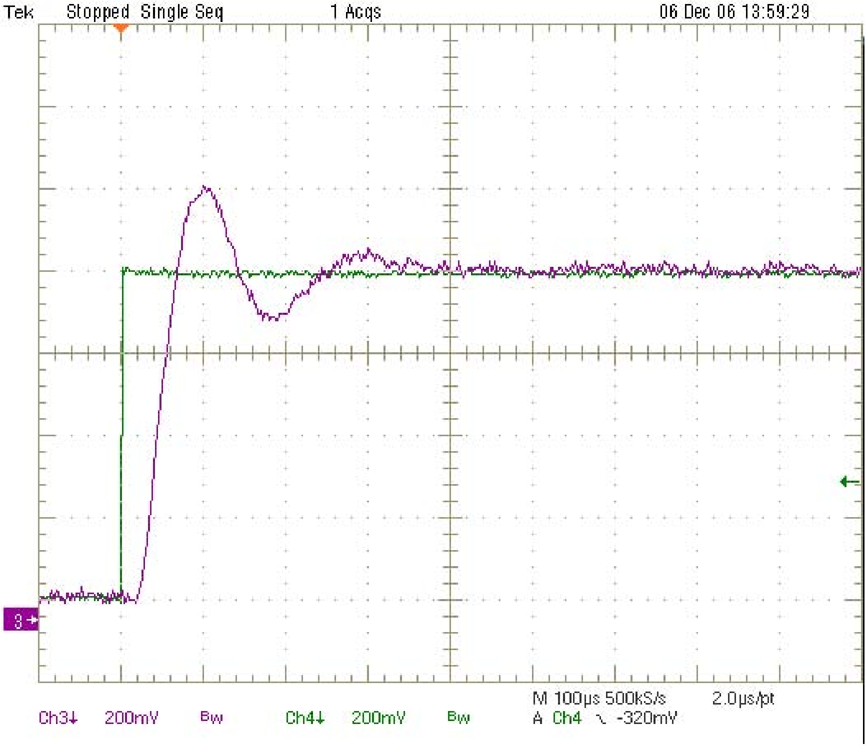

Figure 5-32: 800-millivolt (20 μrad) step response using optical feedback and controller.

40 microseconds rise time, 650 microseconds settling time, and 40% overshoot. Again, the rise times are consistent with the natural frequency in the step response plots.

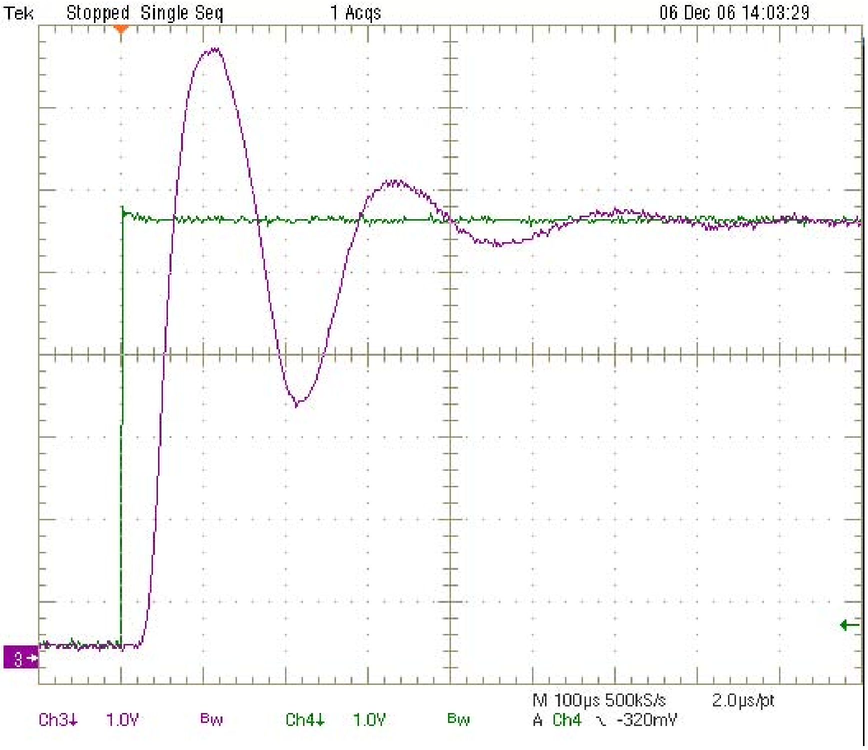

Figure 5-33: 5-volt (125 μrad) step response using optical feedback and controller.

5.3 Dual Axis Characterization and Performance

With the optical and capacitance probe compensator loops designed and running in the elevation and azimuth axes, I was in a position to characterize the performance of both axes running simultaneously under closed-loop control. However, in order to do so I needed two axes' worth of control electronics. Given the myriad of breadboards needed to produce controlled operation of a single axis, I thought it prudent to integrate all the electronic modules onto a single board before performing the dual-axis testing described in this section.

5.3.1 Integrated Electronics Board

While I was working on testing of the breadboard AFSM hardware, Mike Boulet provided a great deal of assistance in the layout and build of the two-axis board. Mike performed a spectacular job with the intricate wire-wrapping on the underside of the board (see layout diagram in Figure 5-34). Despite the presence of hundreds of individual wraps, the board checked out flawlessly after only one diagnostic session, where we discovered a short caused by a small wire clipping laying across two pins.

The board after final build is displayed in Figure 5-35. The major circuit modules are also indicated in the figure. Each compensator module has two I/O channels, while the amplifier compensator has four I/O channels for the sense resistors and power amplifiers. Also, the capacitance probe rotation electronics accepts the four probe inputs and outputs two angles to the probe compensator circuit. Each channel and module are implemented identically to the schematics presented in Chapters 4 and 5, and individual modules are switchable for easy configuration between open- and closed-loop operation and feedback sensors. For diagnostic purposes the controllers may be bypassed entirely.

5.3.2 Axis Coupling Measurement

After testing out the individual board modules to ensure they performed identically to their breadboard counterparts, I proceeded to test the AFSM in dual-axis operation.

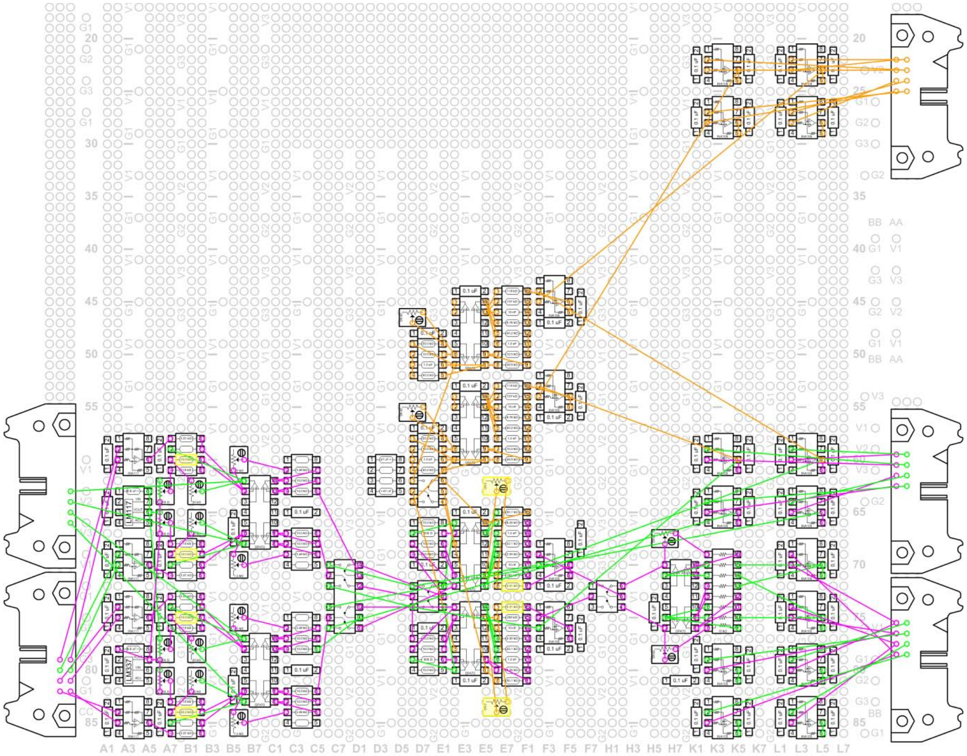

Figure 5-34: Layout diagram of the wire-wrap connections on the underside of the integrated electronics board.

The first step was to characterize the coupling between the two axes. Given the fact that I had designed the mechanical hardware for minimum coupling and that I had successfully closed individual loops around the two axes, I did not expect significant dynamics in these measurements.

The cross-talk plots are shown in Figures 5-36 and 5-37. In each case, the dynamics shown are those of the measured axis while driving the opposite axis in closed-loop using the capacitance probes as feedback. The command signal is ±1 volt–a signal which is deliberately large in order to obtain measurable amplitudes in the off-axis.

As can be seen, the coupling dynamics are basically identical for each axis, as would be expected for the symmetric system. Both plots have a relatively low magnitude of -35 to -40 dB from DC to 100 Hz, which rises to approximately -20 dB just after the resonant peak before falling off at high frequency.

I concluded from the plots that the coupling effects would likely produce small magnitude or phase errors near the resonance, but in general were not cause for

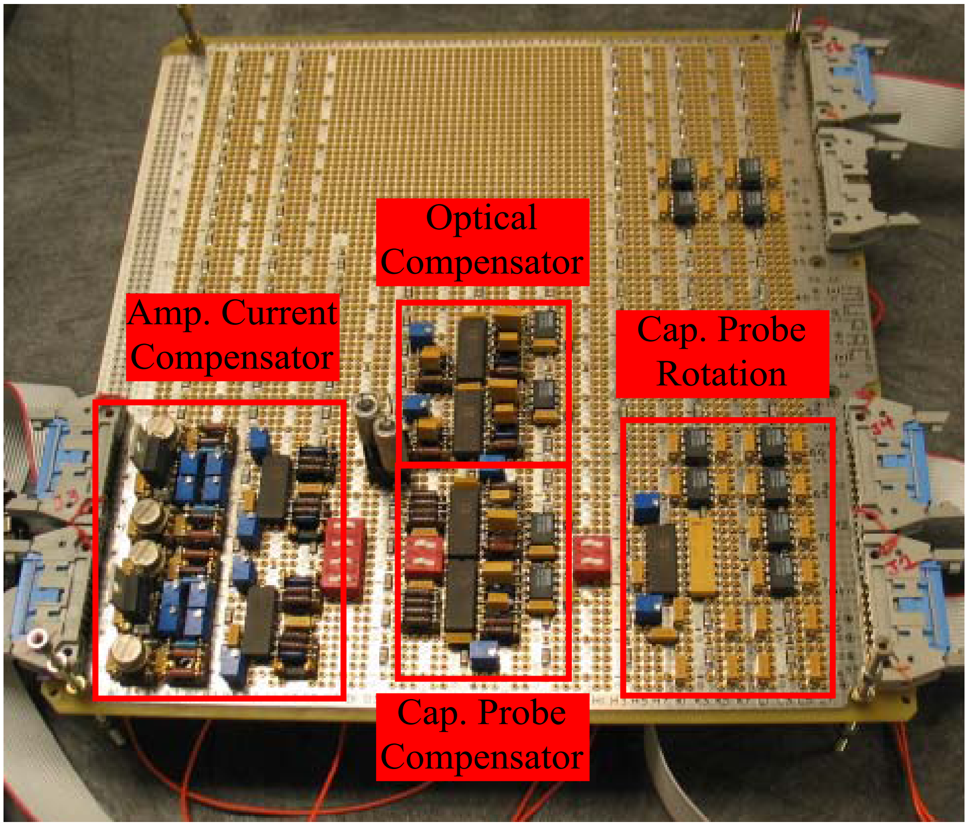

Figure 5-35: Dual-channel wire-wrap integrated electronics board for two-axis testing.

concern. This proved to be the case in the two-axis operation.

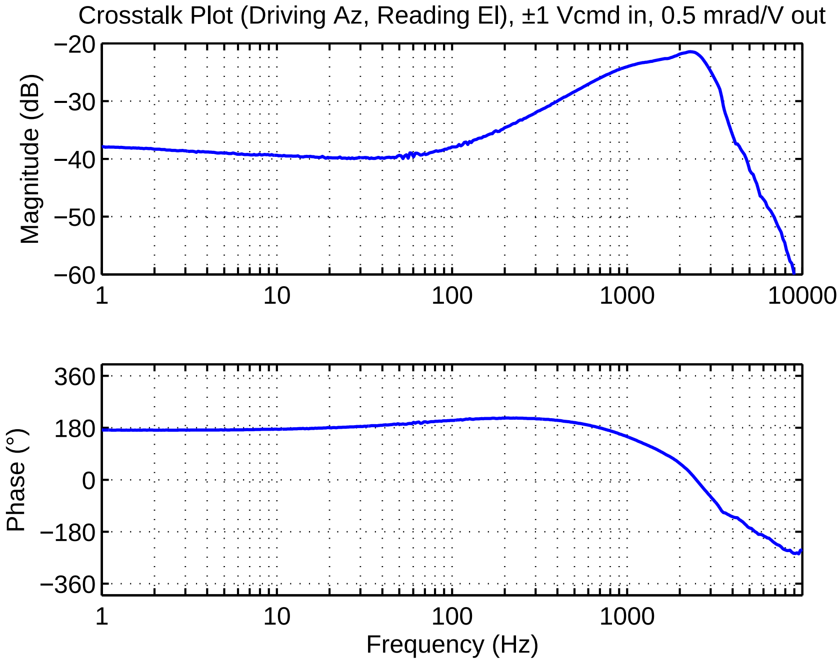

Figure 5-36: Elevation cross-axis plot showing coupling dynamics while driving azimuth.

