

WIDEBAND ACTIVE VIBRATION CONTROL SYNTHESIS AND
IMPLEMENTATION ON UNCERTAIN RESONANT STRUCTURES

CORY M. PAPENFUSS

Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Electrical Engineering

William T. Baumann, Chair
Chris R. Fuller
Marty E. Johnson
Douglas K. Lindner
Jason Lai

April 4, 2006
Blacksburg, Virginia

KEYWORDS:

Active vibration control, heirarchical control, adaptive control, vibration control synthesis,
vibration control implementation

Copyright © 2006, Cory M. Papenfuss

WIDEBAND ACTIVE VIBRATION CONTROL SYNTHESIS AND IMPLEMENTATION ON UNCERTAIN RESONANT STRUCTURES

Cory M. Papenfuss

(ABSTRACT)

Large, “rigid” structures made of interconnecting beams and tendons represent a type of construction that has many engineering benefits. Lightweight, strong, and generally constructed of standard-shaped metal framework, they also incur significant challenges to control vibration. The construction that makes them strong, light, and “rigid” also makes them have very small inherent structural damping, and highly-complex modal structure. A myriad of control techniques have been developed to work on this problem with success usually in very small bandwidths related to a specific frequency “hot-spot” around a few specific modes.

This work describes the design, analysis, and implementation of a novel controller configuration applicable to broadband vibration suppression on a large, uncertain resonant structure. The measurement, identification, characterization, and modeling of a large, flexible, lightly-damped test structure with in excess of 1000 modes in 50-5000Hz range was used as the basis for choosing a control configuration. This choice leverages the relative benefits of different control types to obtain one with a combination of the best features of all of them. High-order and low-order feedback, and feed-forward controller configurations were all used in different frequency ranges. Real design tradeoffs such as computation complexity, model accuracy, and available actuator technologies were fundamental to the design choices.

Measured individual modal reduction was as much as 15dB for feedback control, 20dB for feed-forward control, and 4dB broadband over the range of 50-5000Hz.

ACKNOWLEDGMENTS

I would like to thank all those who assisted me in this research endeavor. In particular, I wish to thank the Vibrations and Acoustics Laboratory in the Mechanical Engineering Department at Virginia Tech for providing me with the support, facilities, and insight to apply active control to such a complicated structural system. Also, I would like to thank the supporters of the original project, ESI/VASCI for providing the research opportunity and financial support.

I would like to express my gratitude to Dr. William Baumann who graciously provided me with continual guidance throughout this project. His valuable theoretical insight grounded in realistic limitations provided a spring-board for refining the research goals of my work. He also provided me with invaluable assistance in proof reading, editing, and refining the documentation of the work presented in this dissertation.

Finally, I would like to thank my parents and siblings for helping me keep my bearings and for giving me the resolve to plunge forward.

TABLE OF CONTENTS

TITLEPAGE	i
ABSTRACT	ii
ACKNOWLEDGMENTS	iii
TABLE OF CONTENTS	iv
LIST OF FIGURES	vii
LIST OF TABLES	xi
1 INTRODUCTION	1
1.1 Active Vibration Control	1
1.2 Contributions and Applications	4
2 LITERATURE BACKGROUND	7
2.1 Structural Modeling	8
2.2 Model-based Control	9
2.2.1 Low-order feedback control	9
2.2.2 High-order feedback control	10
2.3 Dissipative Control	12
2.4 Feed-forward Adaptive Control	12
3 PROJECT BACKGROUND	15
3.1 Control Problem Defined	16
3.1.1 Performance requirements	16
3.1.2 Implementation limitations	17
3.2 Experimental Test Setup	18
3.3 Model Identification	22
3.3.1 Identification results	23
3.3.2 Measurement artifacts	28
3.3.3 Analytical model characterization	29
3.4 Approximate Modeling	34
3.4.1 Static shaker	35
3.4.2 Shaker with dynamic load	37
3.4.3 Plate bending wave equation	38
3.4.4 Modal overlap	42

3.5	Summary	44
4	WIDEBAND CONTROLLER SYNTHESIS	46
4.1	Proposed Control Strategy	47
4.2	High-frequency Synthesis	53
4.2.1	Collocated structural response	53
4.2.2	Realistic actuator	58
4.2.3	Compensated feedback	62
4.2.4	High-frequency summary	66
4.3	Low-frequency Synthesis	68
4.3.1	Synthesis model	68
4.3.2	Practical experimental synthesis	73
4.4	Mid-frequency Synthesis	75
4.4.1	Secondary path model	76
4.4.2	Time-domain algorithm	77
4.4.3	Filter length	78
4.5	Summary	79
5	HIGH-FREQUENCY PERFORMANCE OPTIMIZATION	81
5.1	Problem Specification	82
5.1.1	Setup	87
5.2	Optimization Results	91
5.2.1	Unity plant	92
5.2.2	Modeled plant	98
5.3	Summary	105
6	AGGREGATE CONTROLLER ANALYSIS	106
6.1	System Identification	107
6.1.1	Plant aliasing	112
6.1.2	Plant aliasing with feedback	116
6.2	Discrete-time Rate Changing	121
6.2.1	Decimation	121
6.2.2	Interpolation	124
6.3	LMS Optimal Solution Computation	127
6.4	Total System Modeling	133
6.4.1	Model extension for feed-forward	134
6.4.2	Filtered-error LMS	137
6.4.3	Filtered-error, filtered-reference LMS	142
6.4.4	Model extension for low-frequency feedback	144
6.5	Stability Analysis	149
6.5.1	Disturbance path stability equivalent	153
6.6	Summary	155

7	EXPERIMENTAL RESULTS	156
7.1	High-frequency Results	156
7.2	Low-frequency Results	160
7.2.1	Controller variations	162
7.2.2	Performance	165
7.2.3	Implementation artifacts	169
7.3	Mid-frequency Results	172
7.3.1	Transient performance	176
7.4	Aggregate Results	177
7.4.1	Broadband performance measurement	177
7.4.2	Unmodeled disturbance	183
7.4.3	Unmodeled structural variation	187
7.5	Results	191
8	CONCLUSIONS AND CONTINUED RESEARCH	193
8.1	Contributions	193
8.2	Further Research	195
	APPENDIX A DIGITAL CONTROLLER	196
A.1	Hardware	196
A.2	Control Code	197
	APPENDIX B ANALOG IMPLEMENTATION	202
	REFERENCES	205
	VITA	213

LIST OF FIGURES

Figure	Page
1.1 Representative LFSS	2
3.1 Test structure to represent a large, flexible structure	19
3.2 Actuators to error accelerometers instances (gray) and mean (solid).	21
3.3 SOCIT toolbox system ID (actuator path)	25
3.4 SOCIT toolbox system ID (disturbance path)	26
3.5 SOCIT toolbox system ID phases	27
3.6 Example (disturbance path) model-order reduction from high-order model	28
3.7 Hankel singular values of identified model	29
3.8 Measurement coherence for experimental data used in model fit	30
3.9 Actuator path pole-zero map	31
3.10 Disturbance path pole-zero map	32
3.11 Disturbance path RHP zero map of different IDs	32
3.12 Model ID absolute phase (unwrapped)	33
3.13 Measured structural damping	34
3.14 Fixed electromagnetic shaker schematic.	35
3.15 Actual electromagnetic shaker characterized.	36
3.16 Electromagnetic shaker experimental results and model fit.	37
3.17 Electromagnetic shaker schematic.	38
3.18 Bending wavelength of 12.7mm (1/2”) aluminum plate	40
3.19 Electromagnetic actuator footprint	41
3.20 Simple baseplate analytical frequency response	42
3.21 Simple thin plate analytical frequency response	43
3.22 Simple plate analytical modal overlap	44
3.23 Measured structural modal overlap	45
4.1 Multi-bandwidth control regimes	49
4.2 Multi-bandwidth control scheme	52
4.3 Sample undamped beam accelerance	54
4.4 Idealized lightly damped model	56
4.5 Root locus of idealized lightly damped plant (under acceleration or velocity feedback)	57
4.6 Idealized lightly damped model without and with actuator	59
4.7 Root locus of lightly damped plant with actuator (under acceleration or velocity feedback)	60

4.8	Nyquist vs. Nichols plot of pure velocity feedback on lightly damped plant	62
4.9	Nichols plot of simple plant, acceleration, and compensated acceleration feedback	63
4.10	Idealized lightly damped model closed loop	65
4.11	Root locus of lightly damped plant with compensated actuator (using accelera- tion or velocity input	66
4.12	Idealized lightly damped model without and with actuator	67
4.13	Low-frequency H_2 design model	68
4.14	Low-frequency H_2 design model in standard form	69
4.15	Low-frequency H_2 design model order reduction	71
4.16	Low-frequency H_2 design filters	71
4.17	Low-frequency H_2 design loop gain	72
4.18	Low-frequency H_2 example performance singular values	72
4.19	Low-frequency H_2 example performance bode plot	73
4.20	Low-frequency H_2 synthesis experimental setup	74
4.21	Low-frequency H_2 synthesis model order comparison	75
4.22	Basic adaptive feed-forward system	76
4.23	“Filtered-X” adaptive feed-forward system	76
4.24	Filter length vs. structural damping	80
5.1	Analog loop performance bounds	82
5.2	Collocated actuator and sensor	84
5.3	Plant disturbance rejection	85
5.4	Cost function variations	91
5.5	Valid pole/zero/gain (P/Z/K) sets for lenient constraints on a unity plant	93
5.6	Lenient phase P/Z/K mapping by band-limited RMS integral	94
5.7	Lenient phase optimal loop gain	94
5.8	Lenient phase disturbance rejection and performance	95
5.9	Strict phase P/Z/K ranked by band-limited RMS integral	96
5.10	Lenient vs. strict phase loop gain	97
5.11	Lenient vs. strict phase disturbance rejection	97
5.12	Candidate compensated loop gains	99
5.13	Candidate ZPK combinations’ bandlimited RMS performance	100
5.14	Candidate ZPK combinations’ full-band RMS performance	101
5.15	Valid P/Z/K sets	101
5.16	P/Z/K map by bandlimited RMS (3dB GM, 60°PM)	102
5.17	P/Z/K map by bandlimited RMS (6dB GM, 45°PM)	103
5.18	P/Z/K map by peak gain (6dB GM, 45°PM)	104
5.19	Candidate compensated loop gains (6dB GM, 45°PM)	104
6.1	Impulse-sampled, zero-order held signal	108
6.2	Mixed-signal plant ID setup	109
6.3	Mixed-signal plant ID normalized	110
6.4	Baseband and impulse-sampled hold and anti-aliasing filter	111
6.5	Absolute aliased plant residue	114
6.6	Illustration of different α on a sampled plant	115
6.7	Upper bound on aliasing relative error, $\Delta(\omega)$	116

6.8	Upper bound on continuous-time plant model error due to aliasing	117
6.9	Mixed-signal plant ID setup with analog feedback	118
6.10	Relationship of α to α_a at different phase margins	120
6.11	Decimation analysis setup	123
6.12	Decimation aliasing errors	124
6.13	Decimation/interpolation of plant	125
6.14	Interpolation errors	126
6.15	Decimation and interpolation error summary	126
6.16	General estimation filter	128
6.17	Estimation filter with actuator dynamics	129
6.18	No filter FxLMS input frequency response	130
6.19	Optimal filter computations	131
6.20	No filter FxLMS output filters	132
6.21	Broadband FxLMS performance	133
6.22	Analog actuator and disturbance model	135
6.23	Hybrid continuous-discrete feed-forward model.	135
6.24	Minimum-phase representation of P	137
6.25	FxLMS error-weighting	138
6.26	FxLMS error-weighting equivalent	138
6.27	FxLMS error-weighting equivalent consolidated	139
6.28	Error-weighting FxLMS input frequency response	140
6.29	Error-weighting FxLMS output filters	140
6.30	Broadband FxLMS performance with error weighting	141
6.31	FxLMS error-weighting equivalent with reference filtering	142
6.32	Error-weighted, reference filtered output filters	143
6.33	Broadband FxLMS performance with error weighting and filtered reference	144
6.34	Three controller combination configuration	145
6.35	Three controller combination with high-frequency controller equivalent	146
6.36	Three controller combination equivalent with in-path error filter	146
6.37	Optimal filter computations	147
6.38	Three controller combination equivalent	148
6.39	Three controller combination equivalent with $B_r^* = B_e^*$	148
6.40	Actuator path loop gain	149
6.41	Multiplicative uncertainty	150
6.42	Actuator loop gain uncertainty	152
6.43	Actuator loop error mechanisms	152
6.44	Actuator loop sensitivity bounds	153
6.45	Disturbance “feedback” loop from total model	154
6.46	Disturbance “feedback” loop partially unwrapped	154
6.47	Disturbance “feedback” loop unwrapped	154
7.1	Measured continuous-time controller compensator	158
7.2	Measured continuous-time controller performance (actuator path)	159
7.3	Measured continuous-time controller performance (disturbance path)	161
7.4	Low-frequency ID model order comparison (actuator path)	163
7.5	Low-frequency ID model order comparison (disturbance path)	164

7.6	Low-frequency controller gains varying design noise level	165
7.7	Low-frequency actuator loop gain varying design noise level	166
7.8	Low-frequency measured performance with varying aggressiveness	167
7.9	Low-frequency measured performance (actuator path)	168
7.10	Low-frequency measured performance (disturbance path)	170
7.11	Low-frequency reduced order compensator	171
7.12	Low-frequency reduced order compensator measured performance	171
7.13	Mid-frequency ID model order comparison (actuator path)	172
7.14	Mid-frequency disturbance rejection unfiltered vs. filtered	174
7.15	Mid-frequency disturbance rejection vs coherence limit	175
7.16	Mid-frequency transient filtered error	176
7.17	Full-bandwidth measured performance (actuator path)	178
7.18	Full-bandwidth measured performance (disturbance path)	179
7.19	Full-bandwidth measured progressive RS performance (disturbance path)	181
7.20	Full-bandwidth measured progressive RS sequential performance (disturbance path)	182
7.21	Low-frequency measured performance with varying aggressiveness	184
7.22	Full-bandwidth measured performance (unmodeled disturbance path)	185
7.23	Full-bandwidth measured progressive RS sequential performance (unmodeled disturbance path)	186
7.24	Full-bandwidth measured performance with varying structure (actuator path)	188
7.25	Low-frequency robustness to structural variation	189
7.26	Full-bandwidth measured performance with varying structure (disturbance path)	190
7.27	Full-bandwidth measured progressive RS performance with varying structure (disturbance path)	190
7.28	Full-bandwidth measured performance with varying structure (unmodeled disturbance path)	191
8.1	Possible multivariable extension	195
A.1	LMS code block as Simulink object	198
A.2	Hard real-time mode switch as Simulink object	199
A.3	Decimated H_2 block as Simulink object	200
A.4	Complete controller as Simulink object	201
B.1	Analog compensator schematic	203

LIST OF TABLES

Table	Page
3.1 Shaker coefficients.	37
3.2 Materials properties for Aluminum	40
B.1 Analog component values	204

CHAPTER 1

INTRODUCTION

This chapter discusses an overview of the work presented in this dissertation. A brief introduction to active control as applicable to large, flexible structures is discussed, along with the assumptions made and underlying project goals in Sec. 1.1. The original contributions are outlined, and the applicability of the work to other projects is discussed in Sec. 1.2.

1.1 Active Vibration Control

Active mechanical control has been an interesting topic in engineering for a long time. James Watts' flyball governor in the late 18th century provides an example of the desire to control mechanical devices with some form of automated control mechanism. By the early 20th century, electrical systems knowledge progressed to the point of providing a means to implement more complicated control techniques [1], [2], [3]. With advances in linear systems' mathematical descriptions, the "modern control theory" techniques [4] allowed for more sophisticated optimal, robust, and multivariable control [5].

The field of active vibration control relies on a variety of engineering disciplines. Analysis of the system requires the use of structural modeling techniques based in fundamental kinematic equations. These are generally difficult or impossible to solve in closed-analytic form and usually require numerical or finite-element computational methods. By the later part of the 20th century, classical and modern control theory and digital control provided tractable methods [6] [7]. of dealing with large, complicated physical structures [8]. Classical theory provides some relatively simple solutions for some classes of problems if they have sufficient simplicity and structure. Modern control theory can be used for design based on the analysis of sufficiently accurate structural models.

A particularly challenging application for active vibration control is in the field of large flexible space structures (LFSS) [9]. The structures are necessarily lightweight, extremely rigid, and rely on structurally complex geometries. This provides uniquely difficult challenges for the control of attitude and payload vibration minimization. The research field is extremely varied in its scope, and contains no “magic bullets” for implementation. The inherently high-order structures make accurate models difficult to obtain and use analytically.

This work makes use of the complicated dynamic structure of a representative LFSS (Fig. 1.1) to independently design and implement a wideband active vibration suppression controller. By leveraging the strengths of three different control types using frequency separation, design is simplified, while performance is enhanced to a degree not feasible with a single control scheme.



Figure 1.1 Representative LFSS

The goals of the control design are as follows:

- Minimize the vibration of a portion of a large, flexible structure over a very wide bandwidth. This goal represents a reduction in the mean-squared acceleration over a frequency range. In this case, the frequency range was from 0-5kHz. The vibration can be considered to come primarily from a known disturbance source, but the possibility of unknown disturbances must also be minimized.

- A minimum number of control actuators will be used for specific point control. In particular, this can be considered an “attachment” problem where the vibration at a single or a few specific points is to be minimized. This will allow for a minimum of energy transmitted from one sub-structure to another through its points of attachment. In the representative example of Fig. 1.1, the representative point of attachment is at the bottom plate.
- The primary cause of disturbances causing vibration in the structure will be due to a source within the structure. In order to minimize its effect on the target points, it will be located a long distance away from the attachment points. The propagation path of this disturbance will be long and subject to change with minor structural variations.
- Computational complexity for control implementation must be kept to a minimum. Although no *specific* implementational constraint has been placed on the hardware available, it can be assumed that the processor required for implementation will be reduced to the minimum necessary.
- Structure will have manufacturing and operational variability that present significant modeling uncertainty. This will be represented by a variation of the location of a small fraction of the mass of the structure.

Chapter 2 briefly discusses the basis and limitations of different control schemes previously applied to the LFSS problem from a, historical perspective. Chapter 3 explains the details of the experimental structure, its modelling, and various characteristics of it that were instrumental in choosing the control strategy. Chapter 4 describes the synthesis of the proposed control strategy, how it differs from other approaches, and the process of the control design with reference to the test system. The optimization of one of the parts of the synthesized controller implementation is discussed in Chap. 5. Chapter 6 presents an analysis of the various errors introduced in the implementation of the synthesized controller, which requires the development of a hybrid continuous/discrete-time model of the entire control and structural configuration. Measured experimental results on the test structure are shown in chapter 7. Finally, conclusions and suggestions for further research are discussed in Chap. 8

1.2 Contributions and Applications

The structure of this dissertation is mostly presented in a tutorial fashion, since one of the primary contributions of the work is indeed the *methodology* of the synthesis procedure. Other contributions are found throughout, but are summarized as follows:

- **Provide a synthesis procedure for a multi-loop, broadband active vibration control system, incorporating feedforward, feedback, continuous-time and multirate discrete-time control:** The control synthesis utilizes experimentally-derived implementation limitations that drive the design of three different control methodologies. The aggregate controller leverages the benefits of each of the three to obtain the benefits of each while simultaneously minimizing interaction between them. Each type is tailored to a bandwidth that provides the most benefit for the least amount of implementation resources.
- **Provide a frequency-domain optimization technique using simplified analytical models of a highly complex and unknown system to synthesize a low-order continuous-time controller:** The heuristic loop-shaping approach to the low-order, continuous-time design was augmented with an objective optimization procedure. The result was an optimized frequency response that maximized band-limited disturbance rejection.
- **Provide an analysis of closed-loop errors associated with undersampling and aliasing in a hybrid mixed-signal controller configuration:** Aliasing affects from both continuous and discrete sources are quantified on a model of the experimental plant. A bound on the aliasing error is established for both a simply-sampled continuous-time plant, and a plant under the influence of closed-loop continuous-time feedback control in frequency ranges *above* the sampling rate.
- **Demonstrate measured performance of the synthesized controller on an experimental plant that has in excess of 1000 states:** The hybrid continuous-time, multirate discrete-time controller was implemented on an experimental structure. The control computation was performed in hard real-time on current processor technology.

Broadband disturbance rejection was measured, and the individual contribution of each control type was identified.

- **Demonstrate measured performance of synthesized controller on an experimental plant under unknown disturbances and unmodeled plant variations:**

The performance of the control designed under a deterministic model was measured under minor plant variations and shown to remain effective. Also, the disturbance rejection performance for *unmodeled* disturbances was measured. In both cases, the stability of the model was ensured through a combination of choices made in the synthesized control topology, and also through experimentally-measured loop uncertainty.

At the conclusion of this work, the methods described will provide information useful for a reader interested in performing a similar synthesis and implemented vibration control design. The most important question to determine applicability of this control methodology is whether or not the target structural has similar characteristics and design goals. As investigated in Chap. 3, the modal complexity of the structure throughout the bandwidth of interest transitioned from a deterministic regime through to an almost completely statistical region. If the structure were to exhibit primarily deterministic, yet variable behavior, more emphasis might be merited on the high-order model-based control and the robust modeling and synthesis procedures found throughout the literature. Similarly, if the primary disturbance sources causing vibration were due to unknown influences, the application of feed-forward control as used in this work would not be appropriate. The finite computational resource would be better utilized by an increase the model order and sampling rate of the high-order feedback control. If *only* known-disturbance sources are the cause of the vibration however, more emphasis on the feed-forward control may be warranted. A longer filter length on the adaptive controller would allow for better low-frequency performance on a lightly damped structure than a shorter length, computationally smaller feed-forward controller. All of these controllers vie for limited computational resources, so they should only be simultaneously applied if the structural characterizations merit the band-separation approached used in this work.

The continuous-time controller optimized in Chap. 5 must also be considered in context. If the application requires a smaller bandwidth, higher-order compensators could be used to limit the control bandwidth. More aggressive feedback gains could be utilized if the control

bandwidth is small as compared to the actuator bandwidth. For this investigation, the actuator bandwidth was the only practical limit to the desired control bandwidth.

The most important aspect of the control synthesis and implementation considered throughout this dissertation is that while it is tailored to specific implementation limitations, these are more *parametric* values. They do not represent “band-aid” retrofits necessary to implement a theoretical control technique. Instead, they are chosen in the beginning of the control synthesis process and provide *choices* to be made in the synthesis.

Likewise, the control bandwidth choices made by the measured structural characteristics in Chap 3 are not specific to this structure. Any structure with similar characteristics can apply the control synthesis methods presented. The *specific* actuator resonance and roll-off frequencies investigated in Chap 5 are simply parameters used in the optimization algorithm. Another actuator with similar characteristics but different parametric values could just as easily have been used.

The techniques and procedures presented here provide a representative example of how to use measured structural characteristics to choose the controller configuration. This process is general enough to be used for a wide variety of broadband active vibration control problems.

CHAPTER 2

LITERATURE BACKGROUND

The general goal of structural vibration control is to provide a reduction in structural vibration. The source of the vibration can be from an external source or generated from components within the structure itself. The “vibration reduction” is typically quantified as either a velocity or acceleration measured at points throughout the structure.

The research area of active vibration control is extremely broad, varied, and has been investigated for a long time. It is also an extremely difficult problem that combines structural modeling, control theory, and experimental procedure. Because the area is so broad it would be impractical to try to provide a comprehensive literature review of all of the techniques and analysis which relate to or have been applied to the problem. The work described in this dissertation is not an extension of any *one* type of control technique. Rather, it includes a large number of facets relating to structural modeling and characterization, optimal feedback control, adaptive feed-forward control, and an analysis of modeling errors due to implementation constraints. The literature on all of these encompasses a very sizable piece of both mechanical and electrical engineering disciplines. Any practicing control engineer will agree that implementation *specifics* drive a successful control design. The available theory is only as practically useful as the specifics of a problem allow.

The literature survey in this chapter is concerned primarily with providing a list of references that either provide background to a control technique used in this dissertation, or work that describes a specific implementation utilizing one or more techniques for an active vibration control implementation. This attempts to paint the field with broad strokes, while providing an interested reader with specific references as stepping-off points. Where specifically applicable to the background, synthesis, or analysis of the design or implementation of this work, references

are cited throughout. In this way, they are more readily incorporated as related work, rather than being limited to the general overview presented in this chapter.