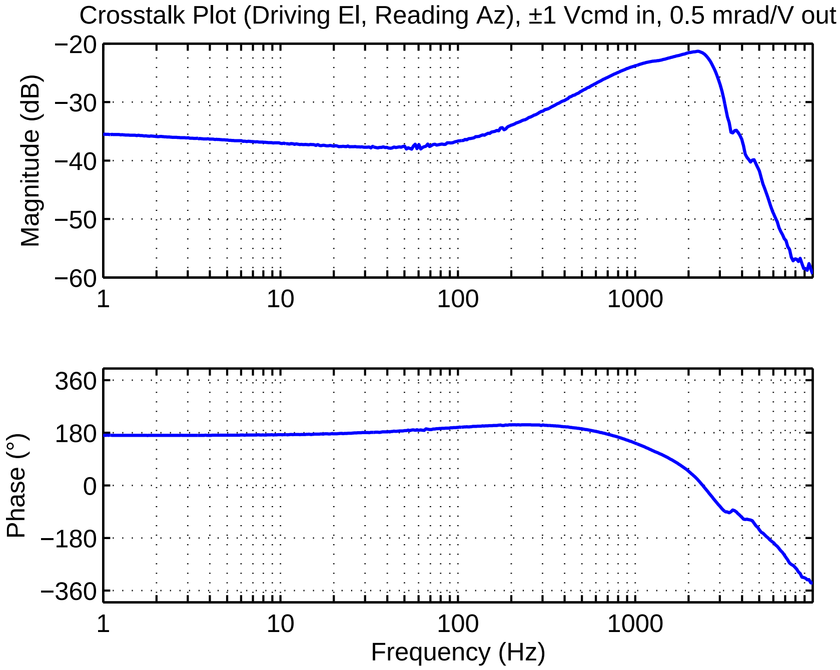

Figure 5-37: Azimuth cross-axis plot showing coupling dynamics while driving elevation.

5.3.3 Dual Axis Operation

During dual-axis testing, I made use of the Matlab xPC system to produce coordinated real-time trajectories as command inputs to the two analog controllers. Here, since the digital system resided outside the control loop, its inherent phase loss was did not inhibit the control performance or stability.

I performed several tests of dual-axis performance to determine accuracy and stability, in both the capacitance probe and optical feedback configurations. As the simplest case, I used the circular trajectory shown in Figure 5-38. The case shown is drawn at 2 kHz using optical feedback, at a radius of 25 μrad.

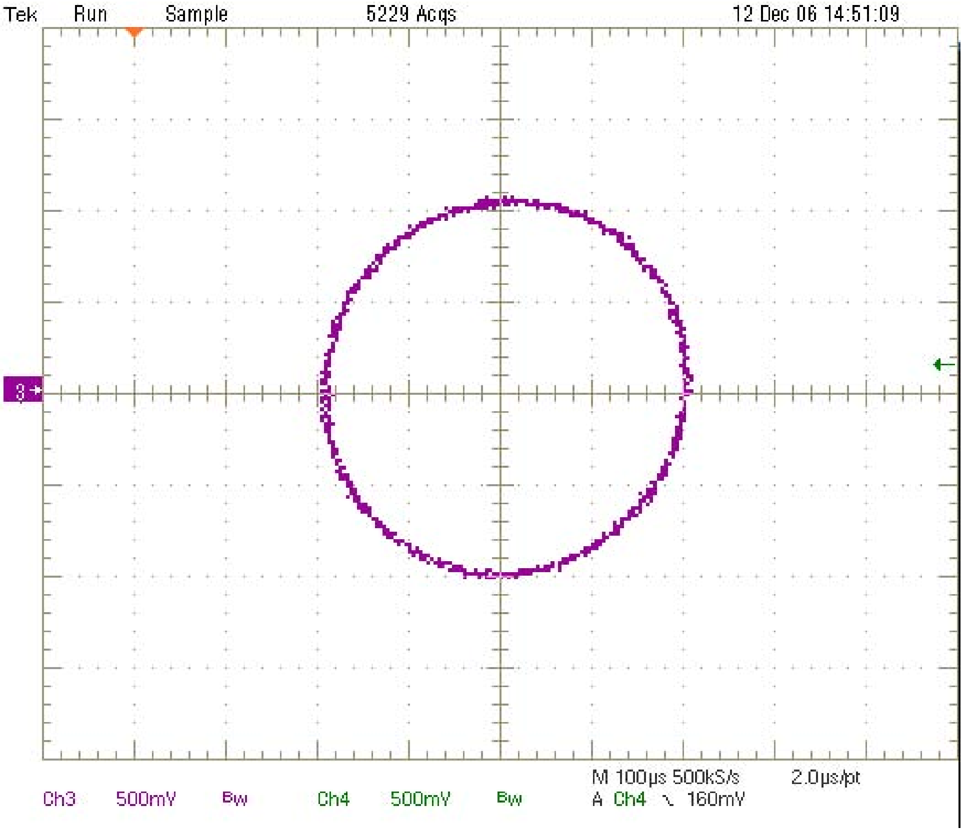

Figure 5-38: Circular trajectory drawn at 2 kHz under dual-axis operation.

Figure 5-39 shows a more complex trajectory drawn using the capacitance probes at 500 Hz. In this case, the trajectory is a Lissajous figure. The frequency and amplitudes of the azimuth and elevation sine waves are chosen to correspond to the Lincoln Laboratory logo: $x = 0.75Asin(ft)$ and $y = Asin(0.75ft)$, where A is the amplitude on the long side, and f is the highest frequency. Note that in this case, the azimuth is drawn at 500 Hz, while elevation is drawn at 375 Hz.

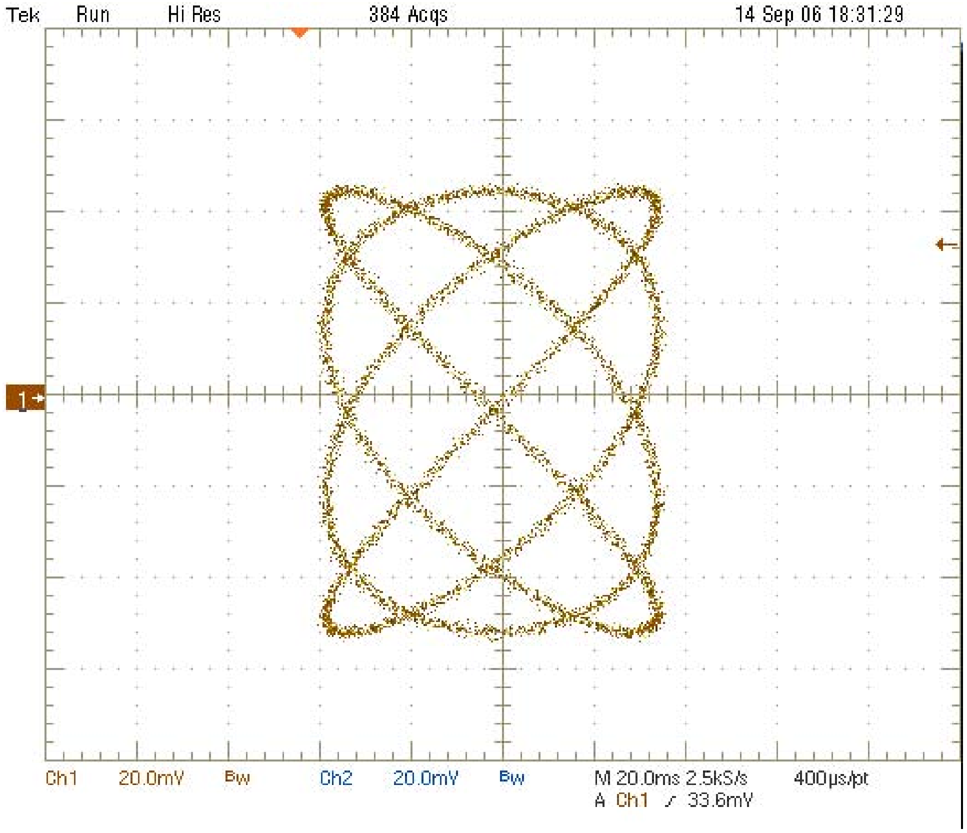

Figure 5-39: Lincoln Laboratory logo lissajous figure drawn at 500 Hz under dual-axis operation.

Unfortunately, I was not able to completely characterize the two-axis performance to my satisfaction because of a hardware failure that occurred late in the testing phase. Yet again, the culprit proved to be the problematic azimuth axis. While exploring a circular trajectory similar to that of Figure 5-38 at increasing frequency, the control loop suddenly went unstable near the closed-loop resonant peak of 4 kHz. I returned to single-axis open-loop operation and ran loop transmission plots for diagnostic purposes. While the elevation axis appeared normal, the azimuth axis exhibited the troublesome dynamics shown in Figure 5-40.

After completing additional closed-loop tests to determine the nature of the new azimuth dynamics, all of which resulted in instability, I noticed that I was receiving virtually no output from the azimuth axis at all. At this point I decided to disassemble

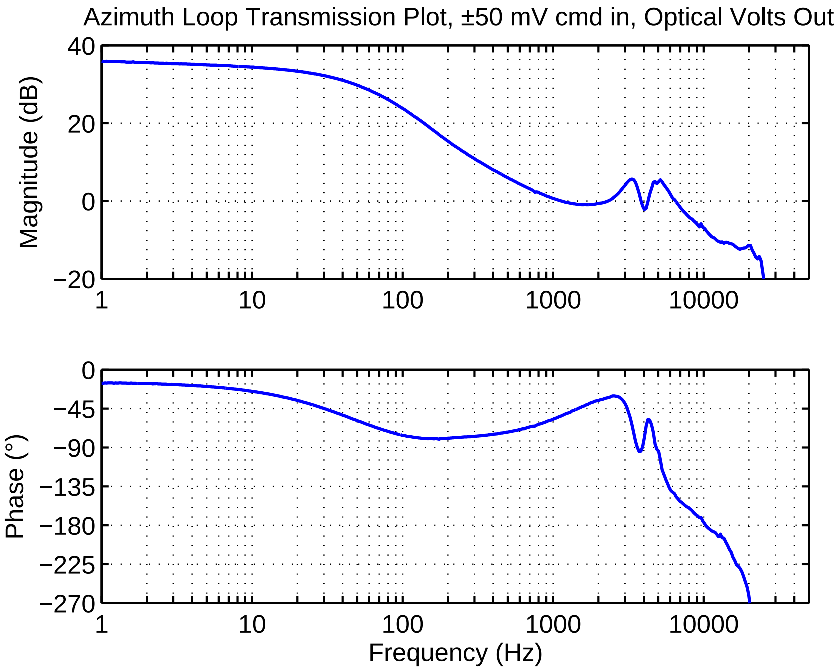

Figure 5-40: Azimuth loop transmission measurement taken just prior to hardware failure.

the azimuth actuators, and discovered that the epoxy bond between the armatures and mirror had failed. A photo of a disassembled actuator half with the magnet and severed armature is given in Figure 5-41.

Since both armatures had been severed, my hypothesis is that one failed first, and subsequently caused the failure of the other due to overdriving in closed-loop operation. Because of the unusual dynamics in the axis, the azimuth loop went unstable many times over the course of my testing, which placed a large number of stress cycles on the epoxy bond joint. These cycles undoubtedly resulted in the initial failure.

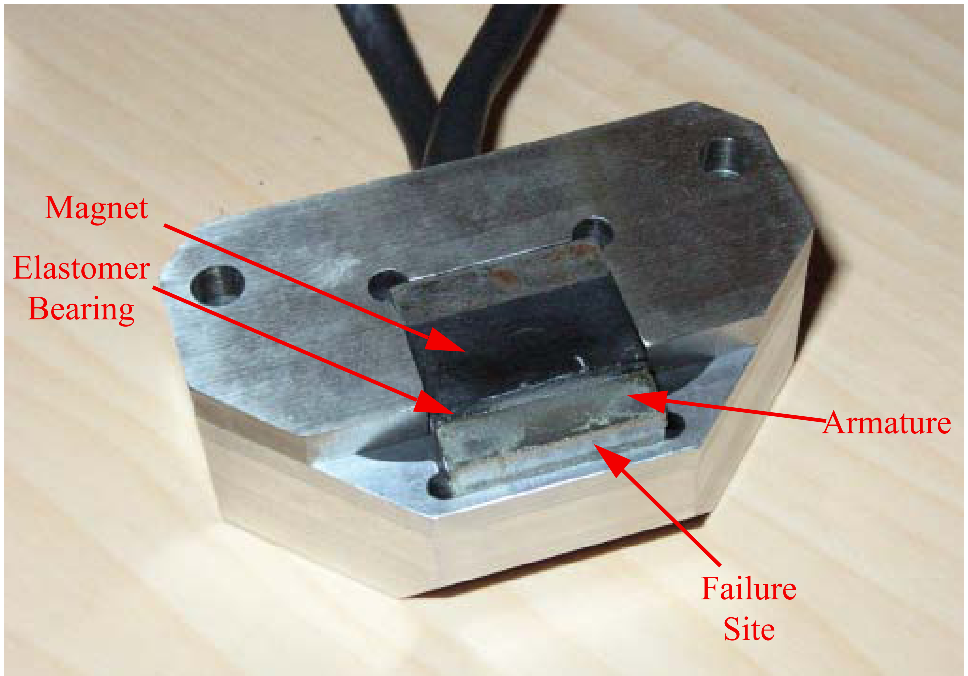

Figure 5-41: One of the failed azimuth axis armatures, shown with actuator half, elastomer bearing and permanent magnet.

5.4 Performance Summary

Collecting the performance data presented in this chapter, we can arrive at a final comparison to the performance goals laid out in Chapter 1. The performance summary is given in Table 5.1, below.

Table 5.1: AFSM Measured Performance Characteristics

Performance Parameter	**Design Goal**	**Measured Performance**	**Comment**
Bandwidth (Single Axis)	5 kHz	10 kHz	1
Bandwidth (Dual Axis)	N/A	4 kHz	2
Angular Range	± 10 mrad	± 3.5 mrad	3
Angular Acceleration	1×10^5 rad/sec^2	1×10^5 rad/sec^2	4
Angular Resolution	8.7×10^{-7} rad	Variable	5
Mirror Aperture	20 mm	20 mm	

Comments:

1. -3 dB, elevation axis only. 5 kHz demonstrated in azimuth axis.
2. Maximum observed prior to azimuth failure
3. See Appendix A
4. Based on 10 kHz bandwidth at 25 μrad amplitude
5. Depends on sensor configuration; quad cell feedback resolution limited by command and drive signal noise floor

As can be seen, the only goal not met by the performance testing was the angular range value of ±10 mrad. As discussed previously, this was due to the unexpectedly high stiffness of the elastomer flexures. This issue is detailed further in Appendix A.

On the positive side, the AFSM exceeded the performance goals in the primary category of small-signal bandwidth. Since for disturbance rejection purposes this is the most important parameter to maximize, exceeding this goal by a factor of two was a significant accomplishment.

With the test results of the AFSM complete, the next task was to evaluate the design against future program needs, and suggest improvements to the technology to meet them. This discussion is the subject of the following chapter.

Chapter 6

Conclusions and Suggestions for Future Work

My work on the AFSM is intended to be a technology demonstrator, and as such, its main usefulness lies in the continuing development and specialization of the technology into real-world applications. The goal of the AFSM project was to build a device in the "middle" of the possible design space, which could be used as a baseline for refinement into mission-specific subsystems. To this end, in this chapter I offer several suggestions for possible excursions to the edges of the design space, focusing on performance, form factor, or flexibility.

6.1 Optimization for Acceleration and Small-Signal Bandwidth

The main goal of the AFSM project was demonstration of a high small-signal bandwidth, and in fact the hardware exceeded the performance goals in this area. However, the program schedule coupled with funding limitations precluded the full exploitation of this capability. Given additional effort, it is reasonable to expect at least a doubling of bandwidth performance to 20 kHz or greater. Do do so, the following characteristics should be explored:

- **Optimization of mirror stiffness and inertia.** As mentioned in Chapter 2, the AFSM mirror's stiffness-to-mass ratio was not optimized in material or geometry. Changing the mirror material to beryllium, silicon carbide, or a composite structure has the potential to reduce the mirror inertia by 50% or more, with resulting gains in acceleration for a given actuator output. In addition, optimization of the structure to remove material but maintain stiffness will result in additional acceleration bonuses. This optimization is preferably done via FEA, but closed-form numerical optimization is also possible. Here, care must be taken not to alter the structural resonance modes, either in shape or frequency, such that they disrupt the controllability or decoupling of the mirror in the actuated DOFs.

  Further benefits would be realized by tailoring the size of the mirror to be no larger than that demanded by the optical application; e.g. use a 20 mm mirror for a 20 mm beam size. Since inertia increases as the square of the mirror radius, the advantages of keeping the mirror small are obvious.

  Finally, minimizing the size and mass distribution of the magnetic armatures will also result in significant performance gains, as discussed in Chapter 2. In particular, relocating the actuators and armatures such that they lie at a smaller radius than the mirror edge would be an interesting design challenge. Here, in addition to the packaging issues, the major tradeoff would involve reducing inertia at the expense of mirror asymmetry. In other words, such a design change would result in a movement of the mirror center of mass away from the center of rotation, which may produce unwanted coupling effects between the controlled axes. However, if extremely high bandwidth performance is warranted, the complexities introduced into the controller design (i.e. the necessity of designing a MIMO system) using such a configuration may be deemed necessary.

- **Maximizing the actuator output force density.** Assuming the pole areas are fixed due to inertia considerations, the actuator output force is proportional to the square of the magnetic flux in the airgap between the pole face and armature. Therefore, producing maximum actuator force is, conceptually, as simple as increasing the airgap flux.

To do this, the first consideration at DC is saturation of the magnetic conductor materials. In the AFSM, the limiting material was the SM-2 HB material used for the armature and backiron pieces. The SM2-based design compromised magnetic performance for structural soundness and manufacturability.

In an advanced design, switching to a material with higher permeability and saturation flux density, such as 50-50 nickel-iron or a cobalt-based alloy such as Supermendur would allow proportional increases in the force density. However, the complication associated with higher operating frequencies is that of power losses in the conductor material due to eddy currents. As discussed in Chapter 3, the normal recourse to mitigate eddy currents is to switch from a solid material to a laminated structure. But, as covered in Chapter 4, it is difficult to structurally attach the armatures to the moving structure without sacrificing either strength or inertia, and a laminated structure complicates this issue even further. Moreover, to be effective magnetically, the laminations must be oriented such that their plane is parallel to the AC magnetic flux paths passing through the armature.

What is required, then, is an armature design that possesses ideal magnetic properties from the standpoints of core loss, flux path, and saturation, yet retains minimum inertia and enough structural soundness to withstand acceleration and impact forces due to over-travel. Clearly, much work can be done to advance in this area–either by exploring improvements in the magnetic properties of a sintered material such as SM-2, or by considering structural improvements to a laminated armature design.

- **Eliminating high-frequency phase loss.** In terms of pure controllability, preserving the phase characteristics of the entire controlled dynamic system at high frequencies will yield maximum bandwidth performance. As discussed in Chapter 5, I identified several factors that contributed to the phase loss characteristics in the AFSM hardware, which in aggregate proved to be the main performance limitation. In a next-generation design, addressing these factors directly in the infancy of the design will be critical. I repeat them here, but with a focus on improvements over the AFSM measurements.

- **AFSM suspension mode.** In theory, it is possible to design the mirror and suspension to be stiff enough to place the desired crossover at a lower frequency than the first suspension mode, thus preserving 180 degrees of phase loss. However, there are several complications associated with this approach. The most significant of these is the difficulty of actually making a macroscopic structure with a suspension mode at such a high frequency (in the present case, greater than 20 kHz). Such a structure may be realizable at a micro- or nano-scale; however, the optical system requirements (most notably, the beam aperture size) may preclude such miniaturization. A second complication is the control effort required to actuate the system. The control and/or power electronics required to drive a plant of such high stiffness would be characterized by very high gains and power consumption, and would also be prone to noise due to their sensitivity.
- **Sensor phase lag.** As I discovered in my experiences with the capacitance probes, excessive phase delay in the position sensing system can prove to be devastating when attempting to maximize bandwidth. The solution to this problem is conceptually very simple: Use a sensor whose bandwidth exceeds the desired system bandwidth by a large margin (i.e. a factor of ten or more), and pay careful attention to the phase characteristics of the unit–particularly with regard to electronic filtering. If the bandwidth margin cannot be met due to technological and/or cost limitations, it may be prudent to explore modifications in the sensor low-pass filters to improve their phase response. In the case of the capacitance probes, this could be done by reducing the order of the Butterworth filter, or switching to an alternative design which preserves phase at the expense of magnitude, such

as a Bessel or Chebyshev design.

- **Elastomer bearing dynamics.** In retrospect, using the elastomer bearings as kinematic flexures both aided and hindered the final performance numbers. On the positive side, they eliminated spurious flexure modes and provided significant natural damping to the system. However, they did contribute to the overall phase delay in a less predictable manner than an equivalent metallic flexure. Provided the issues of low modal damping and the introduction of spurious modes can be addressed, an advanced design may benefit from the use of metallic flexures over the elastomers used in the AFSM
- **Magnetic hysteresis and eddy current losses.** Hand in hand with the issues in the armature design discussed above, the electromagnetic losses in the actuator contribute to the overall phase loss in the system, and therefore must be minimized. Again, this is accomplished by choosing materials and construction with small hysteresis loops and minimum core losses.
- **Electronic system phase loss.** As the mechanical characteristic frequencies of the plant increase, the phase of the electronic system must also be considered. For example, the power amplifier current loops in the AFSM system displayed a sharp resonance at 40 kHz, as discussed in Chapter 4. For a mechanical system bandwidth of 20 or 30 kHz, such a resonance would likely be unacceptable and would require modifications to the electronic system.

6.2 Optimization for Power Consumption, Mass and Form Factor

In space or airborne programs, limiting hardware mass, envelope, and power consumption is often a primary concern. In the AFSM design, however, I paid little attention to these aspects, instead choosing to focus on functionality and manufacturing simplicity. As such, the design may be improved dramatically to optimize these characteristics, should program requirements demand it.

Here, rather than pushing for maximum acceleration and bandwidth, the designer would begin with a fixed mirror aperture and modest performance specifications (for example, a bandwidth in the 500 to 1000 Hz range), and then take all reasonable steps to design a package of minimal form factor that still meets this performance. Interestingly, although it would seem that this approach would require very different solutions than a high-performance design, this is not necessarily the case. For example, reducing the mirror moving mass yields better performance, but naturally it also results in a lighter-weight design.

Some key considerations for a low-mass, low-power, and minimal envelope design include:

- **Optimizing mirror mass over stiffness.** Here, the structural resonances of the mirror are less important, since the bandwidth requirements are lower. Therefore, it may be possible to sacrifice stiffness in the mirror design to achieve lower mass, both static and dynamic.

- **Maximizing actuator force density.** Here lies another commonality with the high-performance design approach. The fixed bandwidth requirements dictate the actuator force that must be provided for a given mirror mass. Once the force is known, the actuator must be designed to provide this force with minimum weight and envelope. Accordingly, designing for a high force density is desired. The considerations here are identical to the previous section, but the emphasis is on mass reduction rather than performance.

- **Reducing coil mass.** As designed, the AFSM is rather inefficient in its usage of the copper coil volume. Only about 30-40% of the coil is actually used to source flux into the magnetic components. The remainder is needed simply to provide a return path for the current. An improved design may be able to take advantage of a larger percentage of the coil volume through more effective geometric allocations.

In addition, an optimized magnetic design should include a study to minimize the coil inductance, because doing so yields improvements in both mass and reactive power consumption. Consulting Equation 3.50, we can see that the inductance is a function of the number of coil turns, pole areas, and airgap lengths in the magnetic circuit. Here, it is clear that the biggest gains are realized via reductions in the number of coil turns N. Doing so lowers the amount of copper needed in the coils, thereby saving mass[1], and also lowers the inductance by the square of the reduction in turns, at the expense of higher power dissipation.

- **Maximize the use of elastomeric flexures.** With the system plant dynamics as a secondary concern, and if environmental specifications allow, an optimized design should take advantage of the mass and space reductions afforded by elastomeric flexures instead of heavier and spatially cumbersome metallic flexures.

- **Make use of lightweight materials throughout.** This is an obvious step in any mass-optimized design. After considering the moving mirror structure, the designer should then proceed to examine the static components of the design, especially the housing and support structure of the device. The use of lightweight aluminum, beryllium, or carbon-fiber will yield substantial mass savings.

At the time of this writing, an effort to perform such an optimization of mass and form factor is underway.

---

[1] This is not a one-to-one reduction, however, because lowering the number of turns generally requires the use of larger gauge wire to avoid overheating.

6.3 Optimization for Optical Flexibility

Despite the fact that the actual design effort for the AFSM hinges primarily on electromechanical considerations, the designer must still keep in mind that the device serves solely to add functionality to the optical system as a whole. As such, it is incumbent upon the designer to consider the optical system requirements first, and if warranted, the mirror design should be altered accordingly.

6.3.1 Reflective Mirrors

Reflective mirrors such as those used in the HBSM and AFSM designs obtain maximum functionality and optical flexibility when they are provided with a clear field of view over as wide an angle as possible, relative to the mirror normal. Any obstructions above the mirror plane will reduce this angle, resulting in decreased flexibility.

The actuators in the AFSM are of a symmetric design for ease of manufacturability and assembly, but the downside of this choice is that the mirror viewing angle is severely limited. For example, when using a 15 mm optical beam diameter, the available included angle is only 60 degrees. This was enough for me to perform optical tests, but may be unacceptable for system design.

The solution is to reduce the actuator profile in order to minimize viewing angle obstructions. This may be done by making the actuators smaller, asymmetric, or both. In addition, an effort to increase performance by placing the actuators and armatures at a smaller radius than the mirror edge may also result in greater optical flexibility.