2.1 Structural Modeling

Structural modeling fundamentally tends to consist of finding solutions to the basic equations of motion. A number of references [10], [11], [12] provide background on the fundamental equations of vibration. These fundamental relationships are partial differential equations in both time and spatial coordinates, and in general provide only Fourier solutions [13]. These *modal* solutions are an infinite summation of frequencies, and provide no closed-form analytical solution. Often, numerical and finite-element methods (e.g. [7]) are used to provide baseline models. Even for simple structures such as a simply-supported beam, the analytical methods typically resort to representing an infinite frequency response with a finite bandwidth [14]. This frequency-domain concept parallels the electrical engineering notion of impedance and is used throughout the structural and acoustics literature. A good summary of the history of these concepts can be found in [15]. Numerical finite element methods have been increasingly used to provide models of complicated mechanical systems without resorting to experimental measurements, but are rarely accurate enough for anything more than the most trivial of experimental implementations [7].

The low-frequency regimes can be successfully modeled with finite-element methods, so long as structural parameters can be accurately described. High-frequency characteristics can be successfully described using statistical methods [16] [17], but the relationship between the low-frequency deterministic characterization of high-order systems and the high-order statistical properties is often referred to in the literature as the “mid-frequency” problem [18]. This transition region between the deterministic and statistical modeling was investigated both analytically and experimentally in [19]. This has led to some hybrid modeling techniques [20] that utilize both finite-element low-frequency techniques and statistical methods together. Recent work [21] has even suggested the “universality” of such statistical modeling related to geometries of substructures such as metal blocks, elastic plates, and membranes. Such “intrinsic” properties of this statistical uncertainty approach are the basis of the frequency-separation control configuration investigated in this dissertation.

2.2 Model-based Control

Active vibration controllers that use models in the design fall broadly into low-order and high-order control. Low-order control assumes some form of average, or worst-case model. These low-order models rely on simplicity or collocation for their robustness. Typical implementations of these are discussed in Sec. 2.2.1 and more generally in 2.3. High-order model control relies on specific, often multivariable models with large numbers of states and is discussed in Sec. 2.2.2.

2.2.1 Low-order feedback control

The simplest form of feedback utilizing single-input, single-output control can be used in some circumstances for structural vibration control. The inherent stability of having a control actuator collocated with a sensor provides unconditional stability under direct rate feedback [22]. Such controls have been successfully implemented [23], and provide significant modal damping when applicable. The single-point, collocated control on a structure has the primary benefit of providing an interlaced pole/zero structure that results in no cumulative structural phase shift [24]. This provides the guaranteed stability under direct rate feedback [22], but realistically cannot be implemented exactly. Collocation limitations were considered in [25] [26] by intentionally introducing small non-collocation between error and actuator placement, as well as nonlinearities. Although some analytical bounds on collocation error were quantifiable [26] and preserved the collocation assumption robustness, they were extremely small and not of a typical type fundamental to the limitation of collocated control on structural control problems. The most restrictive practical limitation on the collocation assumption was that of actuator bandwidth limitations. As expressed in [27], actuators and controllers have a finite bandwidth, but flexible structures have infinite bandwidth due to the wave equations that govern their behavior.

Some more aggressive models were obtained in [28], [29], and [30] to utilize quasi-passive electrical analogs to mimic the structural modes and provide performance with increased controller complexity. Active elements were added in [31] to enhance the order available. They provided good performance, but since these techniques inherently utilize static pole/zero cancellation [29] they are highly sensitive to model mismatch [26]. They also do not work well

when modal density increases due to the additional complexity between overlapping modes. Also, modes present outside the band of interest cause in-band truncation modeling errors investigated in [24] as additional performance and stability-limiting factors.

More recently, experimental implementations that utilize this low-order feedback have been demonstrated in [32]. With collocated piezoelectric actuators and sensors, minimal loop-shaping was required to ensure adequate stability margin. In [33], the same technique was utilized along with layers of structural isolation with their own low-order control to serially reduce vibration at a single point. These examples show that by combining the control together with the structural design, additional performance can be obtained above what is possible with a retrofit design.

2.2.2 High-order feedback control

Despite the complexity of modeling high-order dynamic structures, much work has been done assuming such a model is obtainable [8] [34]. The classical and state-space techniques are considered “model-based” control [26] and as such constitute the bulk of the theory available on active vibration control. A tremendous amount of theory behind state-space-based models developed throughout the 1960’s and 1970’s guaranteed some important properties such as controllability, observability, and synthesis techniques for state-space feedback. Unfortunately, the theory often fell short when implemented. In particular, implementation of these techniques for LFSS’s was often impractical [9] due to the extremely high order required, model truncation and other variable modeling errors [35] [26] [36] [37] leading to the development and application of robust control techniques [38]. Simulation results follow the theory [39], but break down when implemented realistically. Since optimal model-based control techniques inherently rely on pole/zero cancellation [29] for their implementation, modeling the structural dynamics with closely-spaced poles and zeros is extremely difficult to do robustly [34]. The Fourier solution to the wave equation that governs structural dynamics means that modes are *harmonic* in nature, so errors become cumulatively worse as modeled frequencies become higher. Despite the disparity between theory and practice, some systems could be utilized despite modeling and implementational mismatches [37].

The SISO controllers of Sec. 2.2.1 can provide some performance for suitably defined structural vibration problems. If the number of actuators and sensors is increased to where their interaction cannot be ignored, fully-coupled multi-input, multi-output analysis may be neces-

sary. The classical state-space LQR control was generalized into LQG control [40] for stochastic systems. A formal analysis of integrating state observers with the deterministic quadratic optimal control problem was investigated in [41], with practical considerations extended in [42]. A tutorial overview of the LQR design procedure can be seen in [43] in [44], and again in [45].

The stochastic approach of LQG and later-formalized H_2 control was followed by the robust H_∞ control [46] application to the LFSS problem [47] [48] [37]. While the robustness tradeoffs of H_∞ make the stability more assured, the required modeling of the uncertainty often results in poor or no performance [49]. The lightly damped structure has large dynamic range, and very small errors in modeling have been shown to have very large modal contribution [50]. Robust eigenvalue perturbation methods were pursued in [47], but on a modally sparse and low order structure with very small perturbation. As a consequence of this difficulty, mixed H_2/H_∞ methods have been pursued [51] [52] [53], with representative experimental implementations in [54] [55] [56]. These methods attempt to add the robustness parametrization of the H_∞ analysis, with the optimal design criteria of the H_2 synthesis.

The decentralized control [57] [58] [59] problem formulation was applied to the LFSS problem in [60] and [61]. Since the decentralized control is a specialized form of the centralized control problem [62] however, it is an even *more* restrictive special case of high-order model-based control. The centralized problem has robustness difficulties more fundamental [50] than the implementation problems aided with decentralized theory.

The most promising form of high-order model based control combines the high-order control with low-order control. This type of multi-loop control is referred to as HAC/LAC, for High-Authority Control/Low-Authority Control [63]. The high-authority is high-order, model-based control in an outer loop. The low-order control is an inner loop that tries to apply active damping to the structure to make the high-authority controller less sensitive to modeling errors [28] [64] [65]. When the system order and modal density become high enough however, even these methods are unsuitable. Even the increase of modal damping does not sufficiently aid analytical models with hundreds of variable and dense, overlapping modes. The LAC component may still work, but the HAC does not have a sufficiently well-known model to work.

A good summary of the layered approach to control applied specifically to the LFSS problem can be seen in [66]. The typical performance measures are quantified, as well as some of the realistic limitations to the high-order, model-base control. This summary was expanded in [67]

and [32] to include hierarchal, multi-layered control to the LFSS problem. In both cases, local feedback loops were successful in applying significant active damping to a number of modally sparse, lightly-damped structural resonances and high-order control was applied after these loops were closed. In [33], the layered approach was extended to include position feedback and active support actuators. These additions provided impressive performance, but only over a very small bandwidth and acted on just a few modes. The work of this dissertation extends the hierarchal control simultaneously and substantially in bandwidth, system order, measured structural modal overlap.

2.3 Dissipative Control

In parallel with this investigation in formal high-order model-based control development, the robustness problems associated with the LFSS problem were investigated through energy-based and dissipative control. The benefits of control not based on modal truncation or high-order models to provide dissipative control was investigated in [68]. The collocated requirement for dissipative control was investigated in [26] and [63] and shown to be effectively a strict requirement, with only very small allowances made for misalignment and actuator nonlinearity. Experimental implementations of the dissipative control were shown in [69] and [70]. Slight restrictions on the LQG algorithm were shown in [71] to provide a dissipative control result.

While not specifically required to be classified as dissipative control, high-authority/low-authority control (HAC/LAC) often utilized dissipative control for the low-authority loops. Local rate feedback was used in [7] and [50] as a LAC, with an H_∞ robust control as an outer HAC loop. This provided modal damping to facilitate the H_∞ control and was tested through simulation. The primary benefit of the HAC/LAC technique was to provide the robust controller with a plant more conducive to setting standard robustness bounds. The application of dissipative-type limitations on model-based control was investigated in [72] [69] [73] [74] [75].

2.4 Feed-forward Adaptive Control

When the source of a disturbance is available for a vibration control problem, feed-forward control presents a viable means to reduce its effect. This type of control requires a model of the disturbance path in order to provide an equal and opposite effect at the point of measurement.

While static feed-forward systems are possible in theory, for any real system with nontrivial dynamics, feed-forward systems must also be adaptive. Due to the sensitivity of feedback systems to closed-loop stability, feed-forward adaptive systems are often more practical than feedback adaptive systems. Although some work has been done to show the equivalence of adaptive feed-forward controllers as equivalent feedback systems [76] [77] [78], the feedback path of this equivalence is at the *slow* dynamics of adaptation [79]. Thus, for a feed-forward system that has adapted to a stable controller path, the closed-loop system will also be stable. For a feedback system, this is not necessarily the case since the loop-gain of controller and plant is *not* guaranteed to be stable even if the controller is. Even with these limitations, adaptive feedback systems have seen some applications to structural/acoustic modeling [80], but so far only on relatively low-order systems.

Since DSPs have been able to implement adaptive digital filters at real-time speeds, feed-forward control has provided a mechanism for significant performance gains where applicable. The adaptive control theory literature contains rich theoretical results for model convergence rates, sufficiently rich input excitation, and indirect and direct control implementations, [81], but practical considerations generally prevail for implementation purposes. The computational requirements for a number of different online approaches are found in [82], and explain why the gradient-descent approach of least-mean-squares (LMS) is typically used in implementation as an approximation of the recursive-least-squares (RLS). Good tutorial overviews of both single channel and multichannel active feed-forward control can be found in [83] [84]. Another good overview with a number of experimental applications can be reviewed in [85]. Analytical investigations on implementational problems with actuator path modeling errors are investigated in [86] and describe the filtered-X LMS (FxLMS) algorithm most often employed for structural and acoustic control problems.

The performance qualities of the LMS algorithm are the primary reasons why it enjoys such popularity. With a relatively modest requirement that the control actuator path be within 90° of the experimental model [87] [88], convergence is guaranteed. For control actuator paths with larger phase deviations than this, a modification of the simple LMS algorithm known as the filtered-X LMS is often used [89] [86]. The addition of a leakage parameter in the algorithm [90] [91] [92] was used to aid the robustness qualities of the LMS algorithm even further. The transient performance of these modifications was considered in [93] [78] [94]. Comparing the

computational requirements of different LMS variations to other adaptive control types was done in [94] [82] [95] [96]. Extending to recursive filters removes the problem of a truncated impulse response that plagues FIR implementation of LMS. These adaptive IIR filters have seen some limited success [97] [98], but generally are too computationally expensive and unstable to use in general cases.

CHAPTER 3

PROJECT BACKGROUND

Chapter 2 discusses some of the very wide array of control techniques that have been applied to the large, flexible space structure problem. This investigation was to consider a structure that was representative of a complicated, large, flexible space structure. The application could be an orbiting spacecraft with large booms or telescoping actuators that require vibration minimization at the remote end. The applications are not limited to orbiting structures, as large, lightweight, lightly-damped booms have significant uses as manipulators in earth-bound applications as well. For the purposes of this investigation, the control techniques chosen were due to a combination of the original project specification requirements, experimental structural limitations, and an emphasis on the implementation feasibility of the resulting control strategy.

Section 3.1 discusses the specifications for the control problem and the assumptions that were made. They provide the basis for control bandwidths and types. These assumptions also entail tradeoffs with computational requirements, modeling accuracy, types of disturbance rejection, and control system robustness.

Section 3.2 discusses the experimental testbed. By taking initial characterization measurements on the structure, choices about actuator location, orientation, and bandwidths could be made.

Section 3.3 continues the investigation of the experimental structure by including the tools and techniques used to produce an identified analytical model. The data collection and compilation is discussed, as well as the practical attributes of the identification toolbox used. The most important aspect of this was the model order requirements for accurate identification.

Section 3.4 includes some of the simplified models that were required to obtain low-order approximations to different facets of the actuators and structure. It includes an analysis of

the electromagnetic inertial actuators used for control under ideal conditions. Also, analytical models of simple plate structural vibrations were used to provide a baseline estimate of the expected modal complexity of the structure. The modal density of the structure was measured and discussed as a means to separate the control into different bandwidths.

3.1 Control Problem Defined

The potential control aspects of a structure as complicated as the one of interest here were endless. The research goals for the control were broken down into performance requirements, assumptions made about the configuration of the structure, and limits associated with the implementation. This set of conditions and limitations made implementation feasible. It was not a trivial laboratory demonstration, nor was it so completely general and optimistic as to be unimplementable.

3.1.1 Performance requirements

The performance requirements for the structural control were essentially to provide the largest amount of disturbance rejection possible over the widest possible bandwidth. This specification was a bit unusual. Normally, certain bandwidth requirements are specified at the onset of the project. In this case, it was a “more is better” requirement. With that as a driving choice, the control bandwidth was primarily dictated by the bandwidth of the actuator. Electromagnetic inertial actuators were chosen for their robustness, large displacement, and relatively low-frequency response capabilities. Unfortunately, mass limitations restricted the actuator to a physical size that allowed for approximately 100Hz resonance. This provided the lower end of the bandwidth range. The upper end of the bandwidth requirement was ultimately determined by the actuator’s high-frequency performance. It could be extended to some degree with the controller, but to avoid destruction from resistive heating, the practical limit was approximately 5kHz.

Another performance requirement was that the resulting controller be robust to small variations in the system. The types of variations that were to be modeled were consistent with small manufacturing tolerances. Differences in bolt torques, weld irregularities, structural damping variations, and placement tolerances in the load were all potential sources of variation. These

variations were to be considered “second-order-effects,” and be exhibited as small fluctuations in the magnitude response. This mean and variance to the structure was the underlying goal for the uncertainty. The data measured from the experimental structure was also used to verify hybrid statistical energy analysis modeling techniques presented in [17]. The results of the control implementation were presented in [99].

Another assumption that was made was that a single point of control was the target. This assumption is valid for certain geometries of construction where disturbance energy could be rejected at choke-points. Also, instrumentation loads located in a small area structure could be required to have a level of vibration reduction greater than at other locations on the structure. These critical points on the structure were desired to have the most isolation from any circulating disturbance energy.

Another important assumption that was made about the intended application of the structure was that the disturbances that were exciting it were coming from a known source. Another source of excitation (motor drives, linear actuators, stepper motors, pumps, etc) would cause the bulk of the energy injected into the structure. The location of this disturbance was fixed, although the propagation path would be subject to the same small structural variation from the manufacturing variances. Since the disturbance was known however, it was usable as a decoupled reference signal for the run-time control implementation. The major ramification of this assumption was that adaptive feed-forward control was a viable technique.

In addition to the known disturbance, the possibility of unknown disturbances was not completely discounted. The unknown disturbance would be assumed to come from anywhere on the structure, and have any possible spectral characteristics. Practically, this meant that a feedback type control was also desirable.

3.1.2 Implementation limitations

The desired control characteristics and assumptions in Sec. 3.1.1 tailored the control topology to a certain degree. Implementation limits were also a significant consideration. One big implementation limitation was that the control must be implementable as a non-intrusive addition to an existing structural design. This significantly limited the types of actuators available, since structural members could not be cut and replaced with active members as was done in the layered approach of [33]. Also, no absolute reference frame was available. That practically

meant that either inertial actuators or strain-type piezo-electric patches were the only potential control actuators. Although piezo-electric actuators have significant force/strain capabilities, they have limited excursion so effectively work only at high frequencies where structural wavelengths are smaller. This drove the control design to use readily-available electromagnetic inertial actuators.

The other significantly limiting factor on the control design was a limit to the computational processing power required. High-order control systems are generally only feasible using digital controllers. These have the significant disadvantage of requiring high oversampling rates compared to the control bandwidth of interest. Generally, this is due to required anti-aliasing filters and holds that introduce phase delay into the control loop and jeopardize performance. In the case of this experimental setup, computing a control bandwidth of 5-10kHz would require a sample rate 5-10 times that. This 25-100kHz sampling rate would be too high to compute high-order models in hard real-time. Even if feasible with an ambitious computer, code optimization, and other tricks, the argument could be made that there is *always* a computational budget. More computation allows more bandwidth and/or higher order systems, but a computational limit will always be present.

With this in mind, minimizing controller computations became a significant driving factor in the control design. This led to discrete-time control only for the low and mid frequencies. It also led to minimal oversampling, and dealing with the inevitable aliasing and hold artifacts that entailed. This is investigated in the analysis of Chapter 6


3.2 Experimental Test Setup

The structure chosen to represent the complex, lightly-damped resonant platform is shown in full in Fig. 3.1(a). It was constructed of round tubular aluminum struts 25mm (1in) in diameter for the vertical trapezoidal structure. The box frame shown at the top is constructed from 25mm (1in) square tubing. The total height of the structure is approximately 2 meters. All joints are welded, and the bottom plate is 12.5mm (0.5in) thick. The plates connected to the top box structure are thin aluminum.

Three potential control actuators and a disturbance actuator were placed as shown in Fig. 3.1(b) and 3.1(c), respectively. The actuator/sensor pairs were oriented in the X and Y axis



(a) Complete structure suspended by elastic cord at a single point



(b) Actuators with collocated accelerometers oriented along three different axes in the corners of the bottom plate



(c) External disturbance actuator oriented in the vertical axis on top of the structure

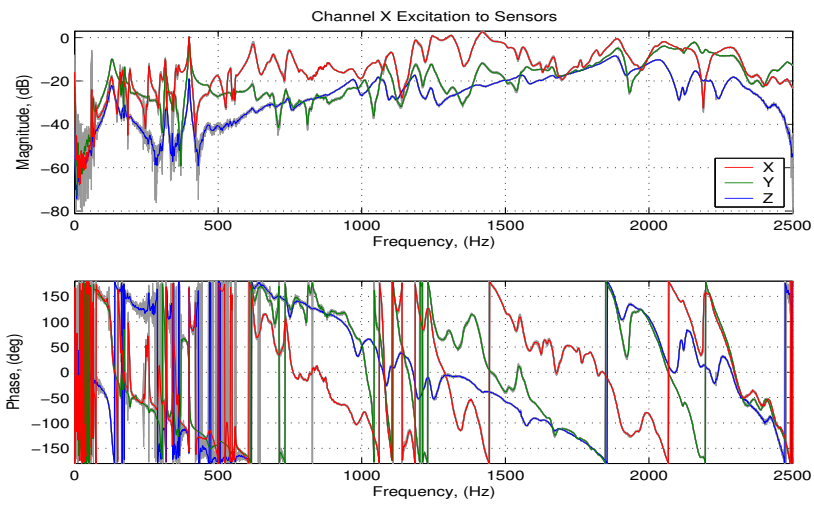
Figure 3.1 Test structure to represent a large, flexible structure

(coplanar on the bottom plate), and Z axis (normal to the bottom plate). The plates on the structure had randomly placed masses and some damping material to simulate structural variation. This “fuzziness” of the plates was varied by moving the masses and taking multiple excitation measurements between the actuators and accelerometers. Due to convention in active noise and vibration control terminology, the path between the disturbance actuator at the top of the structure to the error accelerometers is referred to as the “primary error path.” The control actuators to collocated error accelerometer paths are typically referred to as the “secondary error paths.”


In order to determine how much variation there was in the structure by varying the fuzzy plates, an ensemble of plate variations was performed and the results compared. Because this was a preliminary investigation to get a baseline set of dynamics for the structure, a samplerate of 5kHz was chosen, and all potential actuator locations were excited in turn. The structure was excited by bandpassed white noise through the control and disturbance actuators and the frequency response characteristics were computed at all of the error accelerometer points. The movable masses on the “fuzzy” plates were moved to other randomized locations, and the measurements repeated.

Figure 3.2 shows an example set of experimental ensemble transfer functions between each control actuator and all error accelerometers. Figures 3.2(a)-3.2(c) show the transfer function from each actuator (X , Y , and Z -axis orientation) to all collocated error accelerometers. The ensemble data is shown in light gray and the mean over the ensemble is shown in a solid color. It can be seen that the uncertainty in the secondary path is very small, because the plates have relatively small mass and are physically located a significant distance from the active control area at the bottom of the test setup.

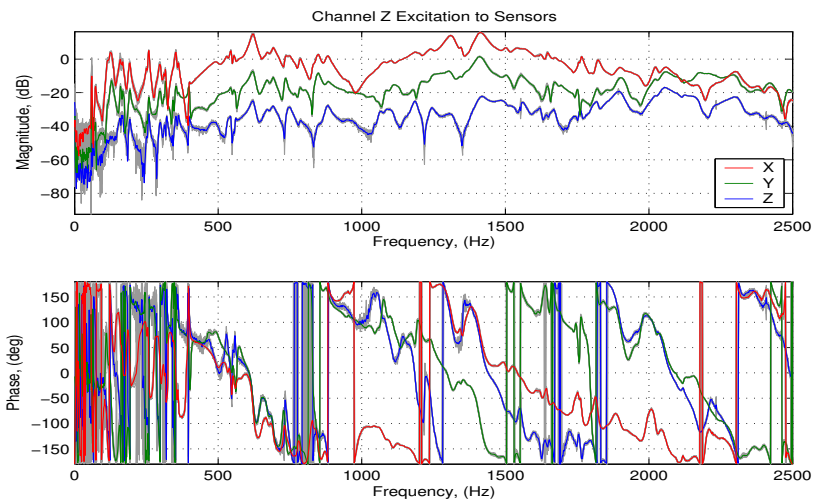
Figure 3.2(d) shows the path from the disturbance actuator to all error accelerometers. Because of the considerable phase attributed to the propagation through the structure, only the magnitude response for the disturbance path is shown. Phase transitions in excess of 20π were identified to all error sensors. It can be seen from this figure that the variation about the mean is greater than for the secondary paths of Figs. 3.2(a)-3.2(c). This is to be expected as the path from the disturbance to the error sensors traverses the entire structure and must propagate through the points where the plates are attached. This would make the interaction



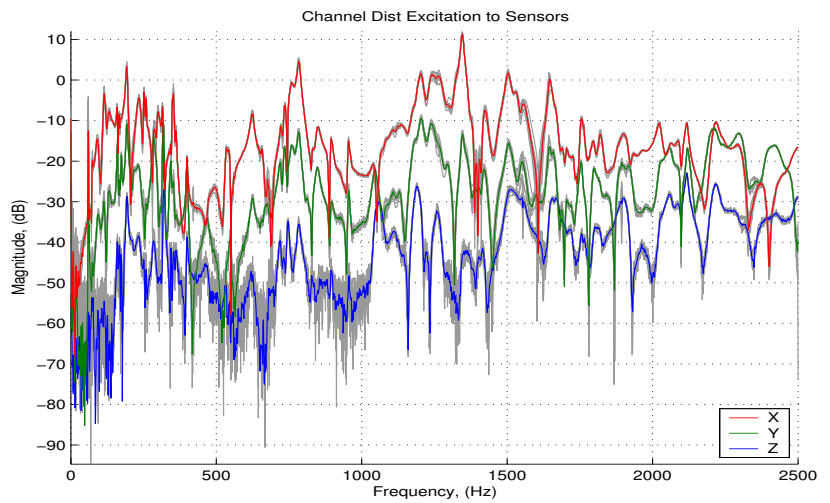
(a) X channel excitation



(b) Y channel excitation



(c) Z channel excitation



(d) D channel excitation

Figure 3.2 Actuators to error accelerometers instances (gray) and mean (solid).

with the "fuzzy" plates more significant. The experimental results from this investigation were used to assist in verifying the modeling analysis presented in [17].