6.3.2 Refractive Lenses

In addition to traditional reflecting elements, it may prove feasible to build a fast steering mechanism around a refractive optical element, i.e. a lens. Here, since the optical beam must pass through the device, the actuators and kinematic supports should be arrayed around the edges. The AFSM architecture lends itself nicely to such a configuration, especially since the beam angle requirements for a refractive element are generally less than for a reflective mirror. The only complication lies with the axial flexure and capacitance probes, which are presently located within the foreseen beam volume. However, the axial flexure may be eliminated in favor of active axial translation control using the actuators, and the sensors may be sensors may be simply relocated without great difficulty, or even eliminated entirely if sufficient optical feedback is available to the controller.

6.4 Conclusions

In this thesis, I have documented a new advanced fast steering mirror design. My work provides a solid foundation for many future applications requiring high-performance stabilized optical platforms. It should be noted that considerable flexibility exists in the AFSM technology, making it amenable to applications requiring one or a combination of high bandwidth, low mass, minimum cost, and modest power consumption. The design demonstrates excellent bandwidth performance as well as relatively low cost. I was able to complete design, fabrication, and testing in about ten months. Given additional time and resources, this technology has considerable potential to advance the state of the art in military, scientific, and civilian optical applications. In the event that such an advancement takes place, I will be humbly gratified.

Appendix A

AFSM Magnetic Analysis Via Maxwell's Equations

I argued in Chapter 3 that magnetic circuit theory was the most direct approach to a reasonably accurate magnetic analysis of machines such as the AFSM. Here I show that the results obtained are identical to those using Maxwell's equations directly, as long as the assumptions used in the analysis are the same.

It should be noted that to obtain a solution with higher accuracy, it is necessary to relax the fundamental assumptions of Chapter 3. The consequence is that the contours and surfaces of integration may take on arbitrary shapes, which may not be describable by analytical functions. The same is true for the various vector field quantities in the integrands. These complications generally result in a full-field formulation of Maxwell's equations which is unsolvable in closed-form. In such a case, a numerical technique such as finite-element analysis may be used to obtain a solution, which is still an approximation, but a much better one than the highly restrictive magnetic circuit solution.

The Maxwell relations are given in Chapter 3, but I repeat them here for convenience:

$$
\begin{aligned}
\oint_C \mathbf{H}\cdot d\mathbf{l} &= \int_S \mathbf{J}\cdot\mathbf{n}\,da + \int_S \frac{\partial \mathbf{D}}{\partial t}\cdot\mathbf{n}\,da && \text{(A.1)}\\
\oint_C \mathbf{E}\cdot d\mathbf{l} &= -\frac{d}{dt}\int_S \mathbf{B}\cdot\mathbf{n}\,da && \text{(A.2)}\\
\oint_S \mathbf{D}\cdot\mathbf{n}\,da &= \int_V \rho dv && \text{(A.3)}\\
\oint_S \mathbf{B}\cdot\mathbf{n}\,da &= 0 && \text{(A.4)}
\end{aligned}
$$

To evaluate the equations with respect to the AFSM configuration, we must define the closed line contours and associated (open) surfaces of integration for use in Equation A.1, as well as closed surfaces for the integrations using Equation A.4. Here Equations A.2 and A.3 are not necessary, but they do come into play in the terminal voltage computation, which has been derived previously in Chapter 3.

The lines and surfaces of integration are given in Figure A-1. Three closed line contours, A, B, and C, enclose three open surfaces, also labelled A, B, and C. The actuator coils penetrate these surfaces, and their current densities and cross-sectional areas intersecting the surfaces[1] are given by J_1, A_1 and J_2, A_2 respectively (note that as before, the direction of positive current is defined as into the page). In addition, two *closed* surfaces, labelled 1 and 2, enclose the armature and permanent magnet, respectively.

We begin the analysis by applying the same fundamental assumptions given in Chapter 3 of infinitely permeable soft magnetic conductors, no fringing, no leakage, and flux densities normal to all pole faces. Next, we evaluate the contour and surface integrals A, B, and C in the directions shown in Figure A-1 using Ampere's Law, Equation A.1, with zero electrical displacement **D**. Beginning with Contour A and Surface A, we have

---

[1] The coil area is defined here as the total cross-sectional electrical conductor area ("copper area") only. It does not include other contributions to the physical cross-sectional coil area, such as insulation, packing factor, etc.

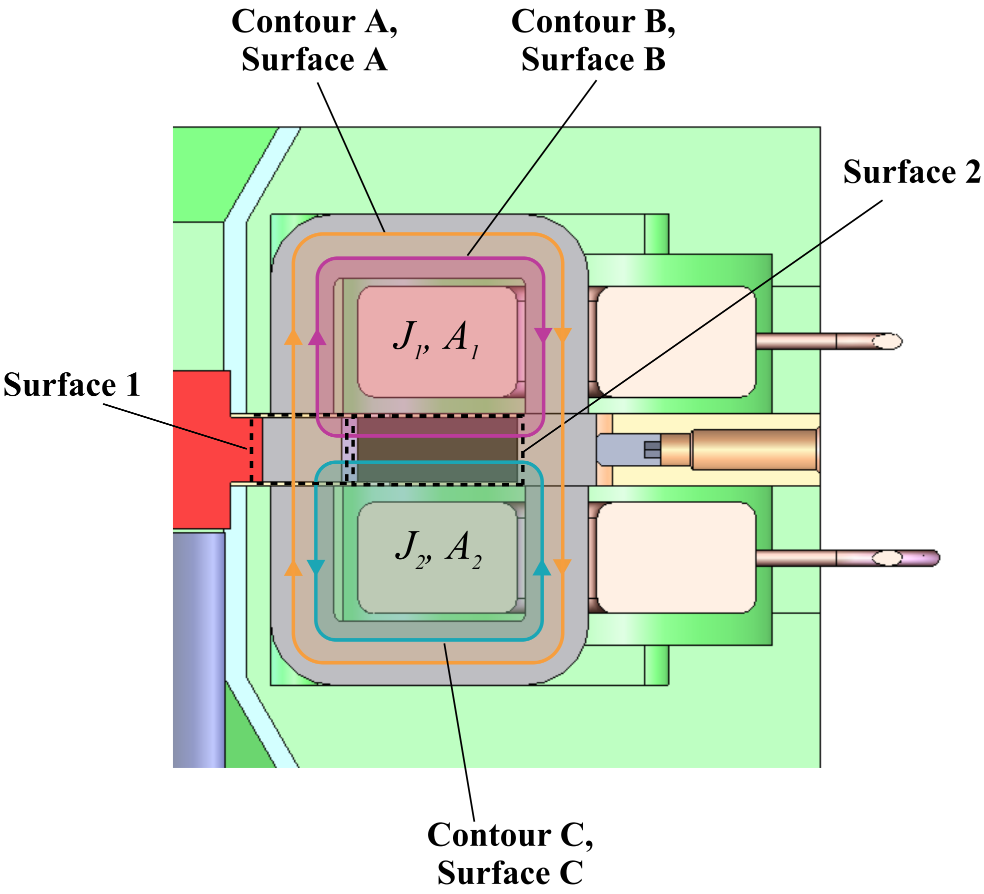

Figure A-1: AFSM actuator with integration contours and surfaces for the analysis using Maxwell's equations.

$$\begin{aligned} \oint_A \mathbf{H} \cdot d\mathbf{l} &= \int_A \mathbf{J} \cdot \mathbf{n} \, da \\ -H_l g_l + H_u g_u &= J_1 A_1 + J_2 A_2 \end{aligned} \tag{A.5}$$

where H_u and H_l are the magnetic fields in the upper and lower airgaps, respectively, and g_u and g_l are the upper and lower gap lengths. Because of the infinite core permeability, there are no magnetic fields in the remaining portions of the contour. Proceeding similarly with Contours B and C, we have

$$H_m l_m + H_n g_n + H_u g_u = J_1 A_1 \tag{A.6}$$

$$H_m l_m + H_n g_n + H_l g_l = -J_2 A_2 \tag{A.7}$$

where H_m and l_m are the permanent magnet field and length, and H_n and g_n are the nonworking airgap field and length. Note that in Equation A.7 the contour direction is out of the page, making the sign of the current term on the right-hand side negative.

Next, we integrate over the closed surfaces 1 and 2 using Gauss's Law, Equation A.4. Recalling that the pole areas of the magnet and the armature upper, lower, and nonworking airgaps are all equal to a_p, we write for Surface 1:

$$\begin{aligned}
\oint_1 \mathbf{B} \cdot \mathbf{n}\, da &= 0 \\
B_u a_p + B_l a_p - B_n a_p &= 0 \\
B_u + B_l - B_n &= 0 \qquad \text{(A.8)}
\end{aligned}$$

where B_u, B_n, and B_l are the flux densities in the upper, nonworking, and lower airgaps, respectively. Note that the direction of positive permanent magnet flux is out of the magnet north pole and *into* the armature nonworking airgap surface, which is in the opposite direction of the surface outward normal. This makes the sign of the B_n term negative. For Surface 2, the result is simply

$$B_n - B_m = 0 \tag{A.9}$$

where B_m is the flux density into the south pole of the magnet. Note that in formulation Equations A.8 and A.9 we have implicitly made use of the assumption of no leakage, which effectively constrains the flux densities through the inactive surfaces of the armature and magnet to be zero.

Ultimately we seek to solve for the flux densities B, which when multiplied by the pole areas a_p give the fluxes Φ. Similarly to Chapter 3, we recognize Equation A.5 as a linear combination of Equations A.6 and A.7, so we are left with four equations in eight unknowns. To obtain the remaining four equations necessary for a solution, we make use of the constitutive laws for the permanent magnet and the three airgaps:

$$\begin{aligned}
B_m &= \mu_o(H_m + M_o) && \text{(A.10)}\\
B_u &= \mu_o H_u && \text{(A.11)}\\
B_l &= \mu_o H_l && \text{(A.12)}\\
B_n &= \mu_o H_n && \text{(A.13)}
\end{aligned}$$

Solving for the upper gap flux density B_u, starting with Equation A.11 and using Equations A.6, A.10, A.9 and A.13 gives

$$\begin{aligned}
B_u &= \mu_o H_u\\
&= \frac{\mu_o}{g_u}\left(J_1A_1 - H_m l_m - H_n g_n\right)\\
&= \frac{\mu_o}{g_u}\left(J_1A_1 - \frac{B_m}{\mu_o} l_m + M_o l_m - H_n g_n\right)\\
&= \frac{\mu_o}{g_u}\left(J_1A_1 + M_o l_m - H_n(l_m + g_n)\right) && \text{(A.14)}
\end{aligned}$$

Using Equations A.13, A.8, A.12, A.5, and A.11, we solve for H_n as

$$\begin{aligned}
H_n &= \frac{B_n}{\mu_o}\\
&= \frac{B_u + B_l}{\mu_o}\\
&= \frac{1}{\mu_o}\left(B_u + \frac{\mu_o}{g_l}(H_u g_u - J_1A_1 - J_2A_2)\right)\\
&= \frac{B_u}{\mu_o}\left(1 + \frac{g_u}{g_l}\right) - \frac{1}{g_l}(J_1A_1 - J_2A_2) && \text{(A.15)}
\end{aligned}$$

Combining Equations A.14 and A.15 gives the result for B_u as

$$B_u = \frac{\mu_o g_l(M_o l_m + J_1 A_1) + \mu_o(l_m + g_n)(J_1 A_1 + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} \tag{A.16}$$

The flux densities B_l and B_m are found in a similar manner to be

$$\begin{aligned} B_l &= \frac{\mu_o g_u(M_o l_m - J_2 A_2) - \mu_o(l_m + g_n)(J_1 A_1 + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} && \text{(A.17)} \\ B_m &= \frac{\mu_o g_l(M_o l_m + J_1 A_1) + \mu_o g_u(M_o l_m + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} && \text{(A.18)} \end{aligned}$$

Multiplying A.16, A.17, and A.18 by a_p gives the fluxes as

$$\begin{aligned} \Phi_u &= \frac{\mu_o a_p g_l(M_o l_m + J_1 A_1) + \mu_o a_p(l_m + g_n)(J_1 A_1 + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} && \text{(A.19)} \\ \Phi_l &= \frac{\mu_o a_p g_u(M_o l_m - J_2 A_2) - \mu_o a_p(l_m + g_n)(J_1 A_1 + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} && \text{(A.20)} \\ \Phi_m &= \frac{\mu_o a_p g_l(M_o l_m + J_1 A_1) + \mu_o a_p g_u(M_o l_m + J_2 A_2)}{g_u g_l + (l_m + g_n)(g_u + g_l)} && \text{(A.21)} \end{aligned}$$

The above expressions are equivalent to the magnetic results given in Chapter 3, Equations 3.33 through 3.35. To see this, we recall the definitions for the gap reluctances and magnet MMF, which were also defined in Chapter 3:

$$\begin{aligned} g_u &= x_o - x && \text{(A.22)} \\ g_l &= x_o + x && \text{(A.23)} \\ \mathcal{R}_{mag} &= \frac{l_m + g_n}{\mu_o a_p} && \text{(A.24)} \\ \mathcal{R}_u &= \frac{x_o - x}{\mu_o a_p} && \text{(A.25)} \\ \mathcal{R}_l &= \frac{x_o + x}{\mu_o a_p} && \text{(A.26)} \\ \Psi_m &= M_o l_m && \text{(A.27)} \end{aligned}$$

Additionally, we recognize that since the upper and lower coils are of identical construction and are wired in series, the current density terms are equivalent to the coil MMF via the expression

$$J_1 A_1 = J_2 A_2 = 1/2\Phi_c \tag{A.28}$$

Combining the above relations with Equation A.19, we can rewrite it as

$$\begin{aligned}
\Phi_u &= \frac{\mu_o a_p(x_o + x)(\Psi_m + 1/2\Psi_c) + \mu_o a_p(l_m + g_n)\Psi_c}{x_o^2 - x^2 + (l_m + g_n)(2x_o)} \\
&= \frac{(\mu_o a_p)^2\mathcal{R}_l(\Psi_m + 1/2\Psi_c) + (\mu_o a_p)^2\mathcal{R}_{mag}\Psi_c}{(\mu_o a_p)^2\mathcal{R}_u\mathcal{R}_l + (\mu_o a_p)\mathcal{R}_{mag}(\mu_o a_p)(\mathcal{R}_u + \mathcal{R}_l)} \\
&= \frac{\mathcal{R}_l(\Psi_m + 1/2\Psi_c) + \mathcal{R}_{mag}\Psi_c}{\mathcal{R}_u\mathcal{R}_l + \mathcal{R}_{mag}(\mathcal{R}_u + \mathcal{R}_l)} \qquad (A.29)
\end{aligned}$$

which is identically Equation 3.34. Similarly, the other two relations are shown to be

$$\begin{aligned}
\Phi_l &= \frac{\mu_o a_p(x_o - x)(\Psi_m - 1/2\Psi_c) - \mu_o a_p(l_m + g_n)\Psi_c}{x_o^2 - x^2 + (l_m + g_n)(2x_o)} \\
&= \frac{(\mu_o a_p)^2\mathcal{R}_u(\Psi_m - 1/2\Psi_c) - (\mu_o a_p)^2\mathcal{R}_{mag}\Psi_c}{(\mu_o a_p)^2\mathcal{R}_u\mathcal{R}_l + (\mu_o a_p)\mathcal{R}_{mag}(\mu_o a_p)(\mathcal{R}_u + \mathcal{R}_l)} \\
&= \frac{\mathcal{R}_u(\Psi_m - 1/2\Psi_c) - \mathcal{R}_{mag}\Psi_c}{\mathcal{R}_u\mathcal{R}_l + \mathcal{R}_{mag}(\mathcal{R}_u + \mathcal{R}_l)} \qquad (A.30) \\
\Phi_m &= \frac{\mu_o a_p(x_o + x)(\Psi_m + 1/2\Psi_c) + \mu_o a_p(x_o - x)(\Psi_m - 1/2\Psi_c)}{x_o^2 - x^2 + (l_m + g_n)(2x_o)} \\
&= \frac{(\mu_o a_p)^2\mathcal{R}_l(\Psi_m + 1/2\Psi_c) + (\mu_o a_p)^2\mathcal{R}_u(\Psi_m - 1/2\Psi_c)}{(\mu_o a_p)^2\mathcal{R}_u\mathcal{R}_l + (\mu_o a_p)\mathcal{R}_{mag}(\mu_o a_p)(\mathcal{R}_u + \mathcal{R}_l)} \\
&= \frac{\mathcal{R}_l(\Psi_m + 1/2\Psi_c) + \mathcal{R}_u(\Psi_m - 1/2\Psi_c)}{\mathcal{R}_u\mathcal{R}_l + \mathcal{R}_{mag}(\mathcal{R}_u + \mathcal{R}_l)} \qquad (A.31)
\end{aligned}$$

which are the same as Equations 3.35 and 3.33, respectively. The remainder of the analysis to determine the coil terminal voltage and actuator force output is then identical to that performed in Chapter 3.

Appendix B

AFSM Stiffness Investigation

The one performance specification not met by the AFSM hardware was the angular range goal of ±10 mrad. I determined during testing that the mirror suspension stiffness was higher than what I had calculated during the design. However, the high stiffness did not greatly affect performance in other areas, and did not preclude a successful controller design. I therefore set the issue aside with the intention of investigating it after completion of testing. I discuss my investigation activity and conclusions in this appendix.

B.1 Stiffness Verification

The first task, part of which I performed during the test phase, was to verify the high stiffness. To do this, I employed three separate methods:

- **Calculation of stiffness at resonance.** The quickest way to determine the experimental stiffness was to simply derive it from the measured suspension mode frequency using the expression $k \approx 4\pi^2 f^2 I$, where k is the stiffness, f is the resonant frequency, and I is the mirror inertia. The expression is approximate because of the modal damping, which lowers the frequency slightly relative to the theoretical undamped case. Using a frequency of 2.1 kHz and the known mirror inertia of $4.1x10^{-6}$ kg-m^2 gives a stiffness at resonance of 721 N-m/rad.

However, based on with work of Barton [17] and Cuff [6], the stiffness at resonance is distinct from the DC stiffness, which was my main metric of interest in the investigation (the only reason for this, however, was to establish a common baseline for all tests and calculations). To determine the DC stiffness, is necessary to examine the AFSM dynamic plots to determine the relationship between stiffness and frequency. Referencing Figure 5-8 (repeated below as Figure B-1 for convenience), the stiffening effect can be seen in the magnitude plot prior to hitting the resonant peak. Note that since the AFSM Bode plots measure *compliance* (that is, the ratio of position output to force input), an increasing *stiffness* (ratio of force to position) manifests itself as a decay in magnitude.

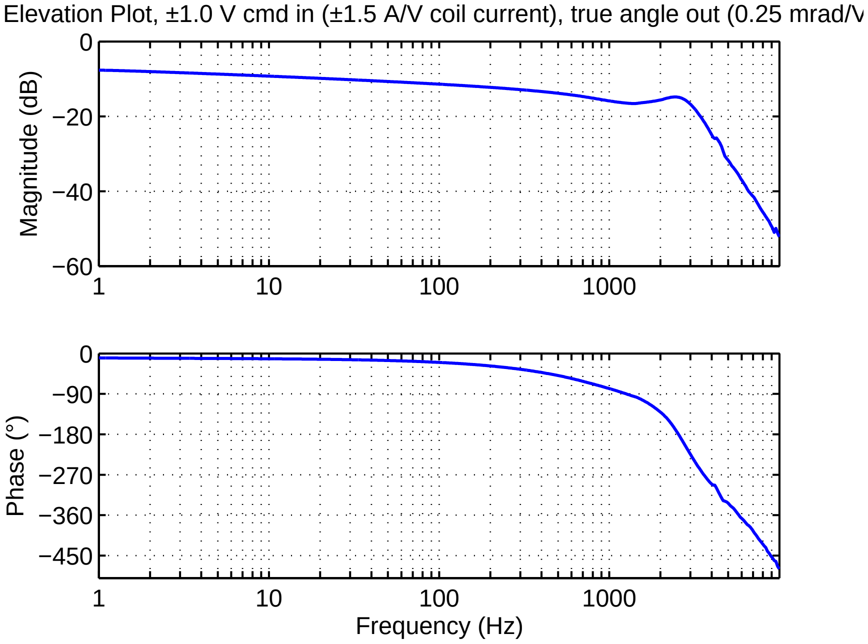

Figure B-1: AFSM frequency response for determining rubber stiffness characteristics.

I determined the slope of the magnitude decay to be approximately -2.1 dB per decade. Using this slope I extrapolated the stiffness at resonance to DC to obtain a value of 322 N-m/rad.

- **Measurement of DC stiffness via actuator current draw.** In the next verification method, I applied various DC current levels to opposing actuator pairs in both axes, and measured the resulting angular displacements via the capacitance probes. I then used the derated theoretical force output from my actuator calculations to derive the DC stiffness. The derating was necessary because the theoretical calculations did not account for the leakage and fringing effects of the magnetics, and therefore resulted in unrealistically high forces. Referencing a FEA magnetic analysis performed by Lu for a similar actuator [16], I found that the actual flux values in the upper and lower working airgaps were about 82% of theoretical due to the loss effects. Since force is proportional to the square of flux, the derating value for force is $(0.82)^2$, or 67%.

  With the force output properly derated, the stiffness result was similar to the value obtained using the frequency measurements: 315 to 350 N-m/rad.

- **Direct Measurement of DC Stiffness.** My final measurement was a rather crude one, but it helped to provide additional confidence in the prior two measurements. I simply used a linear force gauge to apply a known moment to the mirror, and again used the capacitance probes to measure the resulting angular output. By using the force gauge I was able to circumvent inaccuracies in the actuator force calculations; however, there was still uncertainty in the applied moment because I could not determine the precise distance at which I applied the force relative to the center of the mirror. Allowing for a wide error in the moment arm, the results of the test fell in the range of 300 to 400 N-m/rad.

As a result of these three activities, I was fairly confident that the actual DC stiffness value was in the range of 300 to 350 N-m/rad. In contrast, the design value I calculated for the AFSM was 20.5 N-m/rad, which is 14.5 times lower than the measurement.

It should be noted that the design value of 20.5 N-m/rad is a composite stiffness with several contributors: Two elastomer pads in shear along the actuated axis, two pads in torsion along the opposing axis, bending of the axial flexure, and the negative stiffness of the magnetic actuators. I calculated the latter at -99.5 N-m/rad, and the remaining positive stiffnesses at 120 N-m/rad. Although the magnetic stiffness is subject to uncertainty because it is derived from theoretical calculations, notice that even if its value is zero, the remaining stiffness of 120 N-m/rad is still too low to account for the experimental measurements.

B.2 Hypotheses for Stiffness Anomaly

With the high stiffness verified, the next step was to develop hypotheses for the anomaly. After some brainstorming I settled on the following possibilities:

1. **Improper design of elastomer bearings.** My first guess was that I had simply made an error in the stiffness calculations during my initial design–either by using incorrect equations, material or geometry values, or both.

2. **Elastomer over-constrained by surrounding parts.** The stiffness calculations from Lindley [14], which are based on shape factor, rely on the unloaded surfaces of the elastomer bearing being free to bulge into the surrounding volume without constraint. Because of the near incompressibility of the material, such unintended constraints would result in a significant stiffness increase over the design value.

   In the AFSM design, the gap between the elastomer bearing unloaded surfaces and the upper and lower actuator halves was only 300 microns, the same as the working airgap of the actuator. I suspected that when preloaded by the permanent magnet, the bearing may have bulged and interfered with the actuator half surfaces, resulting in the increased stiffness.

3. **Elastomer placed in compression due to AFSM angular motion.** In my original analysis, I accounted for only the shearing displacement of the elastomer bearing resulting from the mirror angular rotation. However, since the mirror has depth and is not a flat plane, its rotation also results in normal displacements of the bearing, placing one edge in compression and the other in tension. At the maximum displacement of 10 milliradians, the normal displacement is about ±19 microns at the edges. Because of the high normal stiffness of the elastomer,

the compressive and tensile reaction forces resulting from these displacements increase the overall contribution of the bearing to the total angular stiffness.

4. **Incorrect elastomer material properties.** As noted in Chapter 2, one of the drawbacks with using elastomers in a precision application is that the material properties may have greater variability than metals. I ordered the neoprene used in the AFSM design from a next-day supply company with no certification of conformance to the advertised material properties. Furthermore, I performed no material testing on the elastomer myself prior to installing it in the AFSM. I therefore suspected that the material itself may have accounted for the anomalous stiffness value.

B.3 Testing of Hypotheses

With reasonable hypotheses established, I proceeded to test each one through various methods, as described in the following sections.

B.3.1 Improper Design

The verification of the design methodology was rather straightforward. I double-checked the equations used by Lindley, Barton, and Cuff, and confirmed that they were identical. I also verified proper calculation of the shape factor using Figure 2-9.

Interestingly, I was able to verify the accuracy of one aspect of the design because of the azimuth failure during my testing; namely, the torsional loading of the off-axis pads. I calculated the contribution of the torsional stiffness to the overall mirror angular stiffness to be quite low–about 9% of the total elastomer stiffness. After the azimuth failure, I was able to verify the minimal effect of the torsion loading by running the elevation axis with the azimuth actuators removed–which removed the torsional bearing contribution. The result is shown in Figure B-2, below. It is virtually unchanged from the plots in Chapter 5.