The overall dynamics of the structure as seen by the three sets of actuator/error sensor pairs has significantly different frequency characteristics. In particular, the Z -axis error sensor exhibits the most interesting characteristics. It is not overly dominated by any particular group of resonances. Rather, it exhibits spectrally rich character throughout the entire frequency range that is not dominated by any one band. It also shows the most amount of variation with respect to changes in the fuzzy plates. This is most likely caused by the direct coupling of the lower plate to the upper frame through the compression/tension in the support members. The X and Y axis are only coupled torsionally through them, so the plate variations would be attenuated through this propagation path. For these reasons, the Z -axis actuator/sensor pair was the one chosen to investigate further as a representative, single-point of control. For the remainder of this dissertation, the Z -axis, collocated actuator/sensor pair is the primary one considered.

3.3 Model Identification

The frequency-response data for the structure provided a baseline for characterization and enough information to choose the desired control axis. An analytical model identification was necessary to perform a modal analysis. It also was required for the control synthesis and analysis of Chapters 4 and 6. This was a very important aspect to the investigation because the system was of such a high dynamic complexity. One toolset in particular was developed to deal with experimental identification of high-order, lightly-damped, flexible structures. The work of [100] was extended and written as a Matlab toolbox for performing frequency-domain identification in [101]. This was further enhanced to include other algorithms with the SOCIT (System, Observer, Controller Identification Toolbox) of [102].

This toolbox provided a number of potential identification algorithms and permutations. The one chosen for the test structure was the Observer/Kalman filter Identification (OKID), as it uses time-domain data alone. It also had the advantage of allowing multiple experimental runs to be concatenated together to improve the model estimate. This feature in particular benefited the multiple-input aspect of the system ID performed. Rather than excite all inputs to the system simultaneously, the stacking feature of the toolset allowed for the excitation

of the control and disturbance actuator separately. The other input was held at zero. The resulting time-domain data was then stacked as though all but one of the inputs was being “excited” with zero. Of course this superposition assumes a linear system, but a linear model was what was desired, and the excitation gains were set so as to minimize structural and actuator nonlinearities.

Practically, the single-channel excitation allowed for improved measurement coherence. The reason for this was two-fold. Excitation one at a time removed the variable of *which* input actuator provided the response at the error sensor. It also aided to keep the structure from being driven into nonlinearities such as “rattling” of the plates. The maximum energy before this nonlinearity occurred was injected by only *one* actuator at a time.

One other improvement on the SOCIT identification toolbox was the addition of the residual whitening iteration term. From [100], this procedure reduces the model’s residual in a least-squares sense and iterates the solution a selectable number of times. It was noted throughout the characterization of the test structure that this aided in both the quality of the resulting model, and the likelihood of causing numerical problems. In particular, the high-order models were less likely to numerically diverge if this whitening iteration was used. After a number of iterations on the parameters for the identification toolbox, a whitening iteration value of 10 yielded satisfactory convergence while minimizing the identification time requirements. Even at this, identifying the 1000-state model using only 83 seconds of time-domain data took overnight on a 2.4 GHz dual processor personal computer.

3.3.1 Identification results

The same procedure was used to excite the system as in the initial characterization of Sec. 3.2, except that only the chosen actuator (Z-axis) and disturbance channels were excited. The anti-aliasing and smoothing filters were both Ithaco model 4302 24dB/octave Butterworth filters set at 4.00kHz low-pass, with a 10kHz sampling rate. The higher rate was necessary to generate a reasonable characterization to dynamics that would eventually be above the sampling rate of the discrete-time computer controller. After stacking the individual excitation data, the input/output time vectors were run through the SOCIT toolbox. The OKID routine with 10 whitening iterations was performed with model orders ranging from approximately 100 to 1000. The whitening variation was included as an additional scalar parameter to the

toolbox and resulted in a more accurate and less numerically sensitive result. The frequency response of each of the resulting discrete-time, state-space models was then plotted against the the frequency response of the original time-domain data. Figure 3.3 shows the quality of the various model orders for the input-output path of collocated actuator to error sensor.

The lower plot in Fig. 3.3 shows the error between the spectral response calculated from the raw data to that of the various model orders. As can be seen, even a model order of 400 states was insufficient to prevent significant deviations in magnitude response. The phase response was equally erroneous with anything less than the large model of 996-states. The disturbance to error path magnitude plot of Fig. 3.4 was similar, although there were many more identifiable structural resonances. Both of these plot the *same* 2-input, 1-output identified model of the plant. Figure 3.5 shows the wrapped phase plot for both the actuator and disturbance path. The actuator path exhibits primarily the phase roll off characteristic of the anti-aliasing, smoothing filters, and hold. The disturbance path exhibits that filtering phase, but also the tremendous phase due to propagation through the structure.

One other attribute that was seen with using the SOCIT toolbox was the identification routines did not perform very well unless given a very large number of states. If the number of states for the ID was increased beyond the target order and then reduced to the target order after running the ID, the resulting model was superior. Figure 3.6 shows a close-up section of the frequency magnitude response for the disturbance path. As can be seen from the frequency response, the over-determined model reduced to 198 states is more accurate than that from the identification engine directly set to 198 states. Note that although Fig. 3.6 shows a truncated frequency response for clarity, the model order reduction was over the entire bandwidth and the resulting accuracy was similar at higher frequencies.

In order to examine this phenomenon further, additional information from the model-order reduction was examined.. The model-order reduction began with a transformation of the system to a balanced realization. This had the effect of ranking the “importance” of the states’ contribution to the controllability and observability [103]. With this new but equivalent balanced realization of the system, only the most important states up to the desired number of states in the reduced-order model were kept. The vector of Hankel singular values (HSV) is a metric of the states’ importance and was also returned by the balanced realization calculation. Figure 3.7 shows the magnitude of these HSV’s for all identified models, as well as the model

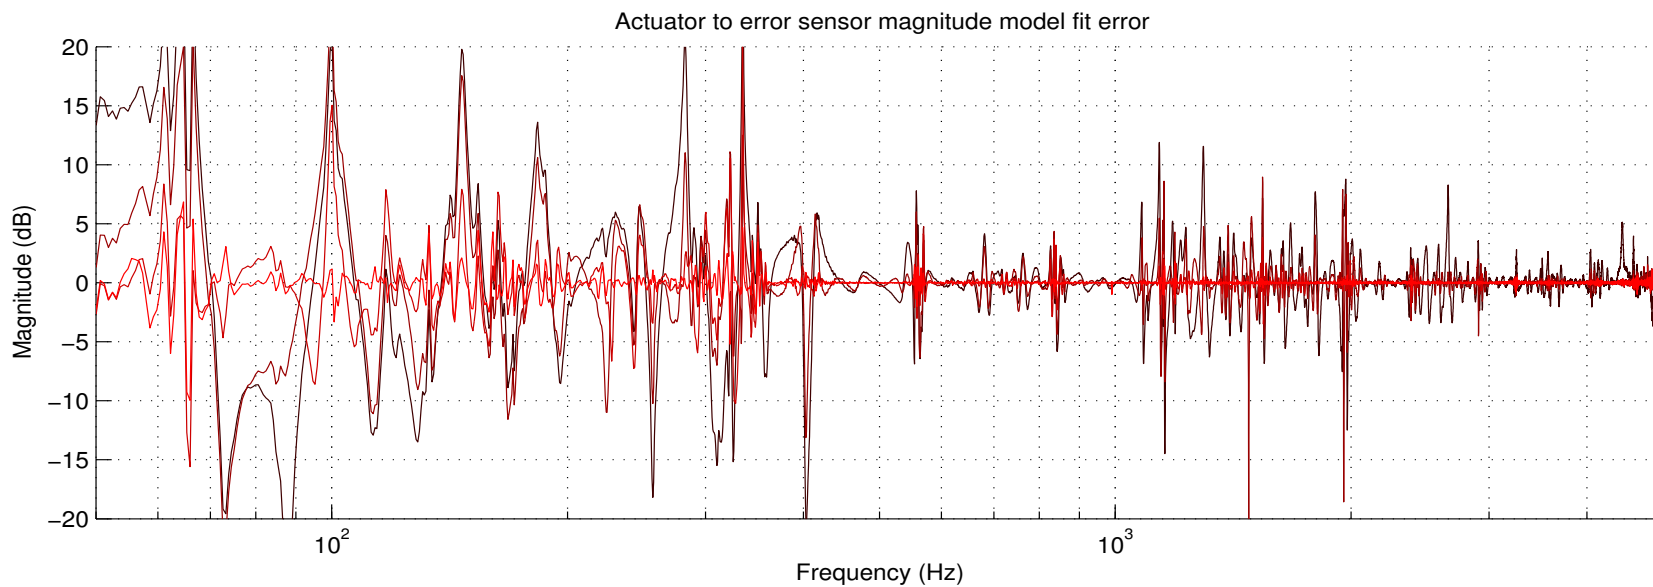



Figure 3.3 SOCIT toolbox system ID (actuator path)

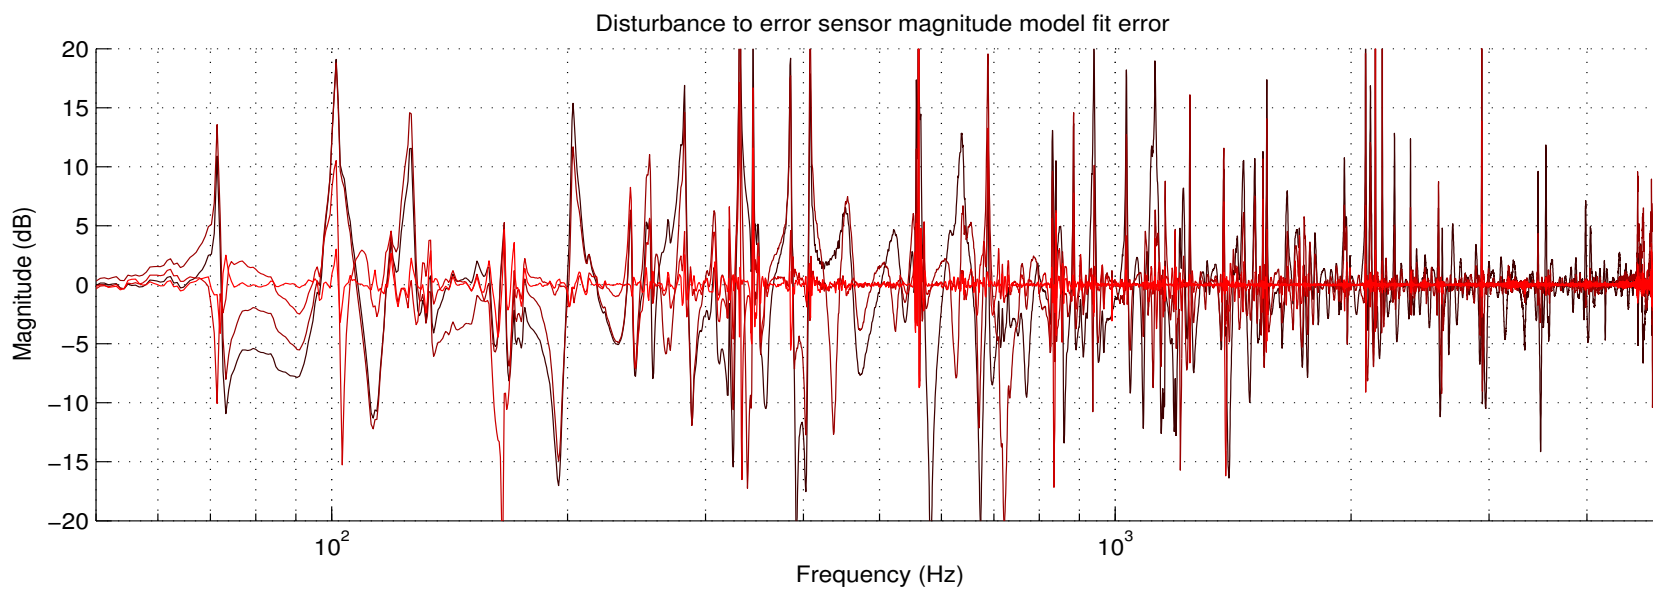



Figure 3.4 SOCIT toolbox system ID (disturbance path)

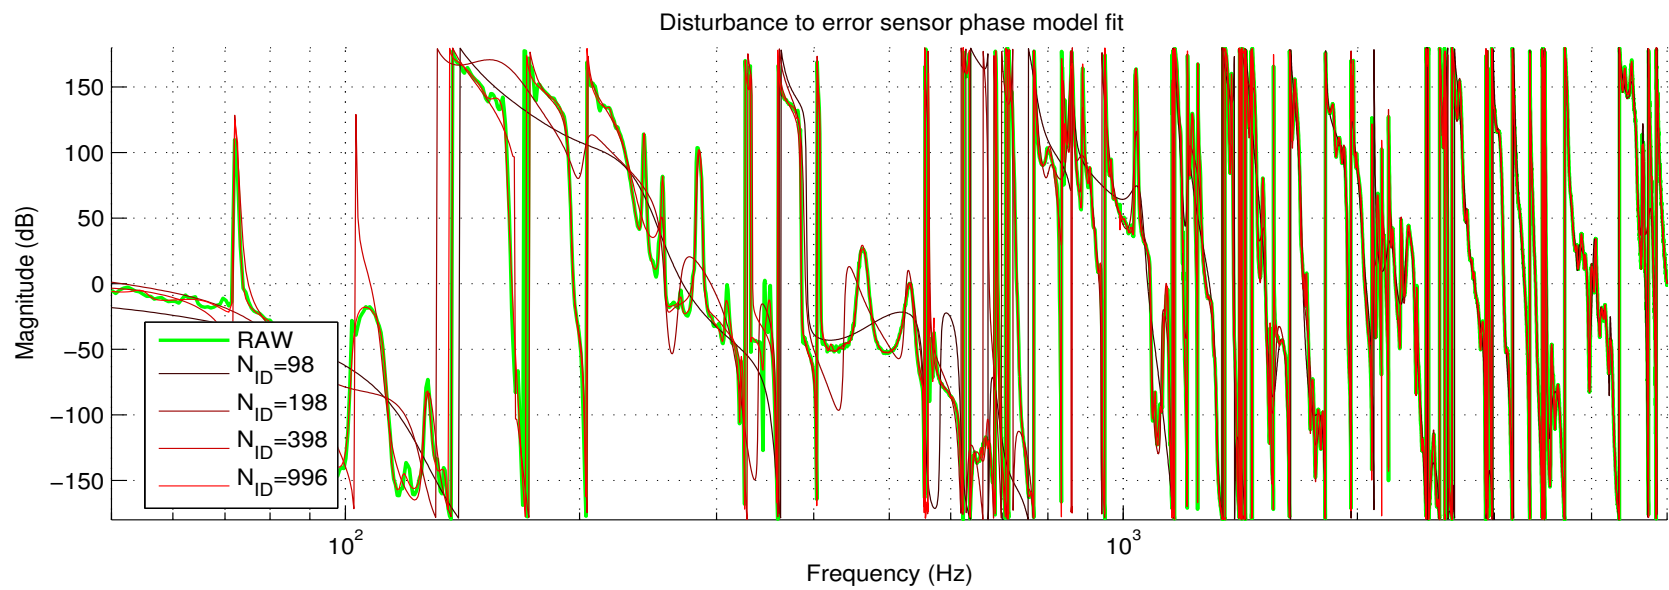
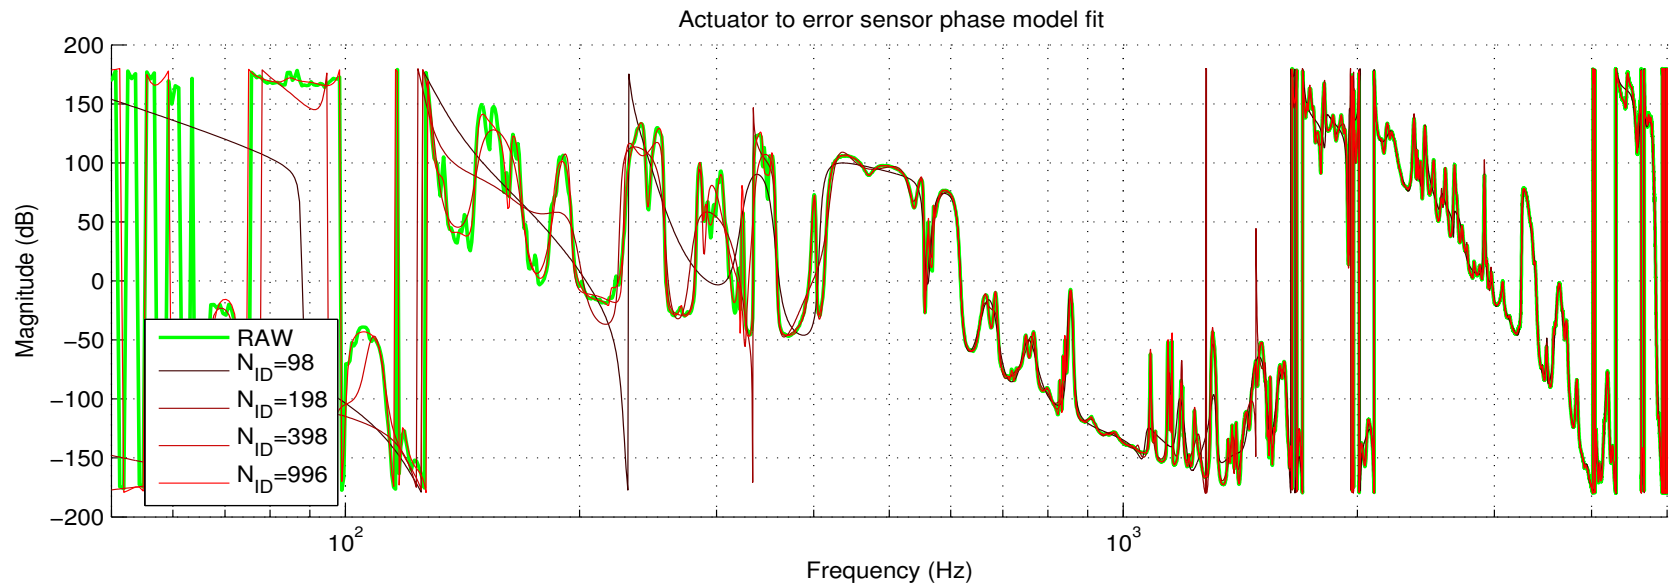


Figure 3.5 SOCIT toolbox system ID phases

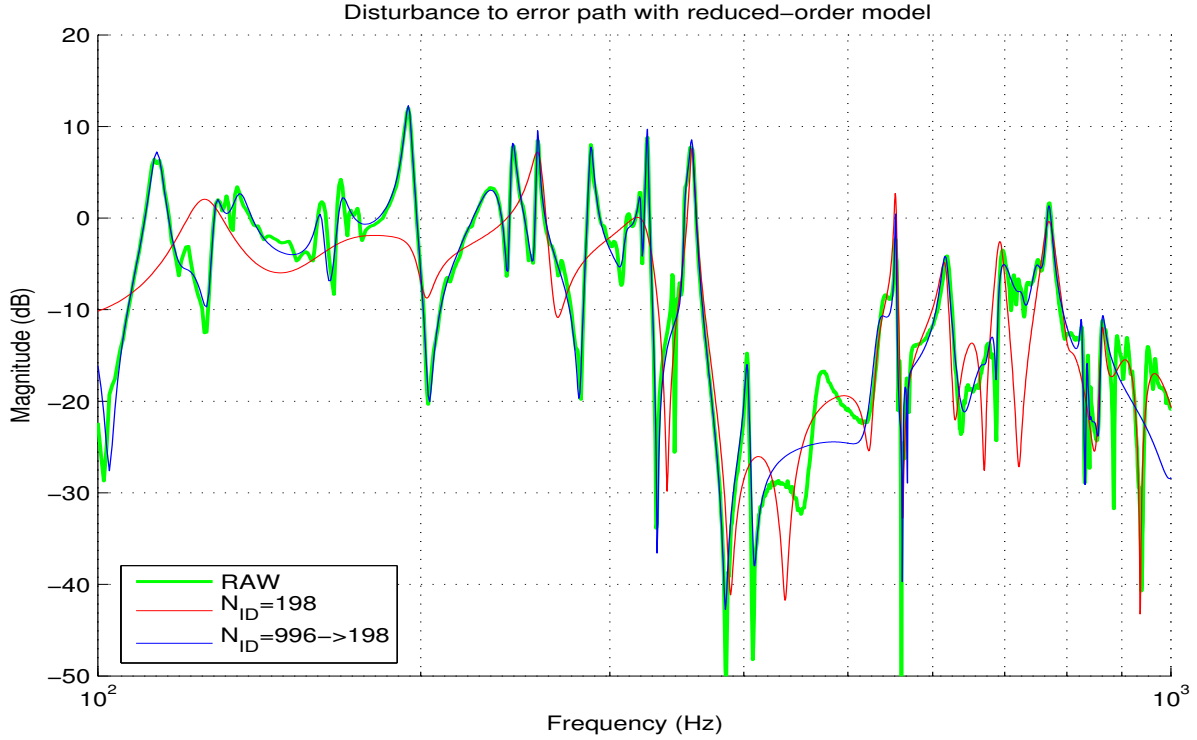


Figure 3.6 Example (disturbance path) model-order reduction from high-order model

reduced from 996 to 198 states. The result is that the reduced-order model stays closer to the original models' HSV through a larger number of states. This would imply a greater accuracy to the experimental data, so long as the higher-order modes were already considered a “better” fit. It should be noted that this more accurate requirement to allow for additional degrees of freedom in the system identification posed a very real limitation to the size of the model. Many attempts at running identifications in excess of 1000 states were run, but resulted in numerically divergent solutions. This meant that a 10kHz samplerate (5kHz Nyquist rate) was the practical limit to the 1000-state ID of the structure.

3.3.2 Measurement artifacts

The quality of the measurements taken of the experimental structure also influenced the quality of the system identification. The measurement coherence is the quantity used to gauge this quality. The coherence of the time-domain data from each experimental excitation to the error sensor was estimated and shown in Fig. 3.8. This computation considered overlapping sections of 5.2 seconds of input data, which consisted of 32 averages of 16384 points at a 10kHz

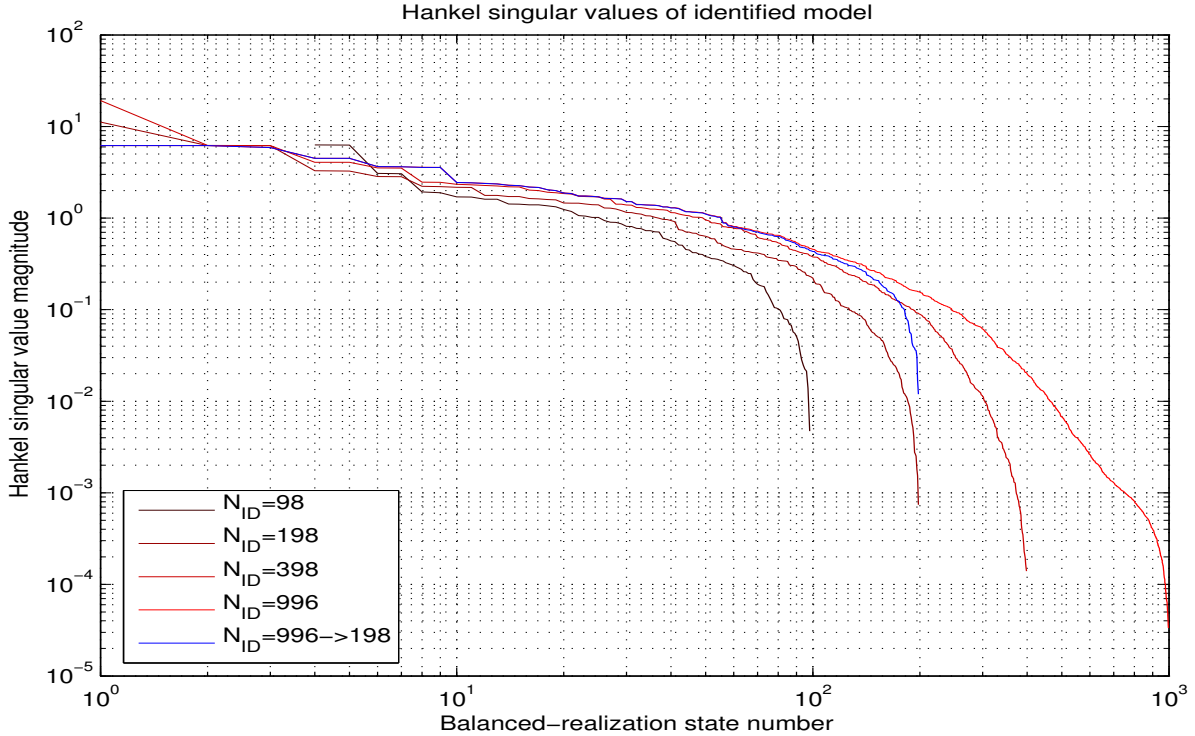


Figure 3.7 Hankel singular values of identified model

sample rate. This provided sub-Hz resolution in the estimate of the measurement coherence. As can be seen from the plot, except at locations of zeros in the transfer path, the measurement quality for each excitation was very good. Comparing the coherence to the measured frequency response plots in Figs. 3.3-3.4, the apparent signal to noise ratio of the test setup was approximately 40-50dB.

3.3.3 Analytical model characterization

Another useful characterization of the analytical model of the plant was to examine the pole/zero locations and structural damping. The specific pole/zero locations would be important in the control design synthesis, but the *trends* in their location give a more intuitive feel to the modal structure. In particular, the pole/zero trends provided some insight into what bandwidths and control structures would be appropriate for the synthesized controller.

The system was ultimately continuous-time, but was measured and identified as a discretized, sampled-data system. To plot the pole and zero locations more intuitively in complex plane rather than the unit circle, it was converted to a continuous-time model. It was con-

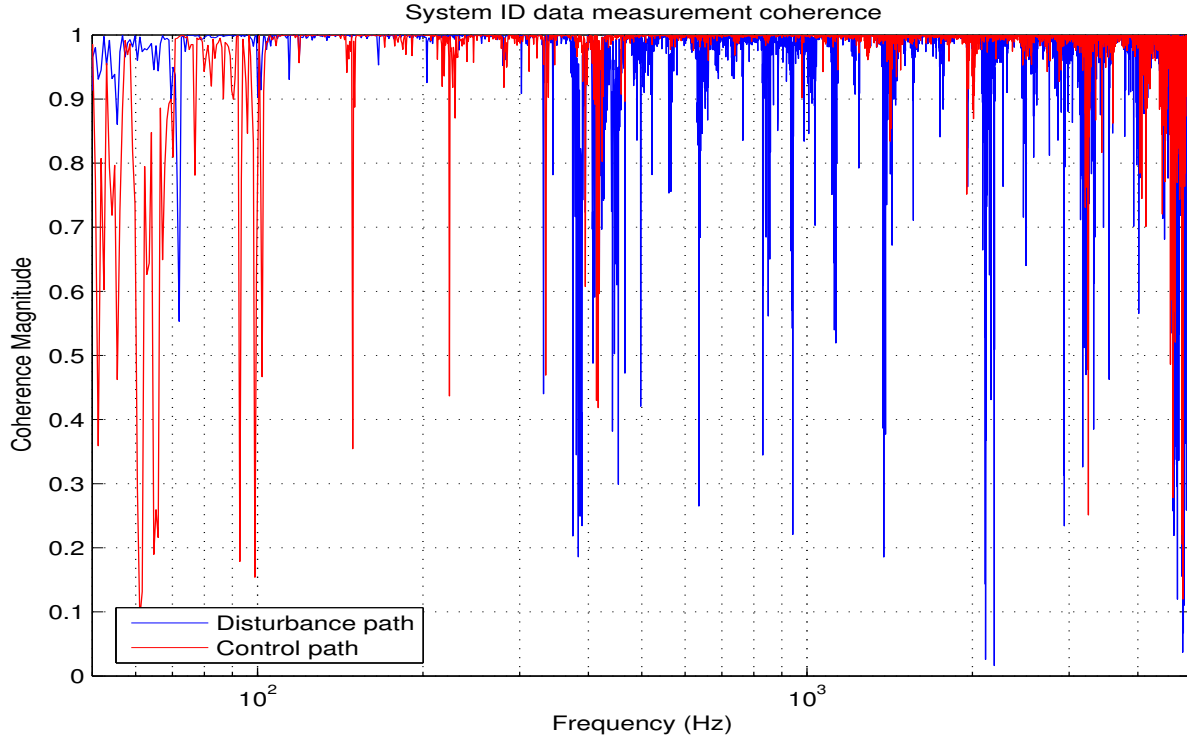


Figure 3.8 Measurement coherence for experimental data used in model fit

verted with a continuous-time zero-order hold equivalent [104] [103] on the 10kHz “oversampled” model. As a general characterization of the structure, the subtle differences that arise from continuous and discrete mappings were not as important as when implementing a control system, so this continuous-time model was considered adequate. The zeros were separated into minimum-phase and non-minimum phase components to be analyzed separately. Figure 3.9 shows the actuator path pole/zero mapping in this framework. Each pole and zero’s real part is normalized to its magnitude and plotted as an “o” or “x” for zero and pole, respectively. The smaller the ratio of real to magnitude, the closer to the $j\omega$ axis the pole/zero lies and the less damped it is.

Another unfortunate artifact of the system identification was that of the non-minimum-phase zeros shown in red. The actuator path exhibits some of these, but the disturbance path in Fig. 3.10 shows many more. It was not obvious whether or not these non-minimum phase characteristics were due to an actual property of the structure, or whether they were a measurement or ID artifact. The 996-state model order was used to compute both the actuator and disturbance pole/zero maps. The high-order model exhibited the lowest percentage of these




Figure 3.9 Actuator path pole-zero map

non-minimum-phase zeros in the actuator path. This feature alone tends to imply that at least in the actuator path the right-half-plane zeros were not due to a real phenomenon. Also, repeated experimental measurements produced different results with fewer, greater, or different-location non-minimum-phase zeros. Figure 3.11 shows the complex-plane mapping of only the non-minimum-phase zeros of the disturbance path at different model order approximations. None of the locations of these non-minimum-phase zeros are consistently identified. This implies that they were due to measurement coherence limits, unmodeled digital delays, or limitations in the phase identification characteristics of the SOCIT toolbox. If the non-minimum-phase zeros were common and clustered around specific points in the s -plane, it would imply a *true* non-minimum-phase zero.

Figure 3.12 also shows the phase difference between all of the different model orders. This is the same plot as Fig. 3.5, only the phase was *unwrapped*. Although the wrapped version shows good phase match wherever the gain was high, the zeros were where *none* of the analytical models matched completely. A relatively clear example of this was at approximately 200Hz on the disturbance path. Figure 3.10 shows a non-minimum-phase zero, where the magnitude




Figure 3.10 Disturbance path pole-zero map




Figure 3.11 Disturbance path RHP zero map of different IDs

plot of Fig. 3.4 shows a deep zero. From these plots of the quality of the system identification, it remained unclear as to whether or not many of the non-minimum-phase zeros were due to physical phenomenon or from a combination of effects in the measurement and modeling processing. Even if they were not physically realized, they presented some challenges to the analysis of the final synthesized controller and are further discussed in Chapter 6.




Figure 3.12 Model ID absolute phase (unwrapped)

The pole and zero locations provide much information, although not necessarily as intuitive as structural damping. Figure 3.13 shows the damping ratio of each of the structural modes in the 996-state identified model. Aside from a few modes, 10% was the upper limit in structural damping, with the majority of the modes above 500Hz exhibiting less than 1%. This very lightly damped character was what was desired as a representative structure, and the 1% measurement fit the requirement well. Although not used directly, this characterization was used to compute the modal overlap of the structure in Sec. 3.4.4.




Figure 3.13 Measured structural damping

3.4 Approximate Modeling

At this point, it was worthwhile to consider *simplified* models of some components of the system. The system identification of Sec. 3.3 provided a full-order, identified model of the structure from experimentally-measured data, but these were too complicated to consider for any classical control framework. Rather, idealized, representative models of simple plates and actuators were considered. For some parts of the control synthesis in Chap. 4 as well as in the optimization in Chap. 5, low-order and asymptotic models were beneficial.

This section deals with *approximate* modeling of some of the components. An analytical model of an electromagnetic shaker both with and without dynamic load is considered in Sec. 3.4.1-3.4.2. Since the electromagnetic shaker was not a true point source, it was also necessary to consider a simplified model of the propagation through the plate it was attached to. This simplified model is considered in Sec. 3.4.3. Finally, in Sec. 3.4.4, this simplified plate model is extended to the large, thin plates on the structure. This was used as a representative example

to illustrate the modal overlap in these plates, and to compare to the *measured* modal overlap in the structure.

3.4.1 Static shaker

Electromagnetic inertial actuators were chosen to provide forces on the structure due to the relatively high mechanical excursion they provided. They also provided a means to provide control force without being intrusive to the base structure. A low-order model of the electromagnetic actuator by itself was desired to ascertain its limitations and baseline behavior. The initial analysis assumed an rigid and infinite mass case, and then the model was extended to include finite case mass and arbitrary external forcing function.

Modeling the electromagnetic shaker was done from a basic electromagnetic and physical model introduced in [105]. This model incorporated a simple $L R$ and v_{emf} driving a mass-spring system, and can be seen in Figure 3.14. The equations governing this system are summarized in (3.1)-(3.5). Applying the Laplace transform to the differential equations and presenting them in typical multi-variable state-space form, (3.6) summarized the system dynamics.




Figure 3.14 Fixed electromagnetic shaker schematic.

$$f(t) = m \frac{d^2}{dt^2} x(t) = c \frac{d}{dt} x(t) + kx(t) \quad \text{Mechanical dynamics} \quad (3.1)$$

$$v_{in} = L \frac{d}{dt} i(t) + Ri(t) + v_{emf} \quad \text{Electrical dynamics} \quad (3.2)$$

$$f(t) = \Gamma i(t) \quad \text{Generated force} \quad (3.3)$$

$$v_{emf} = \Gamma \frac{d}{dt} x(t) \quad \text{Back EMF linkage} \quad (3.4)$$

$$\Gamma \triangleq Bl \quad \text{Flux linkage} \quad (3.5)$$

$$\begin{bmatrix} \dot{i}_o \\ \dot{x} \\ \ddot{x} \end{bmatrix} = \begin{bmatrix} \frac{-R}{L} & 0 & \frac{-\Gamma}{L} \\ 0 & 0 & 1 \\ \frac{\Gamma}{m} & \frac{-k}{m} & \frac{-c}{m} \end{bmatrix} \begin{bmatrix} i_o \\ x \\ \dot{x} \end{bmatrix} + \begin{bmatrix} \frac{1}{L} \\ 0 \\ 0 \end{bmatrix} [v_{in}] \quad (3.6)$$

A custom-built permanent-magnet shaker designed and fabricated in the Vibrations and Acoustics lab at Virginia Tech was used as the testbed for verifying the analytical model. A photo of one of the shakers is shown in Fig. 3.15. Although not identical to the commercial units used on the experimental structure, the characteristics and resonant behavior were similar. The fabricated shakers had the benefit of being easily and non-destructively disassembled, to measure the shaken mass/permanent magnet and spring coefficients. Those parameters that were directly measurable are shown on the left-side of Table 3.1. Other parameters that could not be easily measured were either read from the construction material datasheets (such as the magnetic field strength), or computed from the experimental measurement. This experimental test was performed using an accelerometer on the shaken mass of the shaker of Figure 3.15, with the case of the shaker clamped to a large mass. A current shunt was used to measure the armature current of the shaker, and all of the signals were spectrally decomposed relative to the frequency content of a white-noise voltage excitation. The resulting transfer functions are shown in Figure 3.16.



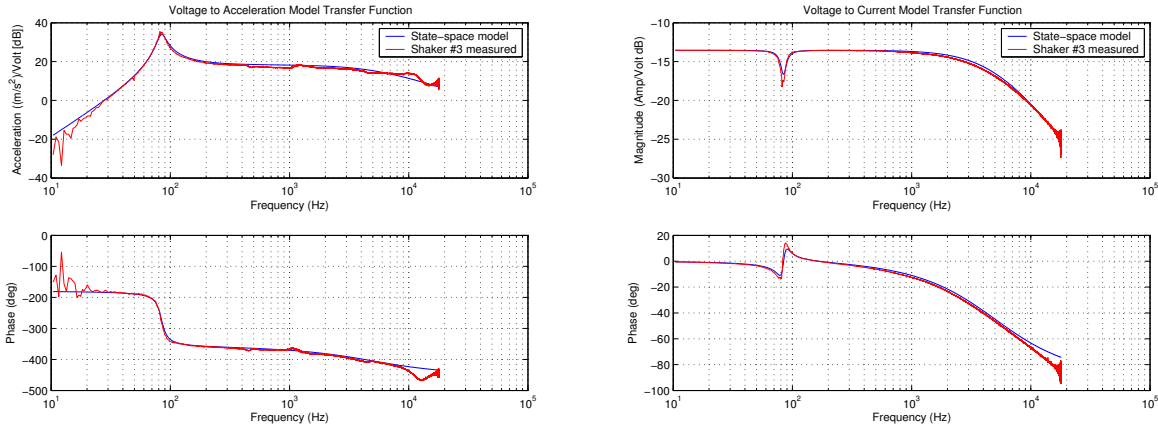
Figure 3.15 Actual electromagnet shaker characterized.

The response of the shaker was as expected, with a mechanical resonance around 100 Hz. Figure 3.16(a) shows the characteristic second-order response, with relatively flat voltage to force behavior above resonance. Electrically, Fig. 3.16(b) shows the single pole roll-off at

Table 3.1 Shaker coefficients.

Direct Measured	Value	Unit	Indirect Measured	Value	Unit
m_m	0.082	kg	c	5	$\frac{kN}{m/s}$
m_c	0.075	kg	B	0.5	T
k	23	$\frac{kN}{m}$	l	10	m
R	4.75	Ω	L	150	μH

approximately 3 kHz. Above this, the winding inductance became the dominant electrical impedance and limited the current, which limited the force response.



(a) Mechanical

(b) Electrical

Figure 3.16 Electromagnetic shaker experimental results and model fit.

3.4.2 Shaker with dynamic load

The experimental results on the fixed shaker of Section 3.4.1 both verified the validity of the model, and provided a means to measure the indirect coefficients of the shaker. With these characteristics known, an extension to the model was included. Since the ultimate goal was to provide force actuation to a moving structure, the load’s dynamics affect the shaker’s performance and must be included in the model. Figure 3.17 shows the modified schematic representation of the shaker that allows for an additional force “input” from an external source. The dynamic equations in (3.1)-(3.5) were modified to include that additional mass, and are summarized in the state-space equation of (3.7). This system allows for connection of the shaker to any arbitrary dynamic load.




Figure 3.17 Electromagnetic shaker schematic.

$$\begin{bmatrix} \dot{x}_m \\ \ddot{x}_m \\ \dot{x}_c \\ \ddot{x}_c \\ \dot{i}_o \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ \frac{-k}{m_m} & \frac{-c}{m_m} & \frac{k}{m_m} & \frac{c}{m_m} & \frac{-\Gamma}{m_m} \\ 0 & 0 & 0 & 1 & 0 \\ \frac{k}{m_c} & \frac{c}{m_c} & \frac{-k}{m_c} & \frac{-c}{m_c} & \frac{\Gamma}{m_c} \\ 0 & \frac{-\Gamma}{L} & 0 & \frac{\Gamma}{L} & \frac{-R}{L} \end{bmatrix} \begin{bmatrix} x_m \\ \dot{x}_m \\ x_c \\ \dot{x}_c \\ i_o \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{-1}{m_c} & 0 \\ 0 & \frac{1}{L} \end{bmatrix} \begin{bmatrix} f \\ v_{in} \end{bmatrix} \quad (3.7)$$

The characteristics of the external structure were now isolated to its mechanical impedance, defined by the external force f , and case velocity, \dot{x}_c . This point impedance for the structure was what was used in the high-frequency synthesis in Sec. 4.2.

3.4.3 Plate bending wave equation

Two different sections of the structure required an analysis of a simple plate's bending response. The variable "fuzzy" plates on the sides of the structure were important for mid to high frequency disturbance rejection, so obtaining a representative model of the dynamics was useful. Also, the baseplate where the control actuator was mounted imposed high-frequency modeling limitations for the actuator loop, so modeling its representative dynamics were also important.

For illustrative purposes, the 1-dimensional bending wave equation was considered initially. The bending wave in thin uniform homogeneous beam was considered from [106], pp. 279-286. Using the shear force on an infinitesimal element of a cantilever beam, the Euler-Bernoulli wave equation of (3.8) results.

$$0 = m \frac{\partial^2 w}{\partial t^2} + EI \frac{\partial^4 w}{\partial w^4} \quad (3.8)$$

This equation assumes materials constants of Young's modulus E , per-length mass of m , geometric relationships of second area of moment I , and transverse direction of motion, w . Substituting a complex exponential solution results in the 4th-order equation of (3.9).

$$0 = EIk^4 - \omega^2 m \quad (3.9)$$

The roots of (3.9) result in bending wavenumber (3.10) and corresponding wave-speed (3.11).

$$k_b = \sqrt{\omega} \left(\frac{m}{EI} \right)^{\frac{1}{4}} \quad (3.10)$$

$$c_b = \sqrt{\omega} \left(\frac{EI}{m} \right)^{\frac{1}{4}} \quad (3.11)$$

The analysis was continued in [106] to extend to a 2-dimensional response of a thin, homogeneous, uniform plate. The substitution of bending stiffness *per unit width* for bending stiffness, and mass per unit *area* for mass per unit *length* are shown in (3.12) and (3.13), respectively for a uniform plate of thickness, h .

$$EI \rightarrow D = \frac{Eh^3}{12(1 - \nu^2)} \quad (3.12)$$

$$m \rightarrow m = \rho h \quad (3.13)$$

The corresponding wavenumber and bending wave-speed are then shown in (3.14) and (3.15).

$$k_b = \sqrt{\omega} \left(\frac{m}{D} \right)^{\frac{1}{4}} \quad (3.14)$$

$$c_b = \sqrt{\omega} \left(\frac{D}{m} \right)^{\frac{1}{4}} \quad (3.15)$$

Since $\lambda = \frac{2\pi}{k}$, (3.16) shows the resulting bending wavelength.

$$\lambda = \frac{2\pi}{\sqrt{\omega}} \left(\frac{Eh^2}{12(1 - \nu^2)\rho} \right)^{\frac{1}{4}} \quad (3.16)$$

Applying the materials constants of Table 3.2 on a 0.5" aluminum plate results in the wavelength of (3.17) and corresponding frequency characteristic of Fig. 3.18.

$$\lambda \approx \frac{28.0}{\sqrt{\omega}} \quad (3.17)$$

With the bending wavelength of the plate in Fig. 3.18, now its application to the actuator site can be considered. Figure 3.19 shows the footprint of the electromagnetic actuator that

Description	Variable	Value	Units
Young's Modulus	E	71×10^9	Nm^{-1}
Density	ρ	2.7×10^3	kgm^{-3}
Poisson's ratio	ν	0.33	—

Table 3.2 Materials properties for Aluminum

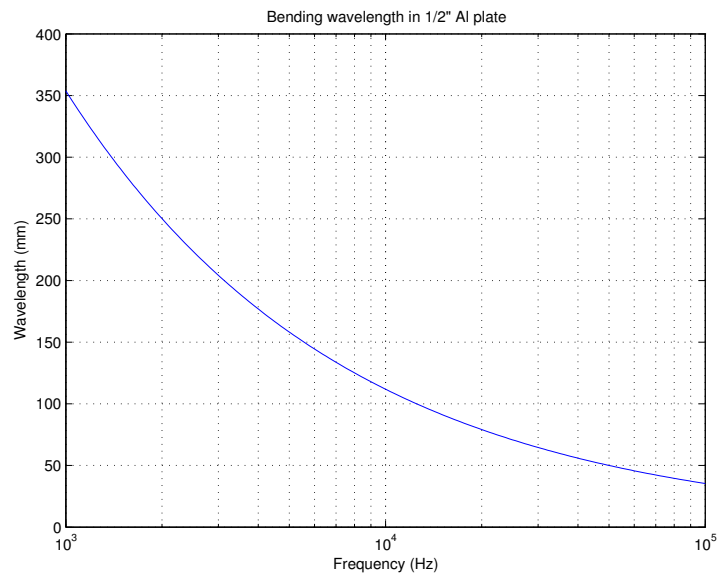


Figure 3.18 Bending wavelength of 12.7mm (1/2") aluminum plate

was used on the baseplate. The force was applied to the plate (normal to the page) in Fig. 3.19, but only around the *perimeter* of the footprint. Thus, at sufficiently high frequencies, there was a bending mode that could fit inside the actuator’s footprint. The error accelerometer was placed in the center of the footprint (circle), so this bending mode would have an effect.

The first mode of the plate would be when the physical size of the actuator becomes an integer multiple of the $\frac{\lambda}{2}$. Figure 3.18 shows the plate wavelength approaching 0.1m at approximately 10 kHz. Given the largest dimension of the actuator is 37.5mm, the first expected mode and its associated phase reversal would be between 10-20kHz. This frequency is well above the control frequency of interest, but its phase contribution at lower frequencies reduced the feedback stability margin.




Figure 3.19 Electromagnetic actuator footprint

Rather than look at just the wavelength on the scale of the *actuator* scale, the bending plate model was then extended in two dimensions to provide a representative frequency response of the bottom plate. The frequency response at a point approximately 0.1m from the edge of a 0.4m x 0.4m simple-supported plate was then calculated. It can be seen in Fig. 3.20. In the experimental structure, the plate was not simply-supported, nor was the actuator a “point force,” but the *general* character of the mid-frequency response in Fig. 3.20 looks similar to Fig. 3.2. Of course the full structure of Fig. 3.2 is much more complicated than the single plate. All of the control actuators (Figs. 3.2(a)-3.2(c)) exhibit resonances in the 400-600Hz region, but the disturbance path in Fig. 3.2(d) shows little response in that same band. These modes are attributed to the lower plate’s response where a few, lightly-damped modes dominate between 200-2000Hz. Above that, the response is much less defined by individual modes.