As a final check, I verified that the geometry of the elastomer bearings matched

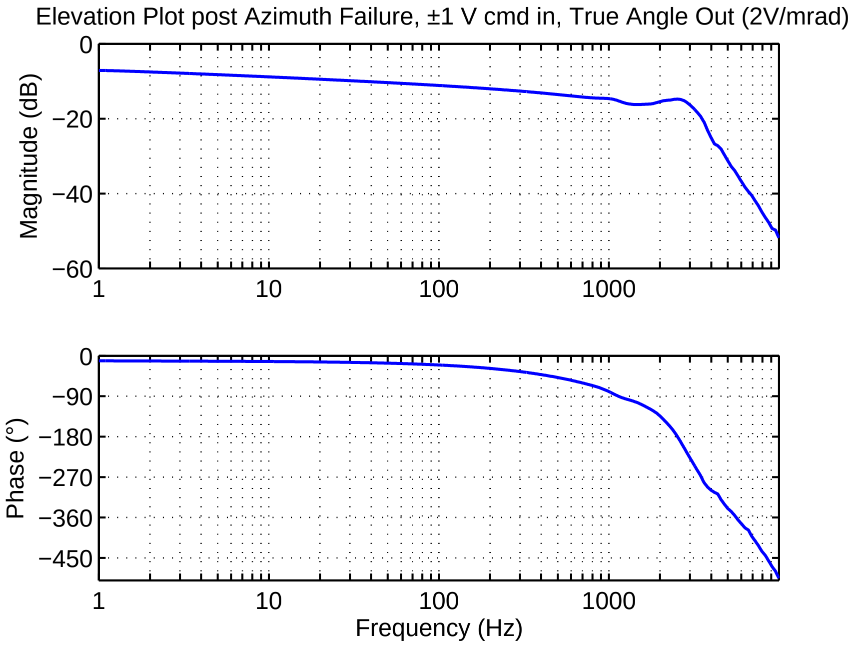

Figure B-2: Elevation axis frequency response following failure of the azimuth axis, with azimuth actuators removed.

with my calculations, and found no discrepancies. The above findings led me to concluded that the bearings were properly designed, and were not the cause of the unusual stiffness measurements.

B.3.2 Over-Constraint of Elastomer

The investigation of the over-constraint hypothesis was similarly straightforward. After the failure of the azimuth axis, I disassembled the AFSM hardware and examined the bearings while they were still preloaded with the magnets. I found no evidence of constraint around the free surfaces of the elastomer, as shown in Figure B-3. Although the image quality is poor, it is still possible to see light on either side of the bearing, as indicated by the arrows in the photograph.

As an additional check, I was also able to operate the still-intact elevation axis through its full stroke while examining the bearings, and found no interference. By virtue of this and the other visual evidence I ruled out over-constraint as a cause of the high stiffness.

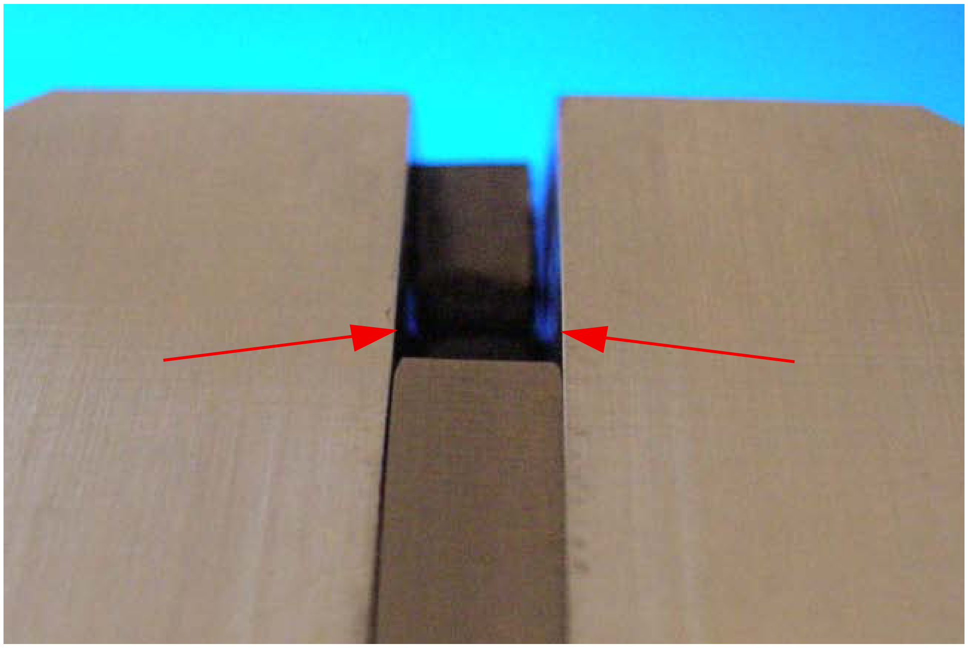

Figure B-3: Backlit photograph of one of the azimuth actuators after failure, showing clearance around the elastomer bearing (arrows).

B.3.3 Elastomer Compression Effects

To investigate the possibility of added stiffness due to unanticipated compression effects in the elastomer, I employed the finite element method. I modeled a neoprene block in ADINA using an orthotropic material model–that is, one which allows for an explicitly specified (rather than derived) value of the shear modulus. In the model I entered a compression modulus of 30 MPa and a shear modulus of 1.7 MPa. I also specified a Poisson's ratio of 0.4995.

To check the validity of the finite element model, I first ran three simpler cases for comparison to my existing hand calculations. In the first simulation I displaced the model in pure shear; in the second, pure torsion; and the third, pure compression. The results of these runs are listed in the first three lines of Table B.1, and they agree with the hand calculations to within a few percent.

The fourth line in Table B.1 is simply the consolidated total flexure stiffness from the hand and FEA results, assuming two bearings in shear, two in torsion, and the small contribution from the axial flexure bending. To this result I compared the stiffness of the FEA model when subjected to the "true" nodal displacements

Table B.1: Bearing Hand Calculations vs. Finite Element Results

Parameter	**Hand Calculation**	**FEA Result**
Shear Stiffness, Single Pad	0.136 N/μm	0.129 N/μm
Torsional Stiffness, Single Pad	4.715 N-m/rad	4.588 N-m/rad
Compressive Stiffness, Single Pad	2.400 N/μm	2.290 N/μm
Total Stiffness, no Compression	120.2 N-m/rad	114.4 N-m/rad
Total Stiffness with Compression	N/A	135.3 N-m/rad

resulting from the mirror and armature rotation. The total stiffness (again assuming contributions from torsion and the axial flexure) is given in the last line of Table B.1, and the resulting reaction forces are shown graphically in Figure B-4. Note that the viewpoint in the figure is looking down the 'long' end of the bearing; that is, the 4 mm dimension is the horizontal dimension in the figure.

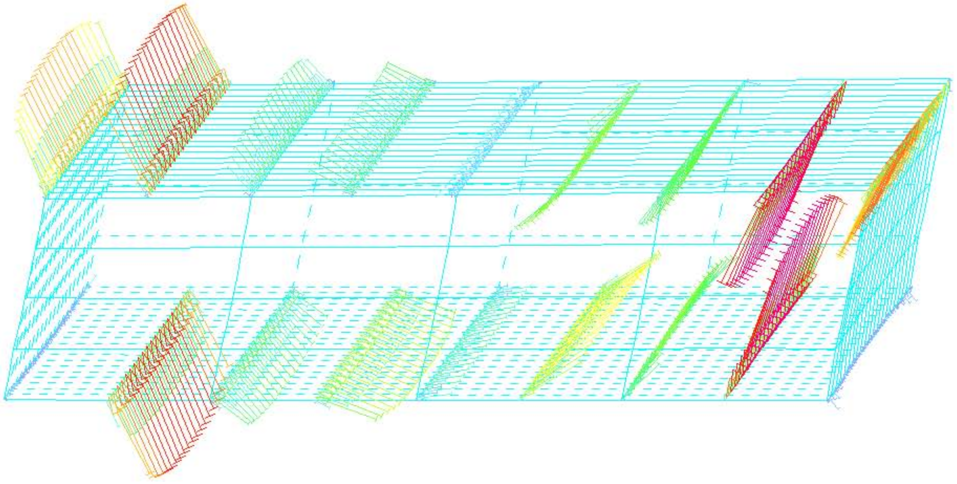

Figure B-4: Finite element analysis results of elastomer bearing showing shear and compressive / tensile reaction forces.

From the results, it is evident that the compression effects do in fact increase the total stiffness. However, the increase is only about 18% over the case with no compression, which by itself is not enough to account for the stiffnesses measured during the AFSM testing. A visual intuition for this result is given by Figure B-4. It is easy to see that the reaction forces in the normal direction are relatively large, but the moment arm over which they act is only 2 mm. The moment arm is simply too small to allow the normal forces to significantly contribute.

B.3.4 Elastomer Material Properties

I next turned to measurement of the bearing material properties. Fortunately, I had a significant quantity of spare material to test; however, the testing itself proved to be more involved than the prior verifications, and yielded some interesting results. I performed two sets of testing–one at Lincoln Laboratory and another at MIT–as explained below.

Material Testing at Lincoln Laboratory

As luck would have it, one of Lincoln Laboratory's resident experts in polymer material properties, Todd Mower, was a member of the committee that granted my funding for the AFSM project. During a committee report I informed him of the stiffness problems, and he offered to help with testing of the neoprene bearing material using a dynamic mechanical analyzer, or DMA. Taking advantage of this offer, I brought several samples of material for Todd to test. I was primarily interested in the test values of shear and compressive modulus for the Shore 70A neoprene sheet that I had used in the AFSM, but also brought along a softer Shore 50A neoprene and a latex rubber of Shore 40A hardness for comparison purposes.

Using the DMA, Todd obtained plots of the material modulus against frequency, which allowed for direct comparison to the AFSM test data. Unfortunately, the frequency range of the DMA is only 100 Hz, which is only a small fraction of the AFSM spectrum. Nevertheless, the data covered a wide enough frequency range to make the comparison meaningful.

The shear results for the AFSM Neoprene sample are shown in Figure B-5. Since the DMA obtains the material complex modulus, which also includes damping (loss factor) data, the value of interest is the storage modulus plotted in blue.

To my surprise, the results indicated a shear modulus of about 1 MPa at DC, whereas my bearing design assumed a value of 1.7 MPa. Thus, the measurements

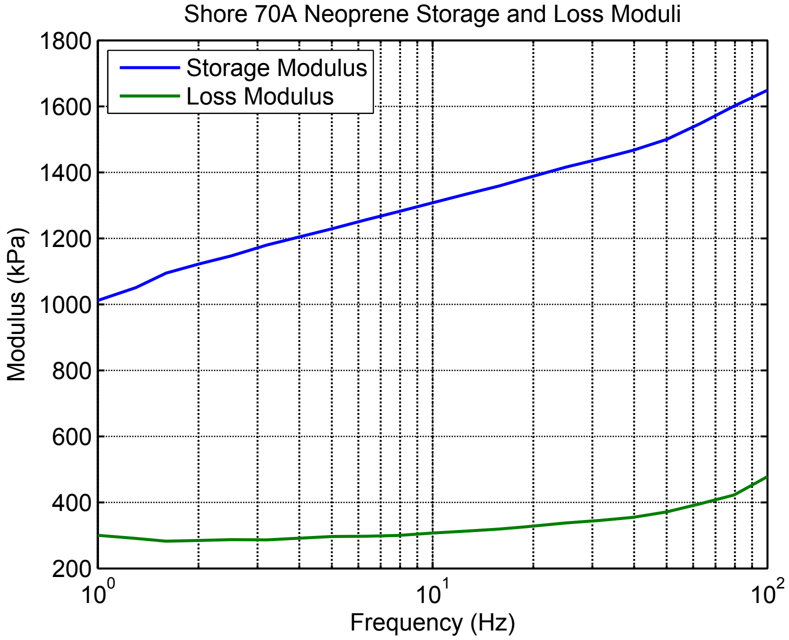

Figure B-5: Measured storage and loss modulus of 1mm thick Shore 70A neoprene sample in shear.

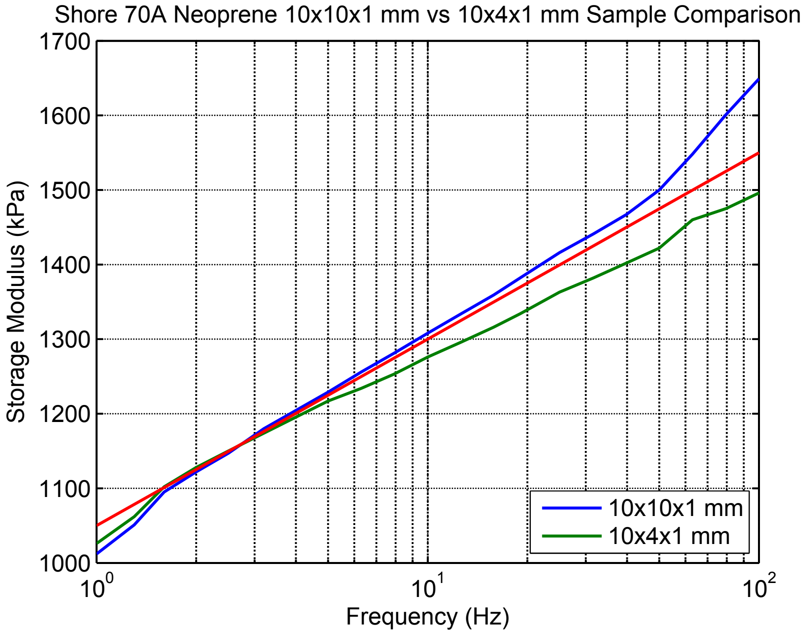

Figure B-6: Comparison of storage modulus results between two sample sizes of 1mm thick Shore 70A neoprene in shear. The red line denotes a +1.7 dB slope.

indicate that the rubber sample is actually less stiff than even the design value! The storage modulus results from another sample of differing size, shown in Figure B-6, confirm that the data is not an anomaly, and also confirm that the material shape factor does not affect the shear properties. In addition, both plots exhibit a logarithmic slope of about +1.7 dB per decade of frequency, which is consistent with the slope observed in the AFSM data. Also, a conversion of the loss modulus data into phase angle, as given in Figure B-7, indicates that the phase behavior of the neoprene is also consistent with the AFSM test results.