The same computation was done for the thin “fuzzy” plates. Because they were only 1mm thick as opposed to 12.7mm, their first modes were much lower in frequency. Figure 3.21 shows a sample result from a 1.1m x 0.6m plate 1mm thick excited approximately 0.1m from a corner. This plot exhibits extreme modal density by only a few hundred Hz. Although small in

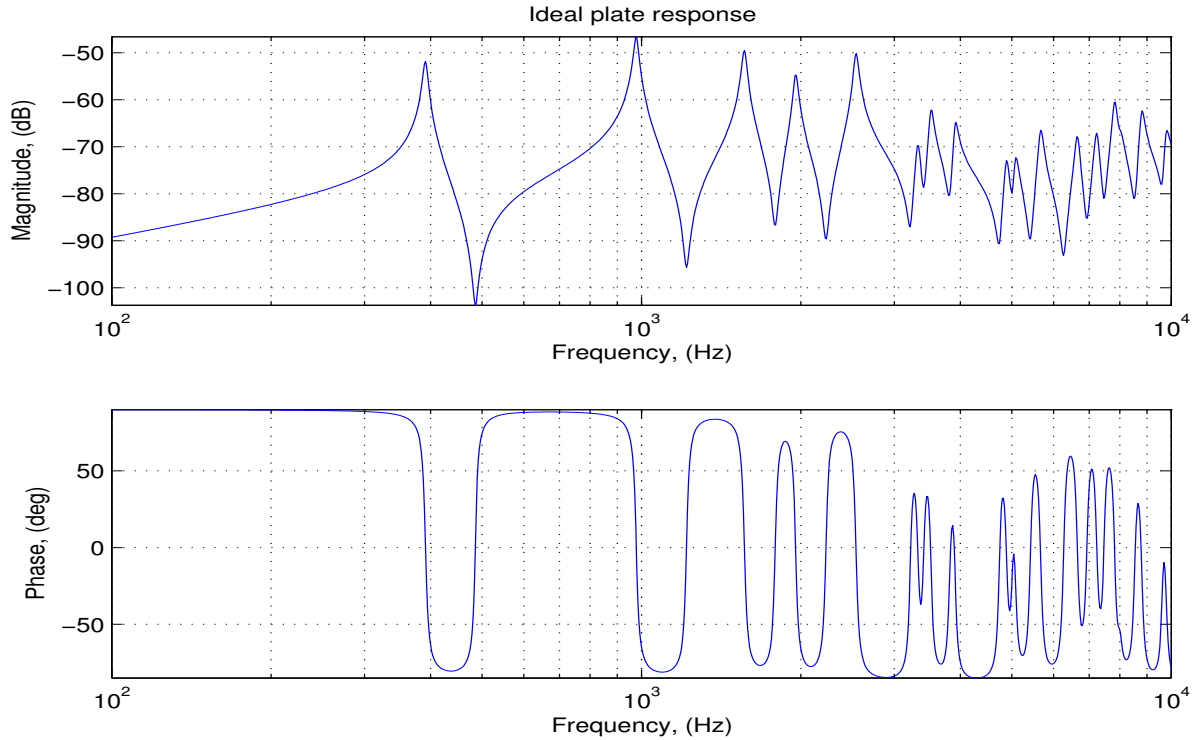


Figure 3.20 Simple baseplate analytical frequency response

magnitude compared to the frame of the structure, the four large plates contribute dynamics like this to the disturbance path.

3.4.4 Modal overlap

With a simplified analytical model of a bending plate in hand, the concept of modal overlap from the acoustics field was used as a method to describe the complexity of the model. As defined in [107], modal overlap is a measure of the extent to which individual resonances within a system are influenced by resonances at adjacent frequencies. The modal overlap definition compares the number of modes within a bandwidth around a single mode to a specific bandwidth. Normally, that bandwidth is the half-power bandwidth of a single mode and is determined by the damping of the structure.

As an example, consider the same, thin-plate model of Fig. 3.21. This model consisted of a 1% damping, so the half-power bandwidth around each mode of the structure is defined as 2ζ , or 2% of the natural frequency of the mode. Plotting the half-power bandwidth against the modal spacing generated by the analytical model results in Fig. 3.22. What this plot

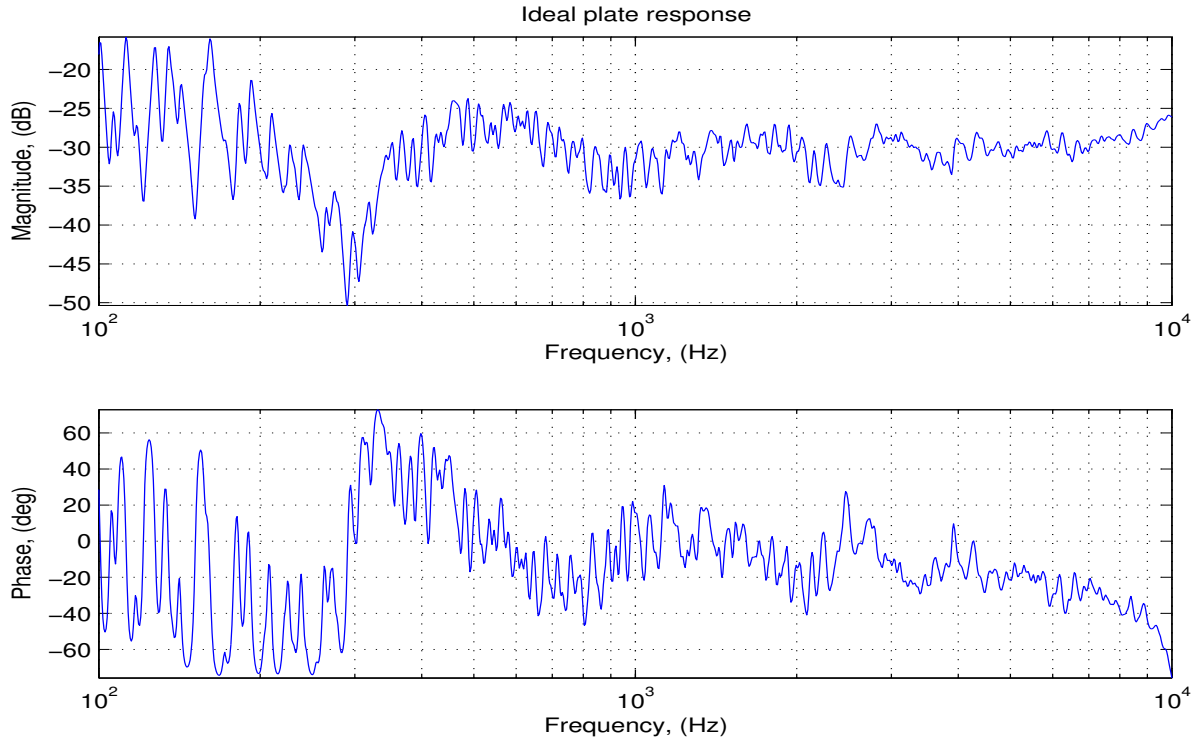


Figure 3.21 Simple thin plate analytical frequency response

shows is where the half-power bandwidth becomes more significant than the spacing between the modes. At 300Hz, the “average” modal overlap would be close to unity. By 2.5kHz, the “average” modal overlap would be approaching 10.

The single-instance analytical model of a plate shows the modal overlap relatively well. For the experimentally-measured system, this plot is not so clean. The half-power bandwidth of the identified models’ eigenvalues was calculated by the product of its natural frequency and associated damping. Realistically, this represents an example that is not as modally dense as an overall characterization of the plate, because not all modes will be visible from a single point. As a means of illustrating a *minimum* level of system complexity however, it has significant merit.

This representation of modal overlap was extended to the identified model from measurements of experimental structure. The difference between the natural frequency of adjacent modes as well as the half-power bandwidth was computed and both plotted in Figure 3.23. For the half-power bandwidth computation, some low-pass filtering was applied to smooth the result

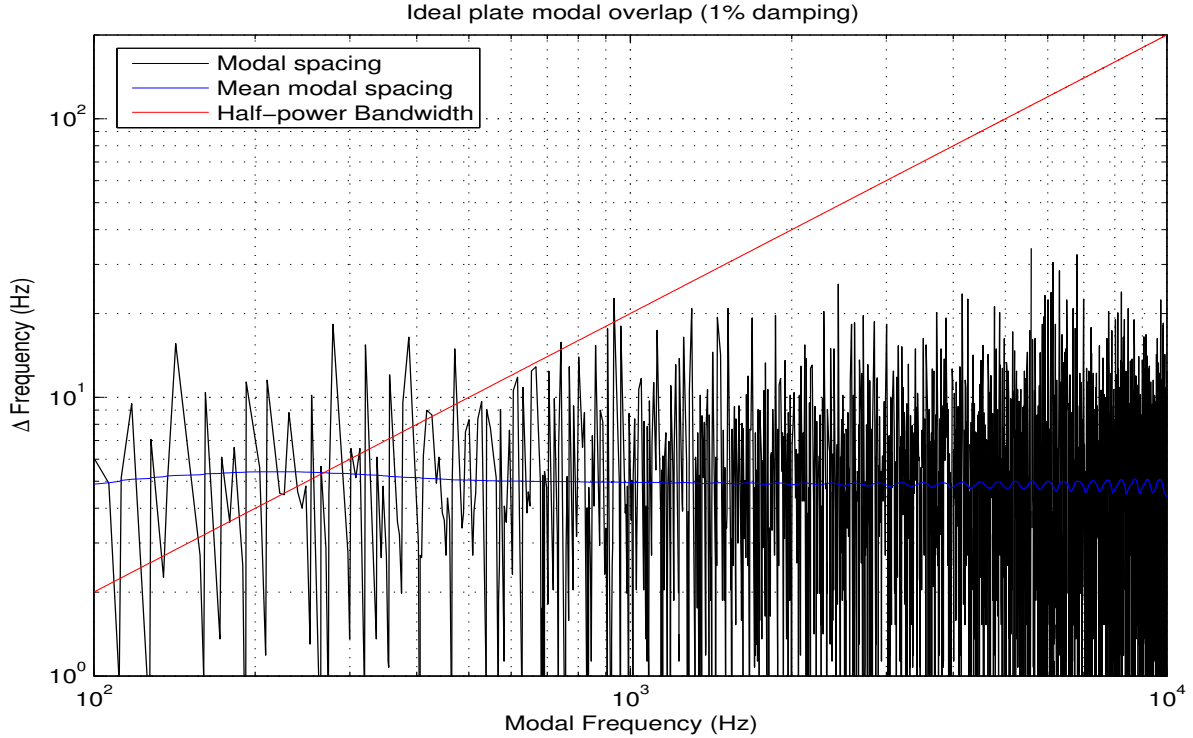


Figure 3.22 Simple plate analytical modal overlap

so that it was more easily comparable to the modal spacing. In each case, the heavily-filtered “mean” shows the general trend in the structural dynamics as a function of frequency.

What the measured structural modal overlap plot shows is that above 400Hz, the mean modal overlap increases above unity indefinitely. This represents a “band” characteristic to the structure. Another trend to observe from this plot occurs above 1-1.5kHz. Above this frequency, the mean modal overlap stays above 2. These two generalized bands provide additional justification for the frequency-separation control bandwidths discussed in Sec. 4.1.

3.5 Summary

The control preliminaries discussed in Sec. 3.1 reduced the number of control design variables to a tractable level. The initial characterization of the structure in Sec. 3.2, looked at its variability, and frequency-domain characteristics, and potential actuator locations.

Also discussed in this section were low-order models of the dominant characteristics of the actuators and the aluminum plates. The models are not intended for direct control applications,




Figure 3.23 Measured structural modal overlap

but rather as another means to characterize the structure as a whole, and to provide justification for the choice of control bandwidths used in Sec. 4.1. Finally, the details of the SOCIT toolbox used to generate the analytical model of the system from experimental data was discussed in Sec. 3.3. This model provided the analytical basis for the error analysis and estimated performance investigations in Chapter 6.

In particular, the focus on *implementability* of the control system was discussed. These provide the justification for the control topologies used for the control synthesis and implementation.

CHAPTER 4

WIDEBAND CONTROLLER SYNTHESIS

No single control technique provided the performance, robustness, and complexity trade-offs for wideband control. Chapter 3 discussed measurements of the test structure and some characterizations of it. Using these as a guide, the control synthesis was tailored to the specific requirements of the test structure. For wideband control, this ultimately required leveraging the benefits of different types of control. This chapter outlines the rationale and procedure for the synthesis of the combination of the different control types used.

Section 4.1 outlines the specific ranges of frequencies and justifications for the control type chosen for each band. This takes into account the measured structural response from Sec. 3.2-3.3, the control problem assumptions of Sec. 3.1, and some sources in the literature survey.

The high-frequency controller synthesis is explained in Sec. 4.2. This included a simplified analytical model of the actuator collocated on the structure. The simplified model's pole/zero characteristics provided insight into the types of feedback that were viable for the structure.

Section 4.3 describes the H_2 design process used for the design of the low-frequency controller. The various model manipulations and design parameters and their tradeoffs are discussed. Also included is the expected performance via numerical simulation. Some of the practical aspects of the implementation are also discussed.

Finally, Sec. 4.4 discusses the synthesis of the feed-forward controller. The high-level configuration of the FxLMS controller is explained. The adaptive algorithm that was implemented in the control computer is also described.

4.1 Proposed Control Strategy

At the center of the concept of splitting the control design into three bands is that the structural dynamics split into low, mid and high frequency regimes [99]. Each regime has properties which are distinct from the other two, and the modeling approach is also different. A convenient metric for drawing boundaries between the regions is the modal overlap, (M) introduced in Sec. 3.4.4. The modal overlap is defined as the number of modes within the bandwidth of a single mode. It is frequency dependent and can be approximated by using macroscopic properties of the structure: wave speed, material damping and surface area, but in this case, the *measured* response of Fig. 3.23 was used as the basis for the structural modal overlap.

The natures of the three regimes are:

- **Low Frequency:** $M \leq 1 \Rightarrow f < 400\text{Hz}$

The dynamics are dominated by well separated resonant peaks. Within the bandwidth of each resonance the structural response can be well approximated by only considering the contribution from the resonant mode. Control in this bandwidth depends on accurate knowledge of the resonant response. A poor model can cause large errors between the actual and modeled phase response [70]. This makes the control design a compromise between performance and stability, but the well-developed state-space optimal and robust control techniques can be utilized [38]. In this regime, the variations in the plant are assumed to be very small because the variable plates will have little effect on such low frequencies.

- **Mid Frequency:** $1 < M < 2 \Rightarrow 400\text{Hz} < f < 1000\text{Hz}$

Accurate modeling of the mid frequency range is a continuing problem for structural dynamicists. The inherent problem is that although the response contains significant peaks due to modal behavior there is also uncertainty in the knowledge of which modes combine. The response at any frequency in this range is due to the interaction of a number of modes. To control response in this range an adaptive controller must be used as it is practically impossible to know the structural dynamics sufficiently well *a priori*. Also, the characteristics in this frequency range are fairly susceptible to operating conditions such as temperature, which can change as the controller runs. Significant performance

benefits have been shown in these uncertain, high-order systems using the feed-forward techniques of [83] [108] where a reference signal was available.

- **High Frequency:** $M \geq 2 \Rightarrow f > 1000\text{Hz}$

The high frequency range is notable in that the complex dynamics of the mid frequency range have now converged to a relatively constant level of response. The variance about this level may be large, and is due to the large number of modes contributing at any one frequency. Very little structure to the modes is identifiable, and although the response may resemble the individually-identifiable modes of the low-frequency regime, this is usually for a very small bandwidth. Statistical Energy Analysis allows a prediction of the mean level. Work such as [19], has allowed the variance about this mean level to be predicted as a function of loading and structural properties. These techniques provide an estimate to the energy transferred in this bandwidth, but are not detailed enough to use for a controller design. The variance of the individual modes that make up the net response is large, so pole-zero cancellation techniques cannot be used. One of the few tractable *active* control techniques is collocation, where actuators and sensors are physically mounted together.

The frequencies at which these transitions occur on the structure can be seen in Fig. 4.1. This shows the target bandwidths of the three separate frequency regimes defined by the modal overlap regions above. As can be seen, the objective measurement of modal overlap agrees with the intuitively “natural” groupings of the structural response.

The technique of collocation has been focused on for a long time [22] where physically collocated, and direct rate-feedback controllers have tantalizingly robust stability properties. This requires direct velocity measurements and the ability to provide a force to the structure at the same location as the measurement, and directly proportional to it. The velocity measurement is not difficult, but providing the force requires actuators that are often referred to in the literature [109] as “sky-hook” dampers. These are actuators that can provide force relative to an inertial “ground,” can apply a single, point force, and have no inherent dynamics. Unfortunately, such “clean-body” actuators are largely mathematical tools that can only be approximated when implemented [26] [25]. *Real* actuators do not have a “grounded” reference from which to apply force, nor can they be exactly collocated or have zero internal dynamics.

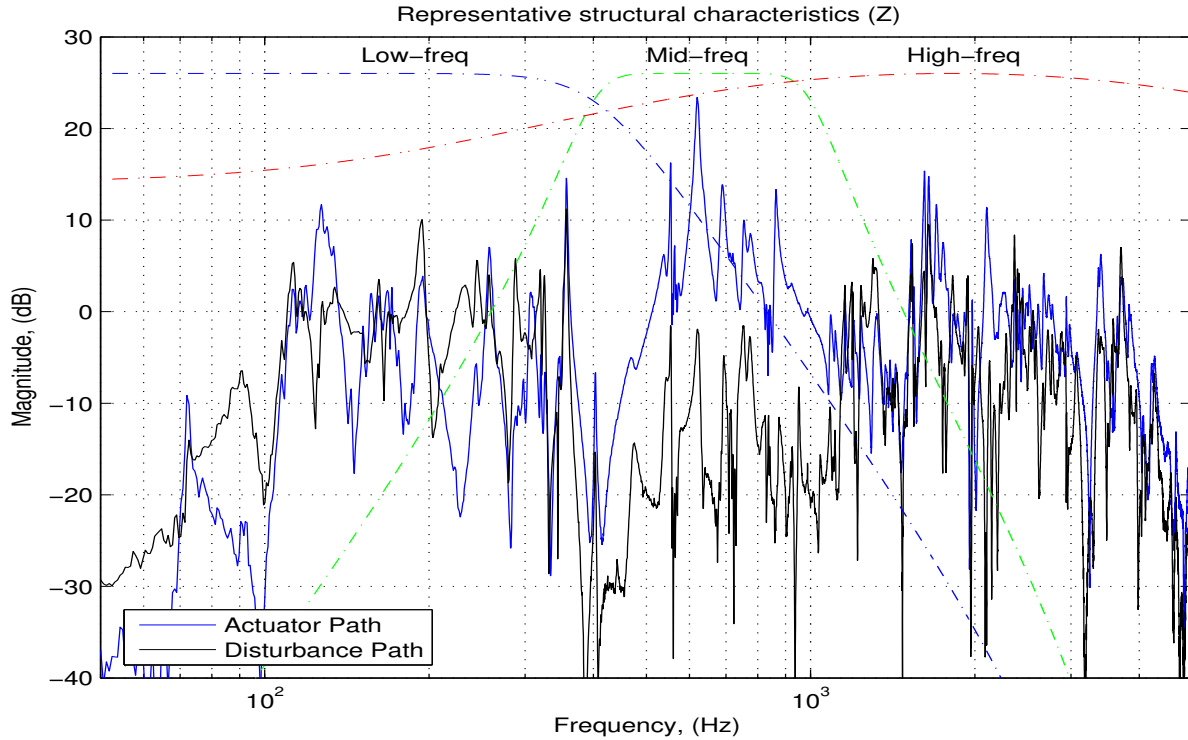


Figure 4.1 Multi-bandwidth control regimes

With these structural regimes in mind, the characteristics of the different control schemes discussed in Chap. 2 must be carefully considered. The biggest control problem with the structure was its uncertainty and complexity. The structural variation shown in the general characterization of Fig. 3.2, combined with a medium-high modal density above 400Hz implied that any sort of high-order model-based control would be impractical above this frequency. Unfortunately, that means that a large amount of control literature typically targeted to “large flexible space structure,” was simply too model-dependent as well as being too computationally expensive to implement above this bandwidth. The model-based controllers and a large amount of optimal/robust control theory cannot be utilized where the pole/zero cancellations [29] they require are too inaccurate [50]. This limits them to modally sparse and relatively well-known structures. The low-frequency regime of the test structure has this characteristic.

For the high-frequency portion of the structure, only the robustness of passive or actively-passive (dissipative) control can be relied on. The model is so unpredictable and potentially variable, that the only viable control is to exploit duality in the actuators and sensors. Collocated actuators and sensors [26], are insensitive to modeling errors [72] such as those at

high-frequencies on the structure. They have very attractive mathematical properties of stability, but in realizable implementations, errors in the collocation assumption present limits to their usefulness. Two very real limitations to the collocation assumption are nontrivial actuator and sensor dynamics, and physical limitation to collocated placement. The HAC/LAC strategy also utilizes this notion of collocation, but generally does so to aid in the robustness properties of high-order model-based feedback. In the high frequencies of interest here, these high-order models have already been dismissed as infeasible.

The high frequency regime allows only the most simplified of modeling assumptions to be made. Robustness must come from collocation. The mid-frequency control presents the largest challenge from a controls point of view. It lies between the relatively accurate model in the low-frequency regime, and the highly inaccurate high-frequency regime. The moderate modal overlap provides significant resonant peaks, while at the same time presenting large uncertainty in the plant model. High-order model-based static controllers cannot provide robustly stable performance in such uncertainty [47], while low-order feedback cannot provide the selectivity required.

The combination of already collocated actuators and sensors required for the high-frequency control, combined with the control problem definition of Sec. 3.1's assumption that a known reference signal was available allows the consideration of feed-forward, adaptive control. This type of control, in the form of the LMS algorithm enjoys much popularity in the active vibration and noise control field [83] [108] where such a reference signal is available. It provides a good tradeoff between computational complexity [82] and performance over other adaptive control techniques [81]. The adaptive implementation allows the controller to follow the variation in a specific structure, or variations on a representative structure.

The main drawback to the feed-forward LMS control is that while it has reduced computational complexity compared to other recursive schemes, it still requires significant resources. The actual control for LMS relies on a *control filter*, which is a finite-impulse-response (FIR) discrete-time filter. This has the significant benefit of guaranteeing stability of the adaptive filter. The disadvantage is that since it is *not* recursive, it can only approximate very lightly-damped resonances since it has a truncated impulse response. The long "ring-time" of lightly-damped modes requires large numbers of filter coefficients, which increases computational requirements.

To reduce the computational requirement, the design choice was made to limit the feed-forward control bandwidth to only the mid-frequency region. This was done primarily to allow the limited number of degrees of freedom in the adaptive filter coefficients to target frequencies out-of-range of the low-frequency controller. A lightly-damped mode at low frequency would require an even *larger* number of filter coefficients than an *equivalently-damped* mode at a higher frequency since the impulse response would require more time to decay for a low-frequency mode.

With all of these tradeoffs in mind, an overview of the proposed composite controller is shown in Fig. 4.2. The various components are:

- G : The “primary” error path. This represents the propagation path from the disturbance on the top of the structure through to the error sensor collocated with the control accelerometer.
- C : The “secondary” error path. This represents the propagation path from the (voltage) input to the control actuator to the collocated error accelerometer.
- \hat{C}_a : The model estimate of the secondary error path. This is required for the variant of the LMS algorithm used. It allows the LMS algorithm to adapt to the primary path, G , taking into account the propagation through the actuator to error path. This estimate must be accurate to within $\frac{\pi}{2}$ in phase. Even with a lightly-damped model, the collocated secondary error path is never more than $\frac{\pi}{2}$ away from an *average* phase between the two extremes. The estimated secondary error path must also be augmented with the contributions from both the high-frequency feedback and low-frequency feedback controllers. It is nominally the same as an estimate of the $\hat{C} \approx C$ but with these minor additions.
- H_A : The high-frequency continuous-time (analog) controller. This takes the error sensor reading directly and feeds it back to the input of the control actuator.
- H_{FB} : The low-frequency, high-order feedback controller. This path actually contains all of the necessary components to implement the digital control at a low-frequency sampling rate. It requires an accurate model of the secondary error path with the analog loop closed.

- H_{FF} : The mid-frequency, feed-forward controller. This path is an FIR filter that is adapted by the LMS block to minimize the measured, weighted error.
- LMS: The least-mean-squares adaption algorithm. This block contains the recursive, adaptive, gradient-descent algorithm to tune the coefficients of H_{FF} . Any spectral weighting terms to adjust the error differently from normal LMS are also contained within this block

The high-frequency feedback consists of a low-order augmented acceleration feedback. This provides high-frequency modal peak reduction for collocated actuators and sensors without requiring explicit pole-zero cancellations. The bandwidth range for the structure of interest for the high-frequency control is approximately in the 800-5000Hz. The design of this controller is illustrated in Sec. 4.3, with optimization detailed in Chap. 5. To the extent it overlaps the LMS frequency range, it reduces the length of the required impulse response.

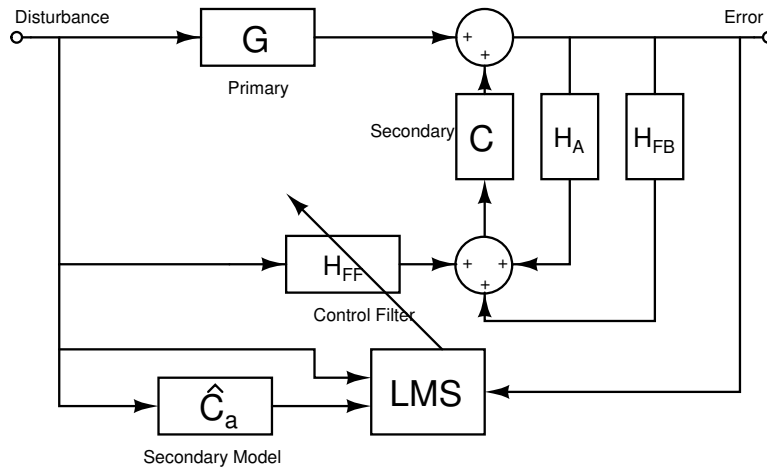


Figure 4.2 Multi-bandwidth control scheme

The low-frequency feedback is implemented as an H_2 controller based on an experimentally-identified model of the plant. This requires an accurate model of the modes of interest, but can provide significant performance with relatively low design complexity. Because it is a feedback controller, it also provides disturbance rejection to unknown disturbances. The design of this portion of the controller is discussed in Sec. 4.3, with the intended frequency range between 100-400Hz.

The mid-frequency control design attempts to bridge the frequency region between the well-known low-frequency modes, and the unknown high-frequency response. It uses the filtered-x variation on the LMS gradient descent feed-forward algorithm to attempt to minimize the output error signal and is further described in Sec. 4.4. The operational frequency range for the mid-frequency control on the structure of interest is approximately 400-1000Hz.

The three control loops were closed sequentially. This resulted in a series of outer loops to be closed around inner loops. It allowed for robustness to out-of-band spillover from each previous controller loop closed. Any changes to the overall system presented by the closure of one loop would be present in the identified model for the subsequent loops. Practically, this removes the possibility of designing an *unstable* inner loop that is stabilized by an outer loop. Such a design could potentially provide some increase in performance, but realistically it would be difficult to ensure stability. Instead, the more conservative design choice was made to ensure each sequential loop closure resulted in a stable aggregate system.

4.2 High-frequency Synthesis

The high-frequency synthesis was simultaneously the simplest and most difficult part of the design. It required the utilization of the information from the modal overlap of the structure of Sec. 3.4.4 and the low-order approximate modeling of Sec. 3.4. The difficulty comes in having to *simplify* the controller complexity as the plant itself becomes progressively more complicated. The mean structural response at these frequencies was used for the baseline design under the knowledge that any specific instance could be significantly different. Instead, the stability attributes of having a collocated actuator and error sensor were leveraged to allow for a reduction in the peak magnitude response, wherever the peak should occur in frequency. This technique was limited only by very real implementation limits.

4.2.1 Collocated structural response

The high-frequency control design was done with the performance of the electromagnetic actuator as a reference. The mechanical resonance provided another mode in the structure, but was at a significantly lower frequency than those of interest for control in the high-frequency regime. The worst case design for such a structure would be an undamped system with infinite

gain at modal frequencies. Previous work [22] provided a simplified model of a reverberant structural plant. The generalized N-mode model can be seen in (4.1)-(4.2), where $\eta(t)$ represents displacements in generalized coordinates, Ω represents the natural frequencies, ζ the damping ratio, and $f(t)$ an arbitrary forcing function.

$$Bf(t) = \ddot{\eta}(t) + 2\zeta\Omega^{\frac{1}{2}}\dot{\eta}(t) + \Omega\eta(t) \quad (4.1)$$

$$y(t) = D_1C\eta(t) + D_2C\dot{\eta}(t) \quad (4.2)$$

The matrix B represents the actuator influence function. The output equation's coefficients are divided between the actuator influence function represented by C , and the displacement and velocity components represented by D_1 and D_2 , respectively. For the purposes of investigating the collocated properties of the reverberant structural plant, the truncated modal displacement model of the undamped plant of (4.3) was described in [25].

$$G_d(s) = \sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2} \quad (4.3)$$

The succinct displacement model of (4.3) was shown in [110] to consist of alternating poles and zeros on the imaginary axis. This presented the displacement frequency response characteristic represented by Fig. 4.3. The phase transitions between 0 and π are the worst-case when there is zero damping.

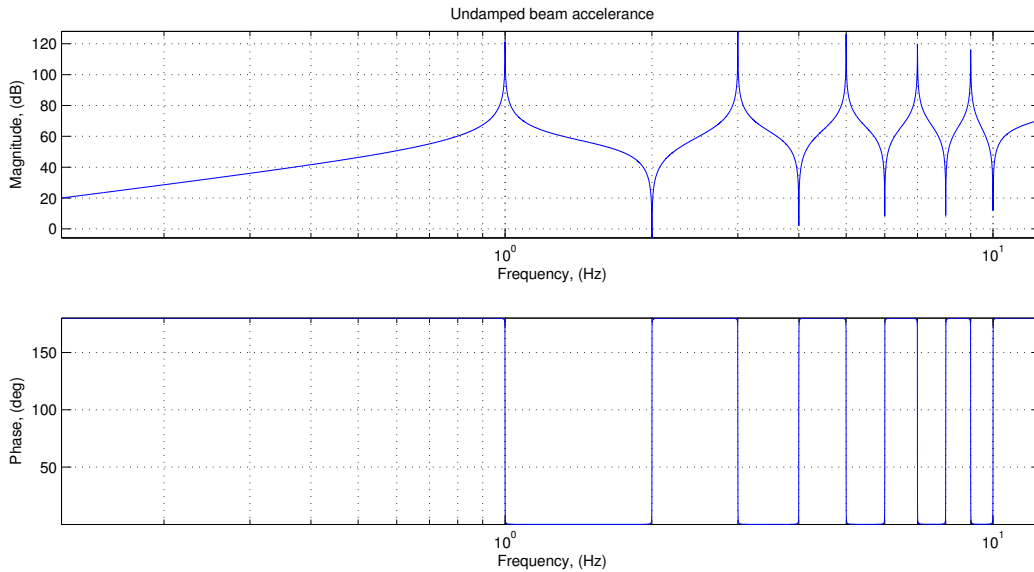


Figure 4.3 Sample undamped beam acceleration

The sensors used for implementation were not displacement sensors, but rather linear accelerometers. Equation (4.3) was rewritten in terms of acceleration by a double differentiation, which results in (4.4).

$$G_a(s) = \sum_{n=1}^N \frac{s^2 a_n}{s^2 + \omega_n^2} \quad (4.4)$$

While succinct as a summation of individual modal responses in (4.4), (4.5) shows the same equation rewritten in a product format.

$$G_a(s) = s^2 \left[\frac{\sum_{n=1}^N a_n \prod_{m=1, m \neq n}^N (s^2 + \omega_m^2)}{\prod_{m=1}^N (s^2 + \omega_m^2)} \right]. \quad (4.5)$$

This representation makes the poles and zeros directly visible, but was unfortunately very numerically sensitive to compute. For the purposes of the high-frequency design, however, only the interleaved *structure* of the poles and zeros was important [110]. Utilizing this fact, without loss of generality, (4.5) can be written as (4.6).

$$G_a(s) = s^2 \frac{\prod_{n=1}^{N-1} a'_n (s^2 + z_n^2)}{\prod_{m=1}^N (s^2 + \omega_m^2)}. \quad (4.6)$$

As an example, consider a representative system in Fig. 4.4. It contains 20 modes between 200-5000Hz of relatively sparse spacing and a 0.5% damping. The interleaved pole/zero structure can be easily seen, as well as the phase jumping between 0 and π . The finite damping ratio was added for clarity and because the real system has a similar amount of damping. Although discrete peaks are still visible, the phase excursions at the higher frequencies become less pronounced as the modal density (modes per damping-ratio bandwidth) increases. The relatively even distribution of modes, relatively low modal density, and constant damping ratio makes the model look dissimilar from the actual colocated plant. Despite this apparent dissimilarity however, the structure of the poles and zeros is still preserved as in the real system. Only the details (magnitude and damping of individual modes) are specific to the real system under study.

For the feedback possibilities of the plant of (4.5), the characteristics of the interleaved poles and zeros were exploited. Previous work [23] [111] [112] utilized direct velocity feedback. In particular [26] describes the stability characteristics of the more general colocated dissipative control. Applied to the plant of (4.3), velocity feedback can be written with the loop-gain of




Figure 4.4 Idealized lightly damped model

(4.7).

$$G_d(s)K_v(s) = s[k_v] \sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2} \quad (4.7)$$

The closed-loop system is represented exactly in (4.8) under the velocity feedback law of $K_v(s) = k_v s$.

$$\frac{G_d(s)}{1 + G_d(s)K_v(s)} = \frac{\sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2}}{1 + k_v s \sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2}} \quad (4.8)$$

$$= \frac{\prod_{n=1}^{N-1} (s^2 + z_n^2)}{\prod_{m=1}^N (s^2 + \omega_m^2) + k_v s \prod_{n=1}^{N-1} (s^2 + z_n^2)} \quad (4.9)$$

$$= \frac{\prod_{n=1}^{N-1} (s^2 + z_n^2)}{\prod_{m=1}^N (s^2 + d_m k_v s + \omega_m'^2)} \quad (4.10)$$

Using the pole/zero alternate form of the open-loop system of (4.6), the closed-loop poles can be seen in (4.9). The relative order of all feedback terms ($\prod_{n=1}^{N-1} (s^2 + z_n^2)$) in the denominator is increased by 1 with the $k_v s$ term. Combining the terms yields (4.10), with additional damping factors d_m and modified natural frequencies ω'_m .

In order to visualizing this and to provide a numerically substantive representation, a root-locus was plotted to see the exact damping and natural frequency movements. This illustrated what would be possible with classical loop-shaping techniques, including direct velocity feedback. The measurements were with accelerometers, so the simplest loop-shaping feedback would

be to feed the acceleration signal back. It does not provide the damping term of (4.10) that velocity feedback would since there is no additional order of s increase. Figure 4.5 shows a root locus of both of these types of feedback with the idealized plant. As expected, the acceleration feedback only changes the natural frequencies, but does not add the k_v -proportional damping factor. Velocity feedback provides the expected potential for increase in damping of the modes.

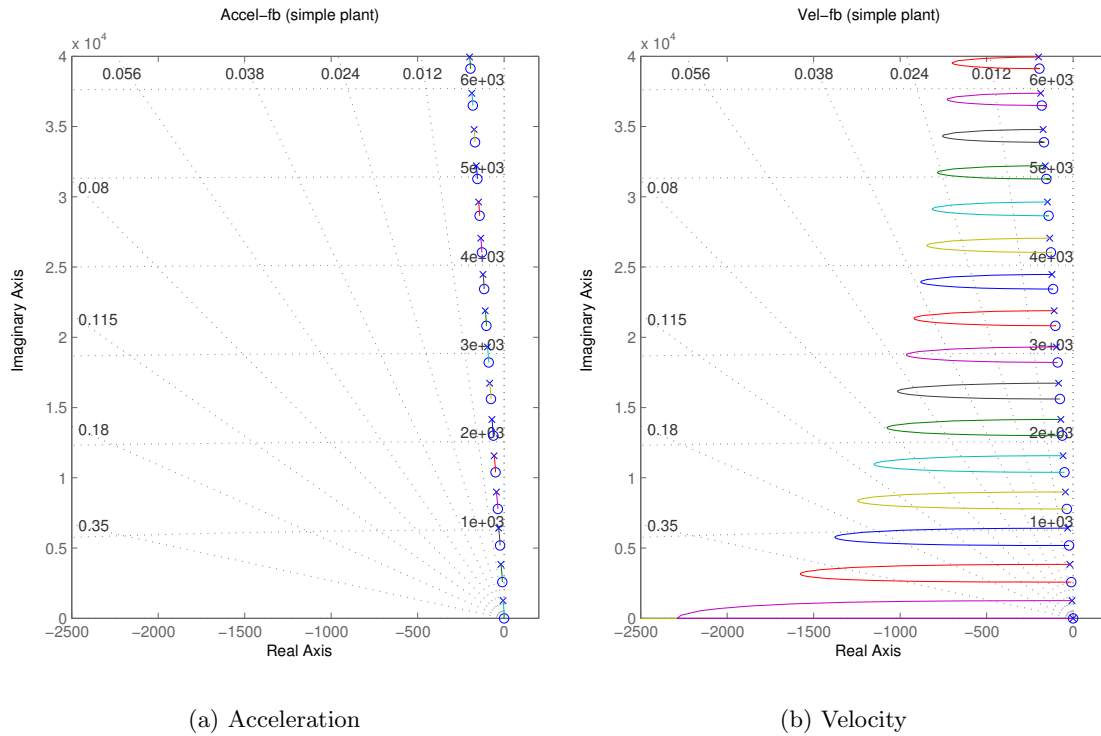


Figure 4.5 Root locus of idealized lightly damped plant (under acceleration or velocity feedback)

The representation of Fig. 4.5 was a useful way to compare the “acceleration” feedback to “velocity” feedback when a measured response is actually neither. Active vibration control literature often applies the notion of velocity feedback almost blindly without care to the the dynamics of the actuator and error sensor. Even when using an accelerometer as an error sensor, the measured signal is often run through an integration to produce a velocity signal before any signal conditioning design is applied. As such, phase roll off in either or both tends to distort the clean “active damping” notion that Fig. 4.5(b) would imply and cause additional stability problems. Indeed, by the time that a realistic actuator is included the root loci look

like those of Fig. 4.7. Under “direct velocity feedback” in this case (Fig. 4.7(b)) *no* feedback gain can provide stable performance. By the time the feedback gain is increased sufficiently to provide any control authority, out-of-band modal peaks combine with high-frequency phase rolloff to make the closed-loop system unstable.

Throughout the remainder of this Section 4.2, the paired acceleration/velocity root loci plot similar to Fig. 4.5 was used. It provided consistency throughout the evolution of the feedback controller synthesis. Rather than apply “velocity” feedback blindly, the resulting phase response of the controller/plant combination resulted in “pseudo-velocity” feedback in the frequency range of interest by measuring acceleration directly. Even though the design utilizes acceleration directly, retaining the “velocity” feedback plot helps illustrate how quickly it would become unstable if it were blindly applied. Comparing the two throughout makes it clearer that the more complexity that was added to the model and compensator, the closer “acceleration” feedback looks to velocity feedback, and the more impractical the use of a velocity signal becomes.

4.2.2 Realistic actuator

By adding an integrator, the simplified accelerance plant of (4.4) can be easily controlled with velocity feedback to provide the damping in the form of (4.10). Unfortunately, the inertial shakers used to provide the control actuation had their own dynamics. They provide control force to the structure based on applying electromagnetic force between the case and its internal mass. There is no external frame of reference to provide the reactive force, so the inertial actuators can only provide effective external force above the actuator’s internal resonance. Above this frequency, the actuator provides a flat accelerance response, and thus a force response, until the electrical $\frac{R}{L}$ pole begins to dominate. The collocated accelerance plant with the addition of the actuator dynamics can be seen in (4.11).

$$G'_a(s) = \frac{s^2}{(s^2 + 2\zeta_a\omega_a s + \omega_a^2)(s - p_L)} \sum_{n=1}^N \frac{s^2 a_n}{s^2 + \omega_n^2} \quad (4.11)$$

The actuator’s second-order resonance behavior is represented with ω_a and ζ_a , and the inductive pole is represented by p_L . A frequency response of the simplified plant of (4.4) and the more realistic plant of (4.11) can be seen in Fig. 4.6. The significant phase transition and ultimate

reversal through actuator resonance can be seen, as well as the gain and phase rolloff due to the electrical pole.




Figure 4.6 Idealized lightly damped model without and with actuator

The actuator resonant frequency was approximately 120Hz. The intended frequency range of interest for this controller was well above this, so the primary *performance* effect of the actuator resonance was to provide a phase reversal for use above resonance. Although important for *stability* constraints, it has little effect on the desired velocity feedback at high frequencies. This allowed the simplification of (4.12) to be valid.

$$G'_a(s) \approx s \left[\frac{s}{(s - p_L)} \right] \sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2} \quad (4.12)$$

The bracketed term is a combination of one of the residual differentiations due to acceleration measurement and the inductive pole of the actuator. As the frequency becomes higher, this looks more like a phase-less scalar.

Figure 4.7 shows the closed-loop pole locus of this system under acceleration or velocity feedback. With the addition of the realistic actuator to the model, the “acceleration” feedback term in Fig. 4.7(a) looks much more like the velocity-feedback term of Fig. 4.5(b) at the higher frequencies. The maximum possible damping ratio of the modes is increased by a factor of 10 at 28krad/s. The closed-loop poles in the “velocity” feedback of Fig. 4.7(b) now tend toward instability at frequencies above 25krad/s.

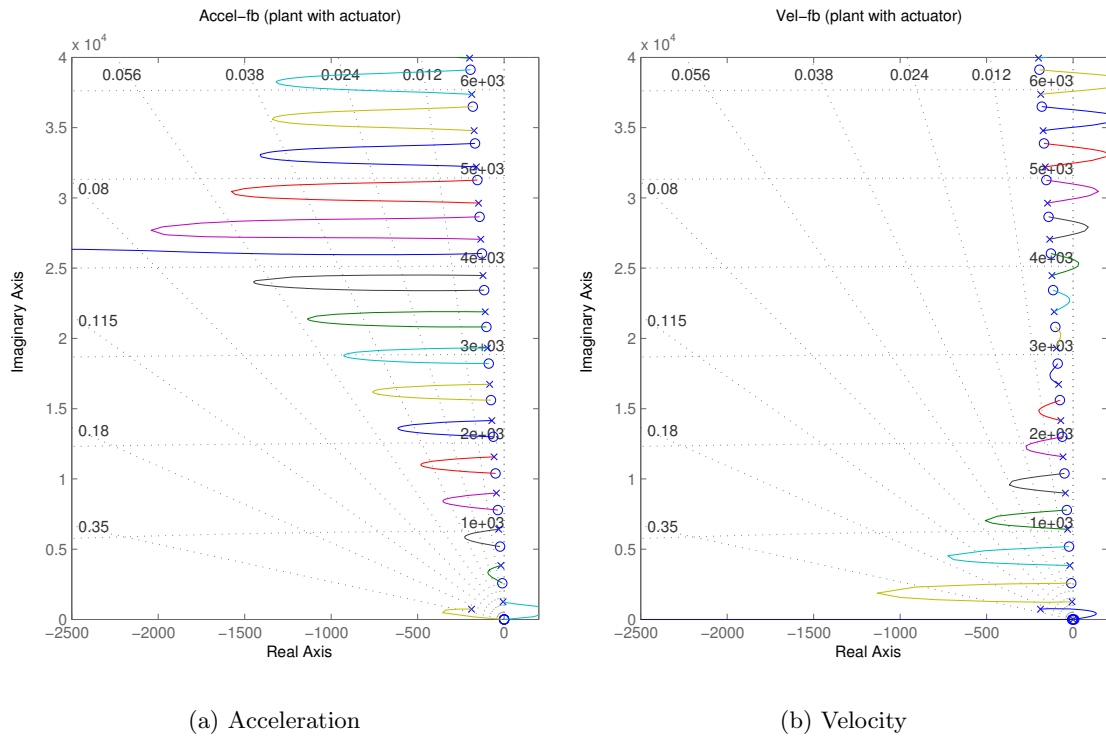


Figure 4.7 Root locus of lightly damped plant with actuator (under acceleration or velocity feedback)

To see this stability information more clearly, a Nichols plot representation of the loop gain was used. The information in a Nichols plot can be thought of as a parametric representation of a Bode-type magnitude/phase plot. The real and imaginary components of the complex-valued function are plotted directly in the complex plane as magnitude (in dB) and phase, while the frequency is swept through the entire range of interest. As a consequence, the *specific* frequencies at which a large gain or phase excursion occur are no longer obvious. What becomes *more* obvious however, are the stability characteristics of the loop. The point at which open-loop loop gain has a magnitude of -1 indicates the point at which the closed-loop system becomes unstable. This point is shown with a “+” symbol on the Nichols plot. The distance away from this point in magnitude indicate the gain margin, and the distance in phase the phase margin. Of course, any integer multiple of $(2N - 1)\pi$ in phase represents the same phase point as -1 , so they are also measures of stability.

It should be noted that for systems such as the one of interest here, a Nichols plot has significant advantages over the more traditional Nyquist plot. A Nyquist plot also represents the magnitude/phase information in a parametric way, but does so in a linear, not logarithmic way. As a consequence, for systems with high dynamic range and wide phase variations, a Nyquist plot ends up effectively showing only the largest of gain peaks. By plotting the magnitude in dB in a Nichols plot, the entire spectrum of dynamic range is easily seen. The stability is also more clearly visible, as both magnitude and phase can be read directly off the plot. Figure 4.8 shows the example of the two types of plots for the same pure velocity feedback system of Fig. 4.5(b). Obviously, stability is not a concern in this example, but the Nichols plot describes the high dynamic range of the system much more clearly.

With the Nichols plot as a good visualization tool, it became easier to see the low-frequency stability problem in Fig. 4.7. For performance, the root loci clearly shows the benefit of the “pseudo” velocity feedback. At the low frequencies however, the actuator’s resonance became a limit to the maximum damping that can be achieved. To more clearly illustrate this, consider the Nichols plot in Fig. 4.9. The dashed black line shows the same idealized modal plant of (4.4) with pure acceleration feedback. The dotted-dashed black line shows the same idealized plant under velocity feedback. Both of these are stable under any feedback gain, since they do not approach the -1 point. The red plot is the idealized plant with the non-ideal actuator of (4.11). This plant’s gain can only be increased to a maximum of what is shown on the plot




Figure 4.8 Nyquist vs. Nichols plot of pure velocity feedback on lightly damped plant

because of the stability constraint at the lowest frequencies. The final (blue) plot is that shown with the compensated acceleration feedback of Sec. 4.2.3 employed. What can be quickly seen is that the both the pure acceleration and compensated acceleration feedback have the most restrictive constraint at the first “loop.” This plot, together with the Bode representation of Fig. 4.6 show that this is the frequency of actuator resonance.

What Fig. 4.9 also shows however, is that the compensation allows for higher gain throughout more “loops” of the Nichols plot than the uncompensated acceleration feedback. Although the plot is shown with enough gain to be only marginally stable, the compensated plant also shows a much steeper slope around the -1 point. This indicates that an increase in gain margin (by shifting the plot downwards) will benefit the compensated plant with more *phase margin* per gain margin than it will for the uncompensated plant. Indeed, for the compensated plant, an increase in gain margin by only a few dB will allow for an infinite phase margin for the first plant resonance. The same type of analysis by inspection would be very difficult for an equivalent Nyquist plot, since angle measurements would be required.

4.2.3 Compensated feedback

The performance of the simple plant with actuator eventually approached that of pure velocity feedback as the frequency increased. In order to tailor the performance to the target range of 5kHz or less however, some additional compensation was used. The goal was to create




Figure 4.9 Nichols plot of simple plant, acceleration, and compensated acceleration feedback

a band-pass filter to provide as much feedback gain within the desired bandwidth as possible. At low frequencies, the limitation was the actuator resonance. At high frequencies, the plant eventually rolls off due to collocation errors [26] and other unmodeled phase delays. The form of the compensator chosen was that of a low-order bandpass filter in (4.13).

$$C_{HF}(s) = \frac{k(s - z_1)}{(s - p_1)(s - p_2)} \quad (4.13)$$

The actual location of the poles and zeros are optimized in Sec. 5.2.2 and depend largely on the modeling of the end constraints and the desired stability margin. For illustration purposes here, the optimized results from Sec. 5.2 of $p_1, p_2 \in -2\pi\{10^{2.9}, 10^{3.7}\} \approx -2\pi\{800, 5000\}$ and $z_1 \in \{10^{2.2}\} \approx -2\pi\{160\}$ were used.



The analysis of the closed-loop poles with the compensator is similar to that of (4.11). The frequencies above actuator resonance are of the ones of interest, so (4.14) can be approximated as (4.15).

$$C_{HF}(s)G'_a(s) = \frac{k(s - z_1)}{(s - p_1)(s - p_2)} \left[\frac{s^2}{(s^2 + 2\zeta_a\omega_a s + \omega_a^2)(s - p_L)} \sum_{n=1}^N \frac{s^2 a_n}{s^2 + \omega_n^2} \right] \quad (4.14)$$

$$\approx s \left[\frac{ks(s - z_1)}{(s - p_1)(s - p_2)(s - p_L)} \right] \sum_{n=1}^N \frac{a_n}{s^2 + \omega_n^2} \quad (4.15)$$

The bracketed term has relative order of 1 which ensures high-frequency rolloff. Also, the phase contribution from the high-frequency compensator pole (5kHz) and electrical inductive pole (3kHz) provide a cumulative phase contribution approximating $-\pi/2$ by 3kHz. This is in the middle of the desired band, so it maximizes the bandwidth for “pseudo”-velocity feedback. Figure 4.10 shows the closed-loop system with the same loop gains and compensator of Fig. 4.9. The increased damping ratio in the 1-5kHz range can be seen, which reduces the modal peaks. Also shown is the nominal aggregate loop-gain phase, including the compensator and inductive pole.