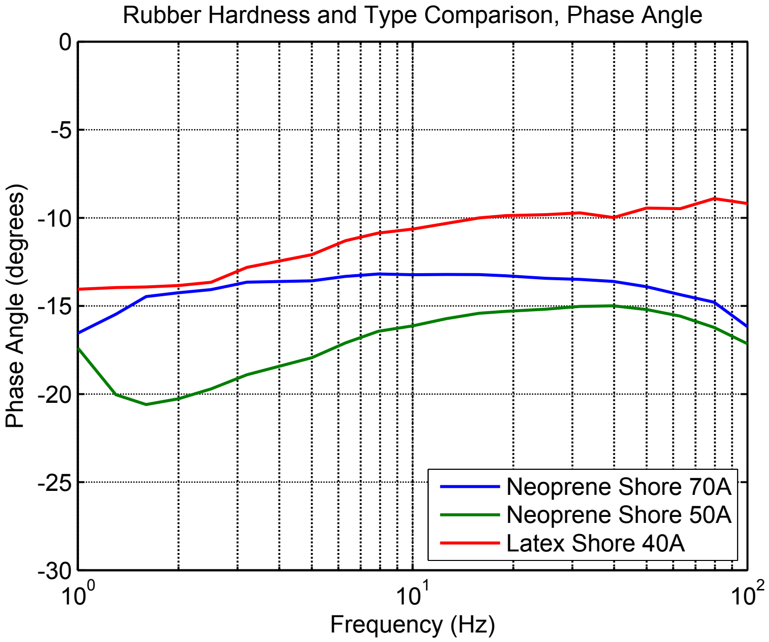

Figure B-7: Measured phase angle of various elastomer samples in shear.

The compression test results were similarly puzzling. Here, we performed only a static test of a sample identical in size to the AFSM bearing. The results of the test are indicated in Figure B-8. Because the sample was not bonded to the test fixture, the initial portion of the data displays a shallow slope as the fixture plates engage the sample. However, the latter part of the data exhibits meaningful, nearly linear results. However, here too the compression modulus test value fell far short of the design value. As indicated by the red line in the figure, the measured modulus is roughly 8 MPa, whereas the design value, based on Figure 2-9 and accounting for the shape factor, was 30 MPa.

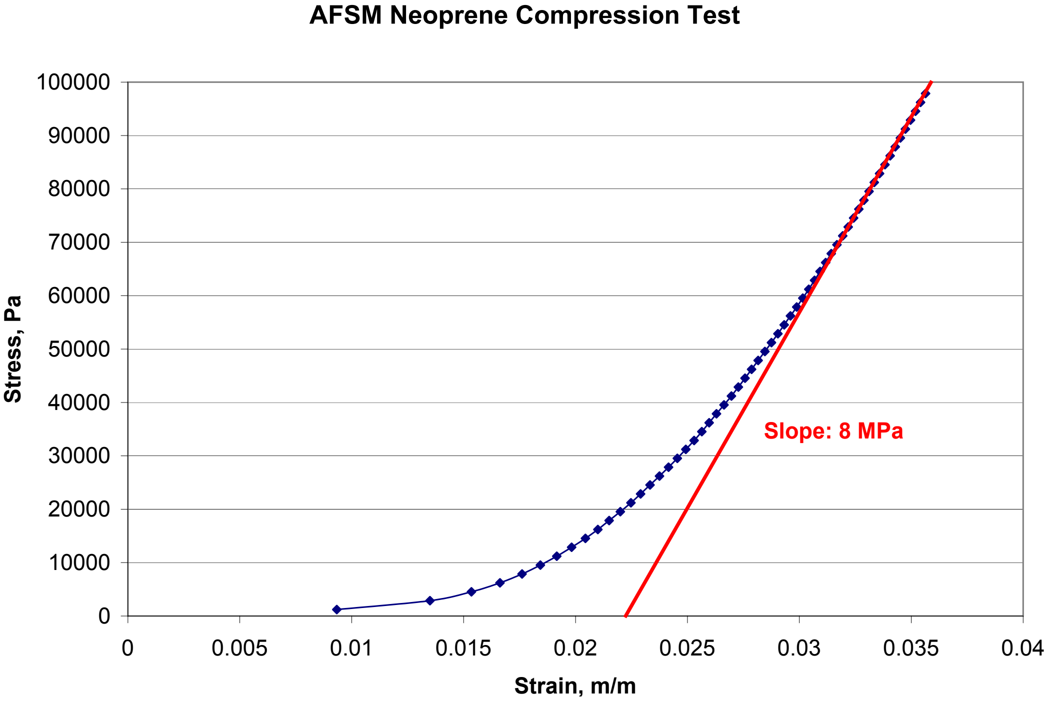

Figure B-8: Compression test of AFSM neoprene sample.

The results of the testing along with the design values are compiled in Table B.2. Surprisingly, the test values indicated that the raw material was significantly less stiff than the design values, which is in direct contradiction to the measured stiffness during AFSM testing.

Table B.2: AFSM Neoprene Static Mechanical Properties Measured Using the DMA

Parameter	**Design Value**	**Measured Value**
Shear Modulus	1.7 MPa	1.0 MPa
Compression Modulus	30.0 MPa	8.0 MPa

Material Testing at MIT

The puzzling results of the DMA testing prompted me to consider an alternate verification of the neoprene material properties which was more appropriate to the AFSM application. Although the DMA is an accurate instrument, the test environment to which it subjected the samples was different from the AFSM environment in several ways. First, the AFSM operated over a much larger frequency range than the 100 Hz limit of the DMA. Second, we did not bond the neoprene samples to the DMA test fixture using adhesive, as doing so would have ruined the fixture hardware. Third, neither the shear nor the compression samples were preloaded during the tests (other than a small clamping force). This is because the test jaws were rather delicate, and applying a large preload would have bent the jaws.

Fortunately, a piece of test hardware is available in our lab which is not subject to these limitations: The shear/compression tester designed and built by Augusto Barton as part of his master's thesis. Although Barton has since left MIT, I was able to find and set up his hardware with the AFSM neoprene samples. For details on the design and operation of the fixture, the reader is referred to Barton's thesis [17].

To perform the test, I set up the fixture in the shear configuration with three pads of Shore 70A neoprene taken from the same sheet as the AFSM bearings. Photographs of the test setup are given in Figures B-9 and B-10.

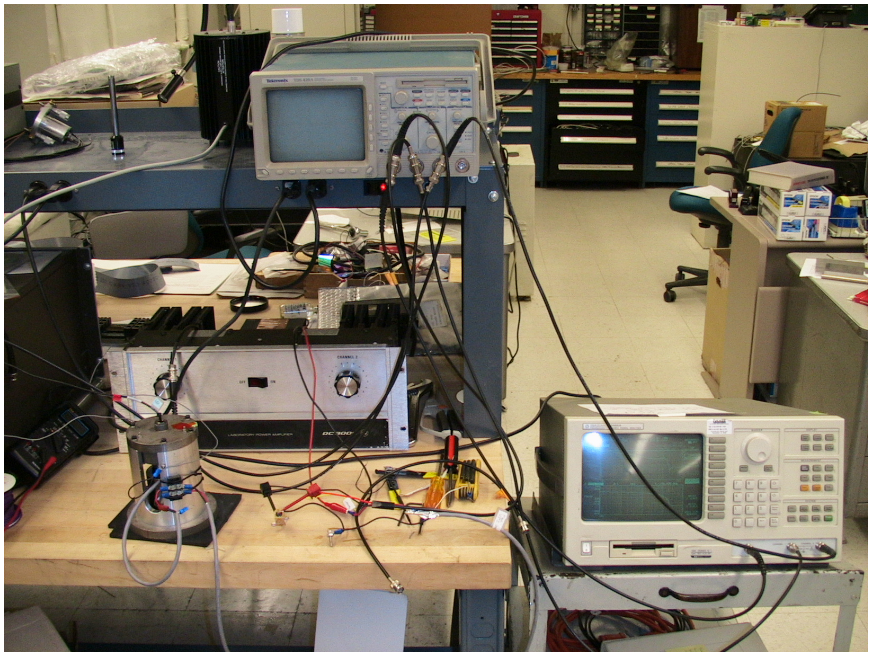

Figure B-9: Measuring neoprene dynamic shear modulus using Augusto Barton's test fixture.

Note that the shear fixture is designed with a clamping mechanism which allows

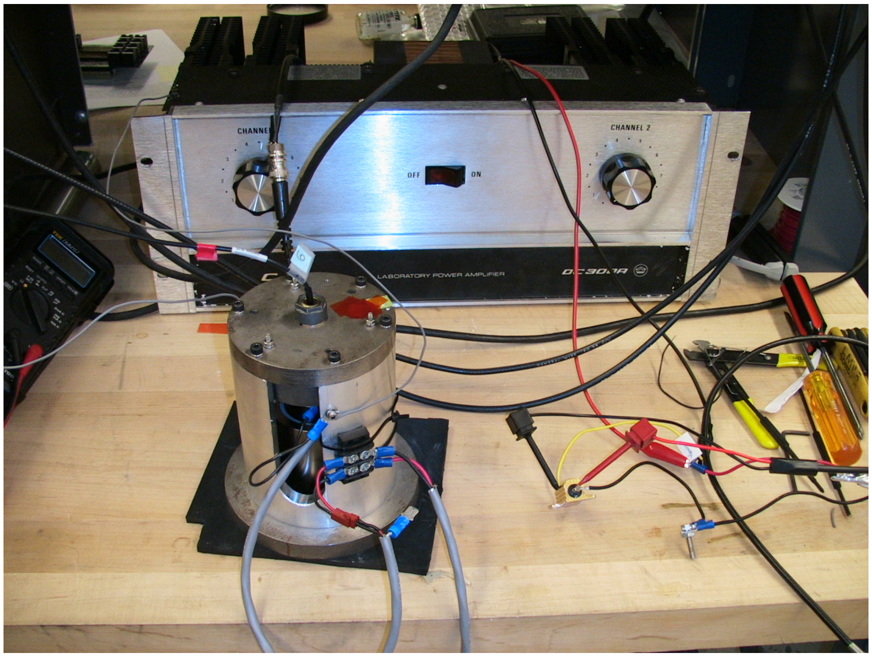

Figure B-10: Close-up of the shear test fixture.

significant preload pressures to be applied to the samples. In the AFSM, the elastomer bearings were subject to a compressive preload of about 0.7 MPa from the permanent magnets and a mechanical preload adjustment built into the back of the actuator. I set up the test fixture to match this preload.

Using the setup, I performed a static test of the shear modulus as well as a full frequency response. The static test data is given in Table B.3, and the frequency response is given in Figure B-11.

Table B.3: AFSM Neoprene Static Mechanical Properties Measured Using the MIT Tester

Parameter	**Measured Value**
Sample Dimensions (three samples total)	16x21x1 mm
DC Force Applied	9.66 N
Displacement at Force	1.59 μm
Shear Stiffness	6.06 N/μm
Shear Modulus	6.01 MPa

As can be seen, the new data obtained using Barton's fixture yielded significant results. First, the static shear modulus is much larger than either the DMA mea-

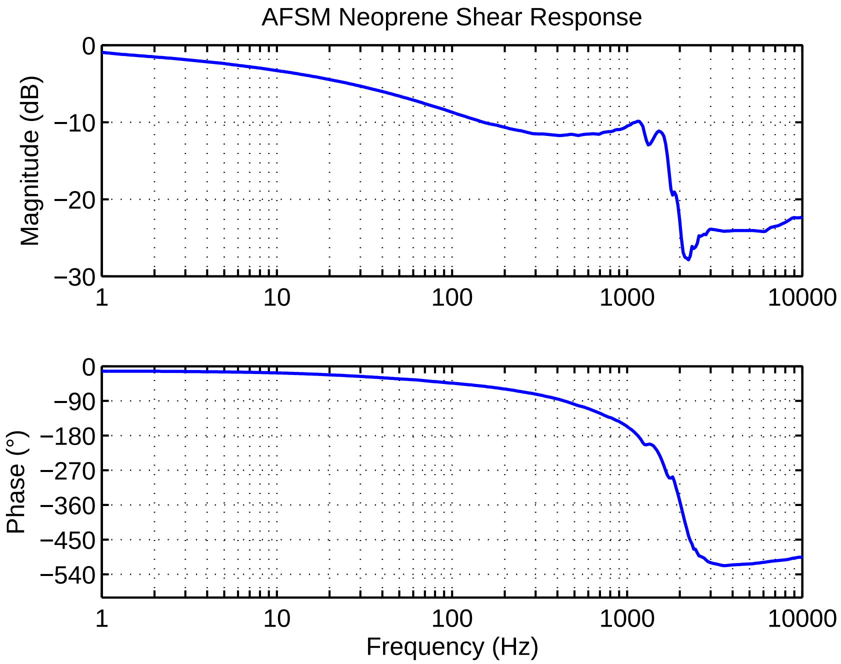

Figure B-11: Measured frequency response of AFSM neoprene in shear.

surements or the design value. In the latter case, the increase is almost a factor of four. Plugging the new value into my original design spreadsheet, I discovered that the new total flexure stiffness resulting from this value was 417.3 N-m/rad–almost a factor of four higher! Furthermore, when combined with the negative magnetic rate of the actuators, the resulting stiffness is 317.6 N-m/rad, which falls exactly within the range of values I measured during my initial verification.