The root locus of the compensated plant can be seen in Fig. 4.11. Again, the “acceleration” feedback of Fig. 4.11(a) is no longer acceleration, but rather has the phase characteristics of velocity feedback. This provides a bandwidth from approximately 1-5kHz for which increased damping on the locus is possible. Also notice that the increased phase contribution of the compensated “velocity” feedback plot of Fig. 4.11(b) results in high-frequency instability at a frequency as low as 20krad/s as the locus tends towards the right-half plane. What this shows



69

Figure 4.10 Idealized lightly damped model closed loop

is that the uncompensated velocity feedback of Fig. 4.7(b) makes even the mild loop-shaping of the high-frequency compensator infeasible. There is not enough bandwidth to work within, since the additional integration to velocity removed $\frac{\pi}{2}$ phase margin on the high-frequency end. This shows how blindly applying “direct velocity feedback” as a control law often does not work at all, or has significant gain limits in practical systems.

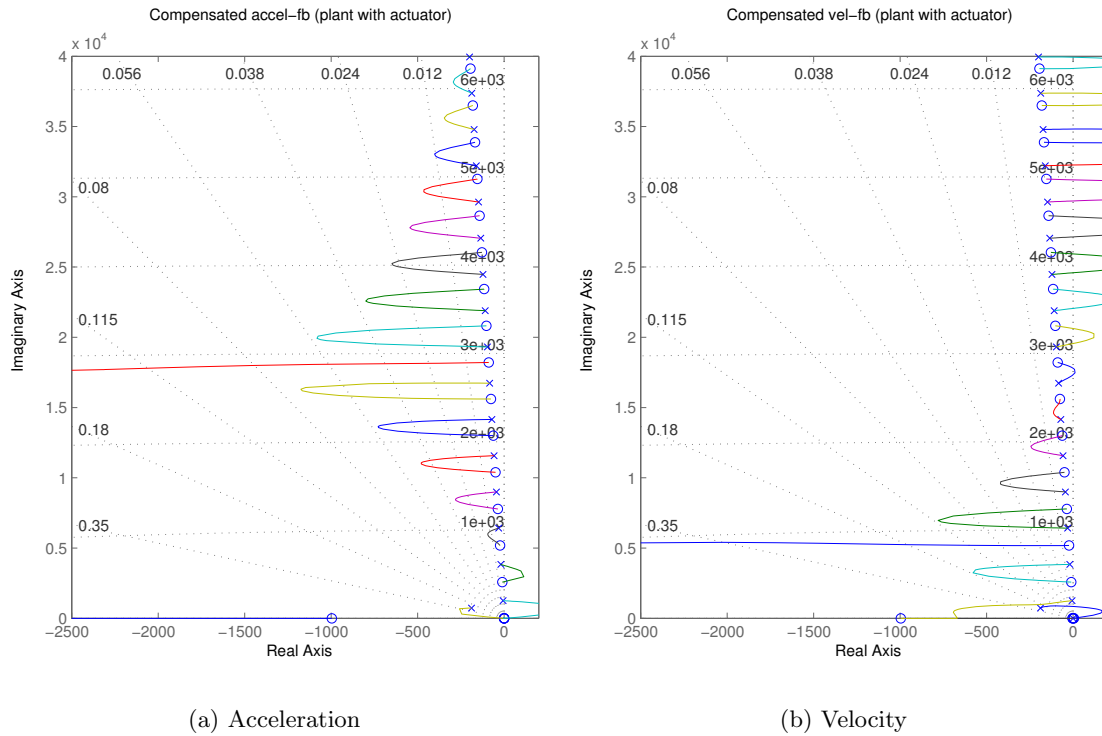


Figure 4.11 Root locus of lightly damped plant with compensated actuator (using acceleration or velocity input)

4.2.4 High-frequency summary

For this feedback example, the actual root locations can be seen in Fig. 4.12. The maximum damping was at approximately 15krad/s and was a factor of approximately 2.5 over the open-loop plant. The uncompensated acceleration feedback may look like it has higher damping above 2.5krad/s, but realistically this form of feedback would be unstable at very high frequencies. The high-frequency rolloff was required. At frequencies above 10kHz, the mismatching in collocation and propagation through the structure become important, yet are very difficult to

quantify. With only the $\frac{L}{R}$ pole of the actuator to roll off this gain in the uncompensated case, the increase of only a factor of 2-3 in frequency yields 6-9dB of gain rolloff. The compensated case has 12-18dB of rolloff for the same frequency range, reducing the sensitivity to high-frequency instability. Simultaneously with increasing the ultra-high frequency stability, the compensated feedback increased the damping in the bandwidth of interest ($< 5\text{kHz}$).




Figure 4.12 Idealized lightly damped model without and with actuator

It should be stressed that although the structure of poles and zeros of the collocated actuator and sensor have been preserved in this example, the actual plant was much more complicated. It has a much more widely varying modal density, damping ratio, magnitude, and distribution. For the actual system, only the structural modes with high amplitude will be reduced, and those with the least damping will have the most active damping added. Indeed, the only frequencies that will be appreciably affected by this feedback strategy are those that have the highest gain within the bandwidth of interest. The controller does not target individual frequencies, since they are assumed to be unknown and variable. What is important, however, is that the interleaved structure of the poles and zeros of the collocated system allow for some performance. This is where the literature claims that theoretically infinite gain could be applied [35] [65] [27]. True, if pure collocated velocity feedback could be implemented without actuator dynamics or collocation errors, the infinite gain could be achieved without concern for robustness. Realistic

limitations on the actuator provide both a low-frequency and high-frequency constraint on the achievable gain. This is investigated further in Chap. 5.

4.3 Low-frequency Synthesis

The low-frequency controller was synthesized using an accurate model of the test structure. A set of Matlab code implementing the system identification routines of [102] was used to process multiple-run time-domain data on the test structure. Unlike the analog high-frequency controller, the low-frequency (and mid-frequency) controller would be implemented utilizing a sampled-data computer controller. For computational complexity, design simplification, and numerical precision reasons, the low-frequency controller was synthesized using a lower sampling rate than that required for the high-frequency controller. The target bandwidth of the low-frequency controller was $< 400\text{Hz}$, so the higher sampling rate required to implement control up to 1000Hz for the mid-frequency control (target sampling frequency $f_s = 5\text{kHz}$) was unnecessary.

4.3.1 Synthesis model

The control design for the low-frequency regime utilized the H_2 design process [113] well-established throughout the control literature [41] [43] [6] [38]. As discussed in [114] and [44], the optimal formalization of the LQG problem allowed for specific frequency weighting of performance parameters to tailor some performance specifications. This approach was utilized as in Fig. 4.13.

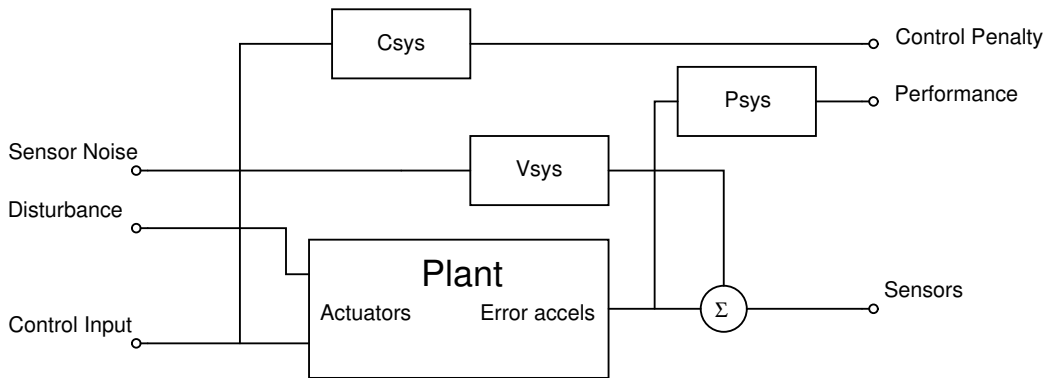


Figure 4.13 Low-frequency H_2 design model

The different components of Fig. 4.13 allowed for various adjustments in the design. The bulk of the “design” process with the H_2 controller became the specific choice of the additional augmentation parameters on the plant. With that done, the actual synthesis became a routine application of the existing **h2syn** Matlab function. The “standard form” for the H_2 design process uses a linear-fractional transformation (LFT) representation as a 2-port system. Figure 4.14 shows how the design model of Fig. 4.13 was considered in this form. The dynamic system K_{H_2} is the form of the resulting optimal controller that minimizes the cost function (4.16) [113].

$$J_{H_2} = \lim_{T \rightarrow \infty} E \left\{ \frac{1}{T} \int_0^T y_1^T y_1 dt \right\} \quad (4.16)$$

It should be noted that this standard form is a bit different from the identified model used throughout this document. The form of Fig. (4.16) has the disturbance modeled as the first inputs, whereas the original system identification used the actuator path as the first inputs.

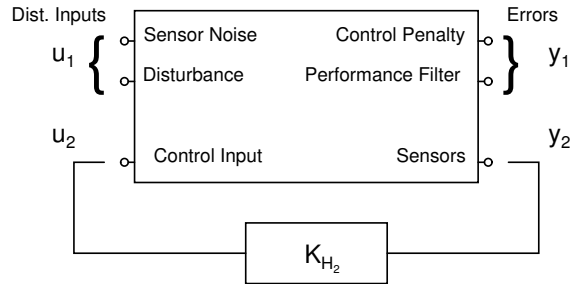


Figure 4.14 Low-frequency H_2 design model in standard form

The application of the H_2 synthesis tool assumes that the spectral content of the disturbance and noise inputs (u_1) are white. In order to adjust the desired frequency response of the closed-loop system, the additional systems present in Fig. 4.13 were added. The C_{sys} system provides a frequency-dependent weighting on the control input. This could be used to reduce the optimal controller’s response to compensate for actuator saturation, enforce high or low frequency rolloff in the actuator drive, or increase control penalties on specific frequencies. Similarly, the P_{sys} system allowed for *weighting* of specific frequencies in the error sensors. This design parameter could be used to arbitrarily weight some error frequencies more than others to counter the effects of a shaped error sensor response, or generally shape the desired error spectrum. Finally, the V_{sys} system provides a means to reduce the “reliability” of the error measurements. The effect of this was to reduce the resulting optimal controller’s attempted control authority by making certain frequencies “unreliable.” Whenever the noise spectrum shaped by V_{sys} becomes

larger than the response of the error sensors, the optimal controller will not weight errors in that frequency range as much since the measurements are unreliable. This effectively gives a parameter to limit the controller’s bandwidth.

The design model of Fig. 4.13 contained all of the pieces necessary to tailor the response to a particular plant model. In order to make reasonable design choices for these weights on the specific system, a model of the plant was required. As discussed in Sec. 3.3, the identification routine did not provide good, low-order models directly. Rather, a higher-order ID was done, and then reduced before synthesizing the controller. Figure 4.15 shows an example of this model order reduction. Due to the multivariable framework of the H_2 synthesis procedure, systems are generally represented as a singular value plots rather than individual input-output pairs of Bode transfer functions. This allows for simultaneous examination of the entire eigenvalue structure of the plant without having to compare many different Bode plots. Shown in this way, the model order comparison plot of Fig. 4.15 shows the singular values of the identified model with 200 states, along with a reduced version that only has 75. As it turned out, 75 states was about the minimum that adequately captured all the relevant dynamics up to the low-frequency sampling rate of 1kHz.

An example set of H_2 design filters with Fig. 4.15’s reduced-order model of the plant can be seen in Fig. 4.16. The open-loop response shows the maximum singular values of the low-frequency plant model. In this example, the control effort penalty C_{sys} was largely unused. What C_{sys} *did* attempt to capture was the reduced actuator effectiveness below 100Hz. It “turned-off” the control where the actuator was impractical to use. Similarly, the example V_{sys} system also had a relatively uninteresting frequency response. In practice however, the value of this scalar was used to change the optimal control “gain.” Increasing the magnitude of the sensor noise reduced the optimal controller’s performance.

The compensator resulting from applying the design filters of Fig. 4.16 to minimize the y_1 error of Fig. 4.14 can be seen in Fig. 4.17. The loop gain is in red and can be seen to have a fairly complicated spectral response. The black line shows the total loop gain for the actuator and H_2 controller path. Although there are only a few frequencies where the loop gain exceeds 0dB, the overall reduction in the plant’s most lightly damped modes is substantial. Figure 4.18 shows the open and closed-loop singular values of the total plant. Appreciable reduction (5-15dB) can be seen in almost all structural modes.




Figure 4.15 Low-frequency H_2 design model order reduction




Figure 4.16 Low-frequency H_2 design filters




Figure 4.17 Low-frequency H_2 design loop gain



Figure 4.18 Low-frequency H_2 example performance singular values

Another way of looking at the performance is with the actual Bode plots. Figure 4.19 shows both the actuator and disturbance path with the same H_2 design. This differs from the singular value plots because it can be used to verify that the actuator path is not driven to saturation or other potential nonlinearities. As can be seen from the lower magnitude plot in Fig. 4.19, most of the actuator path modes are reduced in gain because they are also present in the disturbance path. Since the actuator and error sensor are collocated, this is not too surprising. Interestingly however, some of the modes of the actuator path are *not* decreased, but still result in increased disturbance rejection.




Figure 4.19 Low-frequency H_2 example performance bode plot

4.3.2 Practical experimental synthesis

In order to apply the design synthesis, an accurate model of the *implemented* hardware system was required. This posed a practical problem. For the purposes of analysis as in Chapter 6, practically any reasonably high-order identified model fit to the high-frequency experimentally-measured data resulted in self-consistent analysis. In reality, the system identification routine was unable to reliably identify all of the system dynamics simultaneously. The SOCIT toolbox

[101] used for system identification requires a choice of the number of system modes for an input parameter. This choice must be of equal or greater than the actual number of structural modes present in the identified system [100] [102] for accurate convergence. The necessary amount of time-domain data sampled at the high-frequency often presented numerical instabilities as well as exceeding the memory capacity of the machine running the identification. Fortunately, this was not strictly necessary, as the low-frequency controller design is only minimally affected by all of the high-frequency dynamics of the plant and the additional controllers. These are investigated more in Sec. 6.4.

Instead of trying to produce an extremely high-order model only to reduce its order down by a factor of 10 or more for the low-frequency synthesis, a “hardware-in-the-loop” technique was utilized. This reduced the required complexity of the system identification, and had the additional benefit of directly providing the most accurate plant model possible. It included all unmodeled dynamics such as possible sub-periodic digital propagation delay through the controller’s A/D and D/A stage, as well as the specific delay introduced by the on-line interpolator and decimator. The latter is described in more detail in Appendix A. All of the actual blocks in the system can be seen in Fig. 4.20, and the resulting system identification was used for the *experimental* H_2 synthesis.

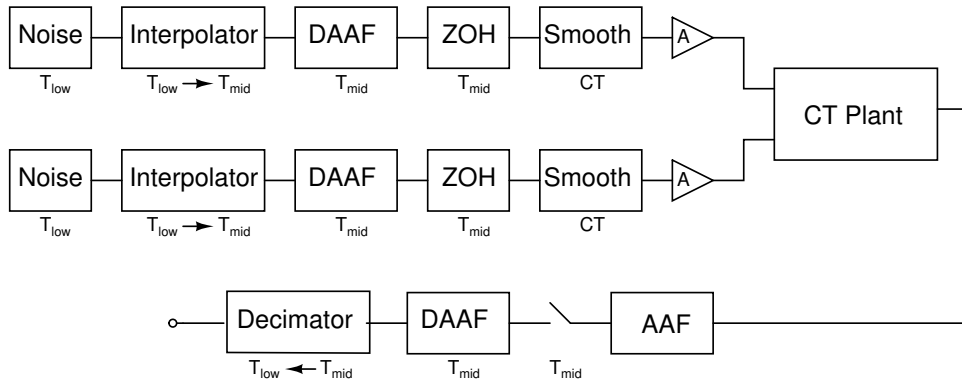


Figure 4.20 Low-frequency H_2 synthesis experimental setup

A comparison of the different model orders can be seen in Fig. 4.21. This compares the quality of the high-order system identification at the high-speed sampling rate ($f_s = 10\text{kHz}$) with that from the low-frequency decimated sampling rate ($f_s = 1\text{kHz}$). In each case, the model order was reduced to 100 states. In the high-frequency ID case however, the high-order

dynamics are no longer part of the model. Good representative frequencies for this are between 120-200Hz. The HF-reduced version has a much smoother and “averaged” response than the LF-reduced version.




Figure 4.21 Low-frequency H_2 synthesis model order comparison

As can be seen, the LF identification at only 100 states contains more system dynamics in the low-frequency region than the HF identification at an extremely high order model of 1000 states. The 197-state LF model was also closer to the original frequency response data computed from the single excitation time-domain raw data. From that, it can be assumed to be the “more accurate” model. As expected then, even when both models were reduced in order to 100 states, the original LF model remains much closer to the actual observed model.

4.4 Mid-frequency Synthesis

The basic structure of a single-channel feed-forward adaptive controller is shown in Fig. 4.22. The adaptation block adjusts the dynamics of the controller block in order to minimize the error. Many different forms of both the adaptation and the controller have been investigated

and implemented [81], but the most popular has been the LMS algorithm. It is a non-recursive controller filter, with a recursive gradient-descent algorithm that tries to minimize the squared error.

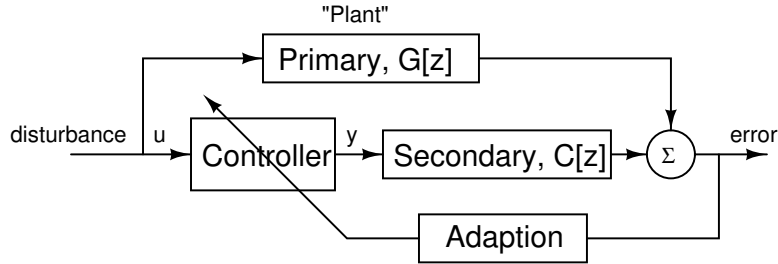


Figure 4.22 Basic adaptive feed-forward system

4.4.1 Secondary path model

Although there are a number of different variations on the basic feed-forward, adaptive control, the algorithm implemented was a filtered-x, least-mean-squares with adjustable leakage [83] shown in Fig. 4.23. The disturbance is the signal that perturbs the plant through the primary transfer function, $G[z]$. The controller sends outputs to an actuator on the plant through the secondary transfer function $C[z]$ in an attempt to drive the error to zero.

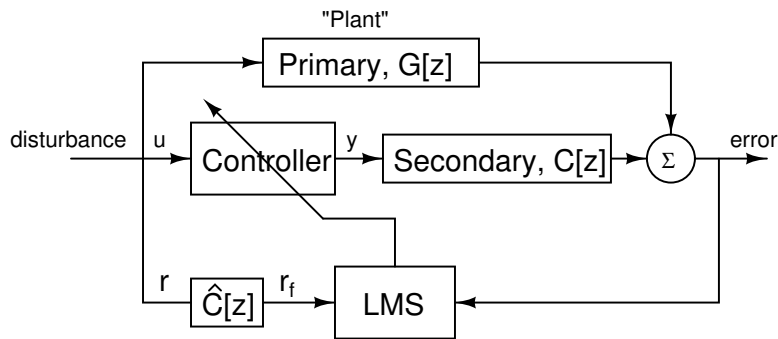


Figure 4.23 “Filtered-X” adaptive feed-forward system

The “design” of the adaptive controller has little to do with the structural model, and more to do with algorithmic constraints. The design constants are discussed in Sec. 4.4.3, but the one direct structural requirement was that of the model of the secondary error path, $\hat{C}[z]$. This required a model identification similar to that performed in Sec. 3.3. The only difference was a samplerate adjustment. The measurement rate of 10kHz was not necessary for the band-limited

implementation of the LMS algorithm. For analysis purposes, the samplerate could be directly changed and any errors introduced would at least be self-consistent. In order to continue to have the ability to see the dynamics *above* the mid-frequency controller’s sampling rate however, the measurement samplerate was maintained. Rather, another decimation similar to Sec. 4.3.2 was done. This time, the decimation factor was only 2, leaving a 2.5kHz Nyquist rate for the effective mid-frequency controller samplerate.

As discussed in [87] and [79], the model of this secondary path did not need to be extremely precise. The only requirement was that the phase modeling stay within $\pm\frac{\pi}{2}$ of the actual secondary error path. Any larger phase errors than this and the gradient-descent algorithm would no longer “descend” on the error surface. Rather, the adaptation would change the filter to *increase* the error and ultimately drive the system unstable. The gain accuracy of the model only limits the transient rate of convergence at that frequency. This has tangible performance limits for rapidly-changing systems, but the assumptions of this investigation do not require this.

The model ultimately used for the secondary path can be seen in the experimental results chapter Fig. 7.13 (pp. 172). It shows the decimated Nyquist rate and the model order vs. accuracy tradeoffs.

4.4.2 Time-domain algorithm

The transient performance of the LMS algorithm has been studied in-depth in [84] [77] [108] to name a few. This algorithm requires a model of the actuator to error path, and provides a filtered reference, $r_f[k]$, in addition to the disturbance reference, $r[k]$. The reference signal and error signal were measured at a time index $k + 1$. These two quantities, as well as the filter coefficients at time k were used to compute the updated filter coefficients in (4.17).

$$\begin{bmatrix} f_0[k + 1] \\ f_1[k + 1] \\ \vdots \\ f_{N-1}[k + 1] \end{bmatrix} = (1 - \lambda) \begin{bmatrix} f_0[k] \\ f_1[k] \\ \vdots \\ f_{N-1}[k] \end{bmatrix} - \mu \begin{bmatrix} r_f[k - 0] \\ r_f[k - 1] \\ \vdots \\ r_f[k - (N - 1)] \end{bmatrix} e[k + 1] \quad (4.17)$$

The filter coefficients are represented by f_0 - f_{N-1} , where N is the length of the filter, a design choice. The other parameters are the scalar gradient gain, μ , and leakage parameter λ . The

gradient gain is the size of the steps taken toward the instantaneous estimate of the error surface gradient. If this parameter is too large, the time-varying character of the adaptive system might interact with the normal dynamics of the plant. The large gain makes the gradient descent “overshoot” the minimum mean-squared-error point and the actual error can oscillate as it is reduced. This can cause instability as the feed-forward path tends toward dominantly *feedback* behavior through the LMS algorithm [77]. If the μ parameter is too small, then the transient adaptation performance suffers. For the purposes of this investigation, a value of $\mu = 0.1$ was eventually chosen. This value provided a good tradeoff between time taken for the filters to converge, and oscillatory/unstable behavior.

The other parameter, λ , was used in the algorithm computation. It is discussed in detail in [91] and [92]. This parameter essentially allows for some “leakage” in the history of the filter coefficients to allow erroneous information to have a finite history. This term is often added to realistic LMS implementations as it has a stabilizing effect when the system has model error or extraneous measurement noise. Too large a value in the λ parameter causes large steady-state error. A vanishingly small value makes the algorithm pure FxLMS. For the experimental implementation performed, although implemented, this parameter proved largely unnecessary. The experimental results of Chap. 7 were all performed with $\lambda = 10^{-9}$, so it was effectively removed.

The output equation of (4.18) shows the computation with the FIR filter coefficients from (4.17) at time-step $k + 1$.

$$y[k + 1] = \begin{bmatrix} f_0[k] & f_1[k] & \dots & f_{N-1}[k] \end{bmatrix} \cdot \begin{bmatrix} r[k - 0] \\ r[k - 1] \\ \vdots \\ r[k - (N - 1)] \end{bmatrix} \quad (4.18)$$

The samples of the incoming *unfiltered* reference signal of length N are multiplied by the filter coefficients. This is a standard FIR formulation discussed throughout many signal processing texts [84], [115]

4.4.3 Filter length

The primary design constraints for this type of control were ones of implementation. The only one directly relevant to the system were the secondary path model discussed in Sec. 4.4.1.