The frequency response data in Figure B-11 is similarly compelling. The data is only valid up to about 1000 Hz due to the fixture resonance shortly thereafter; however, within this range the dynamic characteristics are very similar to the AFSM. Specifically, this includes the -15 degree phase and slowly decaying magnitude at low frequencies, and steeper slopes in both phase and magnitude at higher frequencies.

B.4 Investigation Conclusions

Based on the vast difference between the Lincoln Laboratory and MIT material test data, I concluded from the investigation that the preloading of the AFSM elastomer bearings was the primary cause of the increased stiffness measured in the test data. This conclusion was supported by the fact that the MIT data very closely matched the DC stiffness values that I measured during the actual AFSM testing. A secondary contributor to the stiffness increase was the compressive loading of the bearings that occurred from the rotation of the mirror, an effect for which I had not accounted during my original analysis. It should be noted that the general trends that Barton discovered during his thesis work also support these conclusions.

As a result of this investigation, I have concluded that using elastomer bearings in precision machines must be done with great care. This caution is doubly important for machines such as the AFSM, which rely on precise mechanical properties to achieve the advantageous dynamics which are key to successful performance.

Specifically, the exact loading and operating environment to which the bearings are subjected must be known and accounted for *a priori* during the design phase. Furthermore, it is highly recommended that material testing be done before a machine goes into production. It may also be prudent, if the application warrants, to maintain tight quality control over manufacturer material lots, backed up with periodic in-house tests. Provided such steps are taken, taking advantage of the compactness and inherent damping properties of elastomer bearings should still be possible.

Appendix C

Analog Compensator Analysis

In Chapter 5 I focused on the controller designs and resulting analog circuits without deriving the circuit transfer functions. I perform this derivation here.

Since the controller circuits are composed of separate stages containing either a lead or lag network, it is necessary only to derive transfer functions for the lead and lag stages. By the rules of block diagram manipulation (for example, see [8]), the full compensator transfer function is then just the product of the stages.

In the following derivations, ideal operational amplifiers are assumed. The derivation using non-ideal op-amps is more complicated and accurate, but within the active frequency range of the controllers both derivations produce acceptable results.

C.1 Lag Network

The diagram of the lag networks used in the AFSM controllers is shown in Figure C-1.

We wish to find the transfer function $G(s)$ between the output voltage $V_o(s)$ and the input voltage $V_i(s)$. For the inverting configuration used, the transfer function is given by

$$G(s) = \frac{V_o}{V_i} = -\frac{Z_f(s)}{Z_i(s)} \tag{C.1}$$

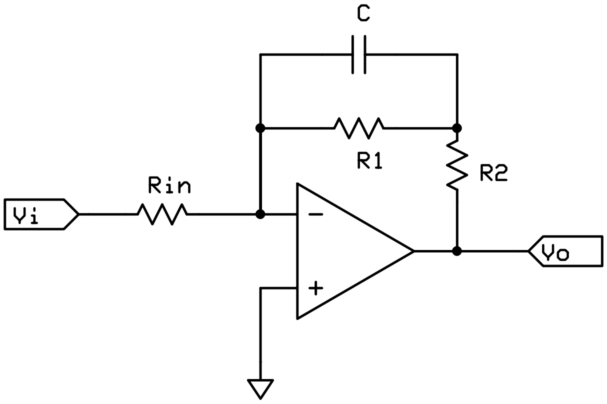

Figure C-1: A general lag analog circuit network.

Where $Z_f(s)$ and $Z_i(s)$ are the complex impedances of the feedback and input legs of the circuit (i.e., the two paths flowing into the inverting op-amp input). Referencing Figure C-1, these are

$$\begin{aligned} Z_f(s) &= R_2 + \frac{1}{\frac{1}{R_1} + Cs} = \frac{R_1 R_2 C s + R_1 + R_2}{R_1 C s + 1} && \text{(C.2)} \\ Z_i(s) &= R_{in} && \text{(C.3)} \end{aligned}$$

Substituting the impedances into Equation C.1 (which for the lag configuration we denote as $G_{lag}(s)$) and writing the result in Bode form (setting trailing polynomial coefficients equal to one) gives

$$\begin{aligned} G_{lag}(s) &= -\frac{\left(\frac{R_2 R_1 C s + R_1 + R_2}{R_1 C s + 1}\right)}{R_{in}} \\ &= -\frac{(R_1 + R_2)\left(\frac{R_2 R_1 C s}{R_1 + R_2} + 1\right)}{R_{in}(R_1 C s + 1)} && \text{(C.4)} \end{aligned}$$

Rearranging slightly to separate the time constant $R_1 C$, we have

$$G_{lag}(s) = -\left(\frac{R_1 + R_2}{R_{in}}\right) \frac{\left(\frac{R_2}{R_1 + R_2}\right) R_1 C s + 1}{R_1 C s + 1} \tag{C.5}$$

By defining the terms

$$
\begin{aligned}
K_1 &= \frac{R_1 + R_2}{R_{in}} && \text{(C.6)} \\
\beta &= \frac{R_1 + R_2}{R_2} && \text{(C.7)} \\
\tau_1 &= R_1 C && \text{(C.8)}
\end{aligned}
$$

it is easy to see that Equation C.5 is simply a traditional lag compensator such as that described in [8]:

$$
G_{lag}(s) = -K_1 \frac{\left(\frac{1}{\beta}\right) \tau_1 s + 1}{\tau_1 s + 1} \tag{C.9}
$$

In this form, it is evident that K_1 is the loop gain, β is the separation factor between the lag pole and zero locations (expressed as a multiple of the pole frequency), and τ_1 is the time constant, the inverse of which determines the lag pole location on the s-plane. By designing a lag compensator using standard methods to get K_1, β, and τ_1, we then use Equations C.6, C.7, and C.8 to select component values for R_1, R_2, R_{in}, and C.

C.2 Lead Network

To obtain a lead compensator, we simply switch the input and feedback impedance configurations of the lag network to obtain the circuit in Figure C-2.

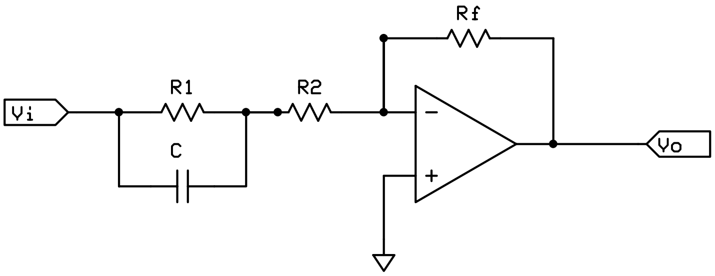

Figure C-2: A general lead analog circuit network.

The derivation for this network is accordingly similar. Here we have

$$\begin{aligned} Z_f(s) &= R_f && \text{(C.10)} \\ Z_i(s) &= R_2 + \frac{1}{\frac{1}{R_1} + Cs} = \frac{R_1 R_2 Cs + R_1 + R_2}{R_1 Cs + 1} && \text{(C.11)} \end{aligned}$$

Substituting into Equation C.1 (this time denoting as $G_{lead}(s)$) and again writing in Bode form results in

$$\begin{aligned} G_{lead}(s) &= -\frac{R_f}{\left(\frac{R_2 R_1 Cs + R_1 + R_2}{R_1 Cs + 1}\right)} \\ &= -\left(\frac{R_f}{R_1 + R_2}\right) \frac{R_1 Cs + 1}{\left(\frac{R_2}{R_1 + R_2}\right) R_1 Cs + 1} && \text{(C.12)} \end{aligned}$$

Defining new constants for the lead parameters, we have

$$\begin{aligned} K_2 &= \frac{R_f}{R_1 + R_2} && \text{(C.13)} \\ \alpha &= \frac{R_1 + R_2}{R_2} && \text{(C.14)} \\ \tau_2 &= R_1 C && \text{(C.15)} \end{aligned}$$

which results in the traditional lead compensator form given by

$$G_{lead}(s) = -K_2 \frac{\tau_2 s + 1}{\left(\frac{1}{\alpha}\right) \tau_2 s + 1} \qquad \text{(C.16)}$$

Here, the inverse of the time constant τ_2 gives the compensator zero location, and α is the separation between the zero and pole expressed as a multiple of the zero's frequency. The loop gain is given by K_2. Component selection for the lead compensator design given by Equation C.16 is the same as the lag compensator.

Bibliography

[1] Michael F. Ashby. *Materials Selection in Mechanical Design*, chapter 4, pages 36–38. Butterworth-Heinemann, Oxford, England, 1999.

[2] Ball Aerospace Corporation. Custom fast steering mirrors. http://www.ballaerospace.com/fsm_precision.html, Accessed Nov 2005.

[3] Klaus-Jürgen Bathe. *Finite Element Procedures.* Prentice Hall, Upper Saddle River, New Jersey, 1996.

[4] Ferdinand P. Beer and E. Russel Johnston Jr. *Mechanics of Materials.* McGraw-Hill, New York, New York, second edition, 1992.

[5] Cambridge Technology. How to choose a closed loop galvanometer based scanner. http://www.cambridgetechnology.com/news/Choosing_A_Galvanometer.html, Accessed Jan 2007.

[6] David P. Cuff. Electromagnetic nanopositioner. Masters of science, Massachusetts Institute of Technology, Department of Mechanical Engineering, 2007.

[7] Dean C. Karnopp et. al. *System Dynamics: Modeling and Simulation of Mechatronic Systems.* John Wiley and Sons, New York, New York, third edition, 1999.

[8] Gene F. Franklin et. al. *Feedback Control of Dynamic Systems.* Addison-Wesley, Reading, Massachusetts, third edition, 1994.

[9] Bruce Hardy and Paul Hauser. Scanning apparatus and method that avoids unwanted reactions. U.s. patent no. 6,612,192, September.

[10] Hermann A. Haus and James R. Melcher. *Electromagnetic Fields and Energy*. Prentice-Hall, 1989.

[11] Larry Edward Hawe II. Control of a fast steering mirror for laser-based satellite communication. Masters of science, Massachusetts Institute of Technology, Department of Mechanical Engineering, February 2006.

[12] William H. Hayt Jr. and Jack E. Kemmerly. *Engineering Circuit Analysis*. McGraw-Hill, New York, New York, fifth edition, 1993.

[13] Daniel Kluk. Modal analysis of a thick plate structure. August 2006.

[14] P. B. Lindley. Engineering design with natural rubber. Technical bulletin, The Malaysian Rubber Producers' Research Association, 1970.

[15] Gregory C. Loney. High bandwidth steering mirror research. Project Report IRP-15, MIT Lincoln Laboratories, Lexington, MA, January 1992.

[16] Xiaodong Lu. *Electromagnetically-Driven Ultra-Fast Tool Servos for Diamond Turning*. PhD thesis, Massachusetts Institute of Technology, 2005.

[17] Augusto E. Barton Martinelli. Rubber bearings for precision positioning systems. Masters of science, Massachusetts Institute of Technology, Department of Mechanical Engineering, 2005.

[18] R.H. Maskrey and W.J. Thayer. A brief history of electrohydraulic servomechanisms. *ASME Journal of Dynamic Systems Measurement and Control*, 1978.

[19] Leander W. Matsch. *Electromagnetic and Electromechanical Machines*. International Textbook Company, Scranton, Pennsylvania, 1972.

[20] James Medbery and Avanindra Gupta. Stable controller design of a six degree-of-freedom magnetically suspended fine steering mirror (msfsm). *SPIE Active and Adaptive Optical Components*, 1543, 1991.

[21] Wouter Monkhorst. Dynamic error budgeting: A design approach. Masters of science, T. U. Delft, May 2004.

[22] Richard Clement Montesanti. *High Bandwidth Rotary Fast Tool Servos and a Hybrid Rotary/Linear Electromagnetic Actuator.* PhD thesis, Massachusetts Institute of Technology, 2005.

[23] Physik Instrumente GmbH. Fast steering mirros with piezoelectric drives. http://www.physikinstrumente.com/en/products/steering_mirror/index.php, Accessed Nov 2005.

[24] Adrian Mathias Rankers. *Machine Dynamics in Mechatronic Systems: An Engineering Approach.* PhD thesis, Philips Center for Manufacturing Technology, 1997.

[25] Richard C. Rice. Metallic materials properties development and standardization. Scientific Report DOT/FAA/AR-MMPDS-01, U.S. Department of Transportation, Federal Aviation Administration, Washington, DC, January 2003.

[26] Herbert C. Roters. *Electromagnetic Devices.* John Wiley and Sons, New York, New York, 1941.

[27] Paresh C. Sen. *Principles of Electric Machines and Power Electronics.* John Wiley and Sons, New York, New York, 1997.

[28] Ming-Chi Weng. *Magnetic Suspension and Vibration Control of Flexible Structures for Non-Contact Processing.* PhD thesis, Massachusetts Institute of Technology, 2000.

[29] Herbert H. Woodson and James R. Melcher. *Electromechanical Dynamics.* John Wiley and Sons, New York, New York, 1968.