That “design” amounted to using the minimum of states necessary to adequately model the actuator to error path. The other parameters were discussed in Sec. 4.4.2. These were iterated upon for the experimental setup.

The final parameter, N , must be taken in context with the sampling rate. Since the LMS algorithm uses FIR filters, their primary limitation on a lightly-damped structure is the length of time they can represent. By definition their impulse response is finite of length N samples. The length of *time* this represents is the product of N and T , the length of a sampling period.

The decimation factor of 2 between the measurement sampling rate and the effective mid-frequency sampling rate discussed in Sec. 4.4.1 gave an effective Nyquist rate of 2.5kHz. This left plenty of transition region for the anti-aliasing filter, given that the target bandwidth for the mid-frequency controller was up to 1kHz.

The time representable by different filter lengths at the effective sampling rate of 5kHz was compared to the plant’s dynamics. The same damping data of structural characterization of Sec. 3.4.4 was used to compute the *time* necessary for an individual mode to decay. Figure 4.24 plots the results of this comparison, with the structural decay time filtered to get a representation of the mean. From the plot, a filter length of $N = 256$ has comparable impulse-response time to the majority of the mean of the plant above 400Hz. Below this frequency, the required time increases. Below 100Hz, the time increases significantly.

Using the information in Fig. 4.24, the value of $N = 256$ was chosen as the target filter length. It allowed for comparable impulse response time to the structure above 400Hz. Obviously, longer filters would be better, but this would present additional load on the computational resources. That would violate one of the design constraints in Sec. 3.1 of keeping computational complexity low. The specifics of the computer implementation are discussed in Appendix A.

4.5 Summary

This chapter discussed the process of the synthesis of the controller. It provided a rationalization for what ended up becoming a relatively complicated control topology. The rationalization included structural measurements from Chap. 3 to make choices regarding the frequency ranges and robustness constraints suitable for the experimental structure. Tradeoffs between




Figure 4.24 Filter length vs. structural damping

computational complexity, control type robustness, and performance were addressed as a means to tailor the different design parameters to the experimental structure.

The high-frequency synthesis utilized collocated pole/zero structure to obtain performance from a very low-order controller. Low-frequency synthesis used noise and error-weighting in an H_2 -optimal design framework to provide control “handles” that trade off performance for implementation stability. Finally, the various parameter choices made in the adaptive LMS feed-forward control were discussed and tailored to the characteristics of the experimental structure.

CHAPTER 5

HIGH-FREQUENCY PERFORMANCE OPTIMIZATION

The fundamental research goal was to achieve wideband disturbance rejection and vibration minimization. For the high-frequency analog controller, the collocated control should theoretically [22] [68] provide no limits on feedback gain. The practical limitations on this type of control [26] stems from the lack of perfect actuators and sensors. Perfect actuators would include infinite bandwidth, and exact collocation.

In order to apply the collocated control configuration to a realistic actuator/sensor pair, some modification to the direct-rate feedback was necessary. This basically required a gain reduction in bandwidths outside the areas of interest. This allowed a reduction in feedback gain *prior* to phase lag introduced from the imperfect actuator/sensor pair and allowed for closed-loop stability. This was examined with the root-loci plots of Sec. 4.2, but a methodology for designing the *specific* compensator was required.

This chapter presents the optimization procedure that was used to design the *specific* compensator for the high-frequency controller given the experimental actuator/sensor pair constraints. Classical loop-shaping techniques only provided a certain amount of guidance when the plant model was so uncertain. *Specific* gain and phase relationships for the actuator/sensor pair as modally dense as measured in Fig. 3.23 were unknown, but their *general trends* were. These trends provided stability constraints, but not a methodology for maximizing broadband disturbance rejection.

A numerical optimization routine was constructed to compute the best compensator to maximize the amount of this disturbance rejection. This optimization procedure removed candidate controllers that did not meet prescribed stability margins, and then ranked them according to one of a set of cost function metrics that attempted to quantify the broadband disturbance

rejection. The *best* candidate controller was the one ranked the highest by the optimization routine, and was then implemented on the experimental setup.

Section 5.1 discusses the setup of the optimization problem. The actuator characteristics are considered from a feedback stability constraint (gain and phase) point of view. The conversion of these constraints to an algorithmic optimization routine is discussed in Sec. 5.1.1, along with the broadband cost function used to determine the optimal solution.

Section 5.2 shows the results of the optimization routine. The results of the optimization for a simple (unity) plant are shown in Sec. 5.2.1. This allows the interaction between stability constraints, and optimal pole/zero locations to be clearly seen. The simplified actuator model was included in Sec. 5.2.2. The tradeoffs between stability margins are not as clear in this case, but it does provide a solution for implementation. The results are shown for two different stability margins on the modeled system.

5.1 Problem Specification

A straight-line approximation of the collocated actuator/sensor path can be seen in Fig. 5.1. This plot shows the overall goal of the optimization problem. The flat response in the center of the magnitude response is the bandwidth that the actuator/sensor pair can be considered “collocated.” This is the region used to provide an approximation to direct velocity feedback.




Figure 5.1 Analog loop performance bounds

At both ends of this nominally flat response, there is a phase limit to what was possible with feedback control. Previous work [116] has used similar low-order band-passing techniques, but with much less restrictive actuator path phase limitations. In this case, assuming a positive feedback convention, in order to remain stable, the phase must be kept away from integer multiples of 360° . At actuator resonance at low frequency, the phase response quickly transitions through almost 180° towards 0° at 0 frequency. At high frequencies, phase lag accumulates from a few different sources to cause a slow transition towards -360° . Both of these restrictions mean that for a real plant with this dominant behavior, some gain excursion outside of the center band will cause the closed-loop system to go unstable *even before the center loop gain exceeds 0dB*. Simply applying the “guaranteed stable” direct velocity feedback discussed in Sec. 4.2 is only guaranteed to either have no performance or be unstable (or both). Some sort of bandwidth limit was required on the feedback compensator in order to allow for in-band loop gain without causing out-of-band instabilities. The shaping of this compensator was dependent on the experimental actuator and structural setup.

The basic model of the electromechanical actuator was discussed in Sec. 3.4.1 and was found to have two critical frequency points. At the low frequency, the inertial mass and actuator spring’s resonance provided a lightly damped set of complex poles at approximately 120Hz. The associated -180° phase transition with the resonance presented a low-frequency feedback stability constraint. Above this resonance, the actuator performed with approximately constant gain until high frequency dynamics become more pronounced.

At the high-frequency, the stability constraint was also viewed as a phase margin constraint. A number of obvious and non-obvious mechanisms for phase rolloff at high frequencies were identified. The most significant and obvious phase mechanism was the inductive winding on the actuator itself. It provided an $\frac{R}{L}$ dominant pole at approximately 3kHz with an additional 45° phase lag at this frequency.

A non-obvious phase mechanisms stemmed from the physical size of the actuator itself. It was an inertial actuator and was physically mounted on the structure so as to provide not only a point for input, but also to support the actuator body completely. The force from the inertial actuator was transferred to its body through the winding’s bond to the outer case perimeter. The “collocated” sensor accelerometer was mounted on the opposite side of the plate from the actuator as seen in Fig. 5.2. Even the plate bending underneath the actuator had a non-

negligible phase contribution. This plate bending wavelength is discussed in Sec. 3.4.3 and implies the first bending mode *under* the actuator footprint in Fig. 3.19 is less than an order of magnitude higher in frequency than the bandwidth of interest. This could certainly provide some phase contribution well below the first mode.




Figure 5.2 Collocated actuator and sensor

One typical phase rolloff mechanism that was *not* present in the optimization is that of sampling delay. The high-frequency controller was implemented in the analog domain, the normal delay due to digital sampling by the computer was not present. This was one of the primary reasons for implementing the low-order control in the analog domain and its decision simplified the optimization process.

The choice of an analog compensator for the low-order high-frequency control, provided a reduction in the practical complexity of the controller that could be implemented. The phase characteristics of the collocated actuator/sensor required the use of a band-selective “filter” for a feedback compensator. The optimization of the pole zero placement had a few driving considerations:

1. Maximize disturbance rejection over a wide frequency band.
2. Provide suitable magnitude/phase constraints to the collocated actuator sensor pair.
3. Limit controller complexity to be reasonable given plant uncertainty.

The disturbance rejection for the system of Fig. 5.3 was computed from the closed-loop simplification of $Y(s) = W(s) + G(s)C(s)Y(s)$. The “output” $Y(s)$ was the collocated error

sensor signal in the physical implementation, and the plant, $G(s)$, was the collocated actuator to error path. The disturbance of $W(s)$ represents the aggregate disturbance path through the structure, but in this case was considered a simple input.

The consideration of #1 above was quantified as the power reduction through the disturbance rejection transfer function. A band-limited form of this integration is shown as the band-limited frequency integral of (5.1).

$$V(\omega_l, \omega_h) = \int_{\omega_l}^{\omega_h} \left\| \frac{1}{1 + G(j\Omega)C(j\Omega)} \right\|^2 d\Omega \quad (5.1)$$

This was a similar metric considered in [117], but in this case for a direct consideration of aggregate disturbance rejection over a limited bandwidth. The limits of integration, $\{\omega_l, \omega_h\}$ were considered to be additional boundary conditions on the optimization problem and determined by the specific bandwidth requirements of the control problem.

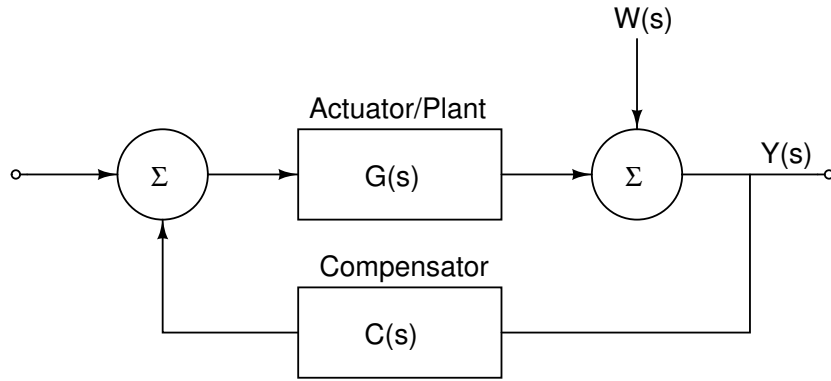


Figure 5.3 Plant disturbance rejection

The considerations of #2 and #3 above required a target compensator form of fairly low order. The chosen target form of the compensator was the lowest-order implementable bandpass filter of (5.2).

$$C_{\mathbb{R}}(s) = K_o \frac{(s - z_o)}{(s - p_{o1})(s - p_{o2})} \quad (5.2)$$

Although $\{p_{o1}, p_{o2}\}$ in (5.2) can completely represent any pole/zero combination, the implementable complex pole form of the compensator was more conveniently described with $\{\omega_o, \zeta_o\}$ in (5.3) when necessary.

$$C_{\mathbb{C}}(s) = K_o \frac{\omega_o^2 (s - z_o)}{(s^2 + 2\zeta_o \omega_o s + \omega_o^2)} \quad (5.3)$$

Even with this simple compensator, the resulting optimization problem was a nonlinear problem with 4 independent variables in \mathbb{C} . While mathematically any arbitrary combination of $\{K_o, p_{o1}, p_{o2}, z_o\} \in \mathbb{C}$ was valid, implementability constraints require $\{K_o, z_o\} \in \mathbb{R}$, and $\{p_{o1}, p_{o2}\} \in \mathbb{R}$ or $p_{o2} = p_{o1}^*$, where in this case, the “*” operator represents the complex conjugate.

The model of the plant was a simplified approximation that takes into account the underlying dynamics of the actuator/plate combination. The actuator resonance and inductive pole were included as from Sec. 3.4.1, and the first plate wavelength resonance from Sec. 3.4.3. The resulting simplified plant model of (5.4) was used to capture the low-order behavior for the optimization procedure.

$$G(s) = \frac{K_a}{(s^2 + 2\zeta_a\omega_a s + \omega_a^2)(s^2 + 2\zeta_p\omega_p + \omega_p^2)(s + R/L)} \quad (5.4)$$

This low-order behavior attempted to capture the *average* model of the modally-dense structure at the higher frequencies. This average model had nominally constant magnitude such as for the modally-dense plate of Fig. 3.21, while still maintaining the nominal phase rolloff at high frequencies. Also note that the active area of operation for the plant was *above* actuator resonance, but *below* plate resonance. Operating between these two critical frequencies as in the straight-line approximation of Fig. 5.1 allowed for *positive* feedback. The actuator resonance provided the required negative gain at mid-high frequencies where the controller would have loop gain.

The other point to note from (5.4) was that it represents a gross simplification of the actual structural dynamics. In reality, the simplified actuator dynamic mode of (5.4) are best thought of as a mean representation of magnitude, and a lower bound representation on the high-frequency phase. Consider the difference between the less-simplified model in Fig. 4.6 (pp. 59). The most negative phase lies *between* the poles and zeros, not actually *at* the points of highest gain. Thus, without trying to match every pole and zero of the structure with its complement in the compensator, the low-order compensator performs as well as is realistically possible. Such pole/zero cancellations were absolutely infeasible on such a plant with 1000 or more states. An ancillary observation to this restriction is that realistically, a compensator of such low order would only be capable of reducing the highest predominant peaks. The optimization procedure

must be viewed with this in mind, that only the points of high gain on the actual structure will be affected by the analog controller.

5.1.1 Setup

The optimization problem was investigated using a numerical optimization routine. The code was organized as follows:

Construct test set: A test set of potential real and complex poles were chosen from the two characteristic compensator forms of (5.2)-(5.3). This test set had to encompass a fine enough grid of the complex plane so as to accurately find the optimal, yet simultaneously remain small enough to be computationally tractable. This was done primarily by limiting the search frequencies to reasonable values. The frequencies of interest for the poles and zeros were chosen to be well below and above low and high-frequency stability margin points, respectively. For the real-only poles, they were chosen to be logarithmically spaced, with values between $2\pi \times 10^{1.5}$ and $2\pi \times 10^{4.5}$ (31Hz-31kHz). The complex poles were chosen to be between $2\pi \times 10^2$ and $2\pi \times 10^{3.5}$ (100Hz-3kHz) with damping ratios between 0.01 and 1. The zero choice was constrained to be real and between $2\pi \times 10^{1.5}$ and $2\pi \times 10^{2.5}$ (31Hz-310Hz). These choices provided a numerical grid of complex choices of poles and zeros subject to the implementability constraints of the hardware system.

Set stability boundary constraint frequencies: The primary boundary conditions required for the optimization were the location of the low and high frequency points of the actuator/sensor path. The low-frequency point was easy to define as the actuator resonance frequency, since it provided the worst-case for gain and phase margin. The resonant peak had the highest gain, and the steepest phase transition from -180° back to 0° so it provided a simple choice for the low-frequency stability point. The high-frequency point was less defined since it depended on a number of both quantifiable and non-quantifiable phase rolloff mechanism of the hardware system. The most dominant one was the inductive rolloff, but the plate bending wavelength began to introduce some phase rolloff at the high frequencies as well. For the optimizations iterations, the high and low stability margin frequencies were set at 120 and 5000 Hz. The low-order of both the actuator model and candidate compensator form required

the margins only need be checked at and outside of these two frequency ranges. The candidate compensators that met these stability requirements also had the benefit of resulting in “sensible” in-band results.

Set boundary constraint margins: This design parameter proved to be the most important for the results of the optimization. The low-order model of the plant used for the optimization captures only the low-order response, and does not attempt to address the hundreds of modes actually present in the physical plant. Because of this, the robustness parameters most suited to limit the performance were the gain and phase deviation at the extremes of the control bandwidth. Ranges in gain margins from 0 – 20 dB, and phase deviations from 0° – 150° were investigated and discussed in the results of Sec. 5.2. Rather than the traditional definition of phase margin as the point at which unity gain occurs, in this context it made more sense to look at the phase at the two critical frequencies. The actuator resonance was at a specific, well-known frequency. At the high-frequency boundary however, the actual stability mechanism was somewhere above 5kHz, but was poorly-predictable. Defining the phase deviation at a single frequency at least provides a fixed point by which to compare designs. This “phase-deviation” constraint was set as a maximum deviation from zero phase offset.

These stability margins were very comparable to those used in traditional loop-shaping techniques, with the added complexity of an uncertain plant. The gain margin provided an adjustment as to how much gain was allowable at the two critical frequency points. Because the plant actually exhibited resonant gain excursions about the nominal model, this gain margin provided a means to remove controllers that were “too optimistic” to be implementable. A similar justification can be made for the phase deviation margins. By providing a maximum allowable phase excursion from unity, the robustness to unmodeled phase variance in the experimental plant could be increased. Since the *specific* model of the plant was not known, the *actual* phase at any point could be widely variable between the established bounds. The phase deviation margin parameter allows the optimization routine to remove controllers that were “too optimistic” from a phase standpoint as well.

Eliminate invalid sets: The stability criterion provided the primary means for eliminating test set members. Any pole/zero combination that did not meet the specified gain and phase

deviation constraints at the frequency boundaries was discarded. The set was then further reduced in size by eliminating any compensator that did not provide sufficient loop gain. In order to reduce the number of variables in the test set to just pole and zero locations, the gain of the compensator was chosen *after* determining a set’s pole/zero combination had sufficient phase deviation margin. The gain was then set to the maximum allowable that still met the gain margin criterion at both high and low frequency points.

The bandpass shape expected from the optimization process yielded a peak magnitude in the middle of the passband. Any otherwise acceptable compensator that had peak gain less than a chosen δ was discarded. The δ variable was typically set to 0 dB as per the common definition of “control authority.” For some particularly strict set of stability boundary conditions, however, reducing this requirement to down to -6 dB peak gain at times allowed for more valid test sets and a clearer picture of the optimization surface. This was not unrealistic since a specific instance of the plant may very well exhibit peak gain excursions well above the ‘average model used here. A loop gain of “less” than unity in the mean could still have more than unity gain at these modes.

Rank sets according to cost function: The cost function of (5.1) provided the basis for ranking the valid pole zero sets that satisfied the stability criterion. The full bandwidth aggregate disturbance rejection was computed with $\omega_l = 0\text{Hz}$, and $\omega_h = 5000\text{Hz}$. Unfortunately, when taken under this extremely wide bandwidth, the cost function was not a good metric of the performance of the *entire* controller. The analog controller was only to emphasize the high-frequency spectrum on the final implementation. Spillover at the low frequency was not a concern. At the high frequencies, the rolloff of the controller was not very abrupt either, so some spillover was inevitable there as well. As a consequence of this, the choice of bandwidth to measure the *band-limited* high-frequency disturbance rejection was between 200-2000Hz. The lower limit was chosen so as to not include the spillover in the actuator resonance in the cost function computation, although it was still considered as a stability constraint. The upper limit was chosen as after considering the full-bandwidth cost function. All of the valid solutions exhibited an inflection point in disturbance rejection performance around 4kHz. The rapid change decreased the usefulness of that frequency as a measure of the *broadband* performance.

The value of 2000Hz was the center of the bandwidth of interest and showed the most discretion between the different solutions.

One other ranking method employed to analyze the optimization was to use the peak gain excursion of the compensator. This of course was related to the frequency separation between the poles and zeros as well as the damping ratio if the solution was complex. The results of this method are discussed primarily in Sec. 5.2.1

Visualize optimization result: The final step in the optimization process was to interpret the results. As mentioned in Sec. 5.1 maximizing the broadband disturbance rejection of (5.1) was the goal. The band-limited computation of this integral provided the rankings of the valid test set, but does not provide any insight into the results. A frequency-progressive modification was included in (5.5) to get a frequency-response representation suitable for plotting.

$$V(\omega) = \left(\int_0^\omega \left\| \frac{1}{1 + G(j\Omega)C(j\Omega)} \right\|^2 d\Omega \right)^{1/2} \quad (5.5)$$

While it would be possible to simply show the final value of the integral of (5.5), plotting it as a function of ω had additional benefits. It allowed the tradeoffs between out-of-band spillover and in-band-performance to be visualized and compared between different optimization results. It also provided the same means of visualizing the aggregate disturbance rejection that was used to characterize the experimental results in Chap. 7. The complexity of the gain response on the real plant was too complicated to get a good “big picture” of the broadband performance from a simple Bode plot representation, so the progressive RMS was *necessary* for interpretation there.

Another modification to the cost function was useful to aid in interpreting the results. The progressive frequency integral of (5.5) added all control performance as a linear function of frequency. When plotted directly, this represented the *power per Hz* in the disturbance rejection function and increased with frequency. In order to visualize the performance through a broad bandwidth on similar scales, a *frequency-weighted* progressive integral was also used. The full-bandwidth version of this is shown in (5.6), and the band-limited version in (5.7).

$$V(\omega) = \left(\frac{2\pi}{\omega} \int_0^\omega \left\| \frac{1}{1 + G(j\Omega)C(j\Omega)} \right\|^2 d\Omega \right)^{1/2} \quad (5.6)$$

$$V(\omega)|_{\omega_l < \omega < \omega_h} = \left(\frac{2\pi}{\omega - \omega_l} \int_{\omega_l}^\omega \left\| \frac{1}{1 + G(j\Omega)C(j\Omega)} \right\|^2 d\Omega \right)^{1/2} \quad (5.7)$$

All four different permutations for a representative band-pass compensator are shown in Fig. 5.4. The full-band progressive integral increases continuously since it represents the aggregate power summed throughout the spectrum. The band-limited version does the same, although it is truncated. The frequency-normalized progressive integral clearly illustrates where disturbance rejection performance increases and decreases on a per-Hz basis.




Figure 5.4 Cost function variations

5.2 Optimization Results

The results of the optimization were relatively difficult to interpret. Of course, the end result provided the relevant poles, zeros, and gains for (5.2) or (5.3). The relative merit *between* possible solutions which resulted in the optimal solution was not necessarily clear. As it turned out, the results of the optimization were almost entirely dependent on the boundary conditions. The chosen phase and gain margin at both critical frequency endpoints were primarily responsible for determining the pole locations (or pole and damping if complex). Whether or

not there were any real or complex values *even in the solution set* depended on the boundary conditions.

In order to understand this a bit further, the first optimization was performed on a trivial unity plant in Sec. 5.2.1. This provided the framework for comparing the optimization results to each other without the frequency asymmetry associated with the actuator path. The tradeoffs between the boundary conditions and the solutions were more clearly visible in the case of both strict and lenient stability requirements.

Section 5.2.2 shows the results on the simplified model of the system. Although they become harder to interpret, the result is a choice of an optimal compensator subject to the set of stability constraints established.

5.2.1 Unity plant

The unity plant model was optimized under two different sets of stability constraints. The primary purpose for this was to determine what set of boundary conditions would yield real poles vs. complex-valued poles. It also allowed for a clear comparison between the compensator and resulting disturbance rejection, since the “plant” was trivial.

Recall from Sec. 5.1.1 that the two stability frequencies points chosen were at 120Hz and 5000Hz. The first set run through the optimization were the lenient constraints of 3dB of gain margin and 60° of phase deviation margin. The set of pole/zero/gain locations that solved all of the constraints can be seen in Fig. 5.5. The set’s index number is plotted on the x -axis. The y -axis shows the $\log_{10}(\frac{-\Psi}{2\pi})$ where Ψ is one of a zero, real pole, real component of complex pole, or imaginary component of complex pole. The gain was increased to the maximum allowable by the most restrictive gain margin constraint was plotted as a simple \log_{10} .

This type of plot provided the ability to directly and simultaneously compare all parameters of all of the optimization solutions. Any given set’s value are vertically lined up with one another. The actual number of solutions was relatively unimportant, but rather was an indicator of how fine of a grid of the complex plane was chosen as the possible solution set, and how stringent the boundary constraints were. The grid was increased to where an adequate characterization of the solution surface could be seen.

Figure 5.6 shows exactly the same solution set as in Fig. 5.5. In this case however, the x -axis is now not the set’s index number, but rather the outcome of the cost function for a

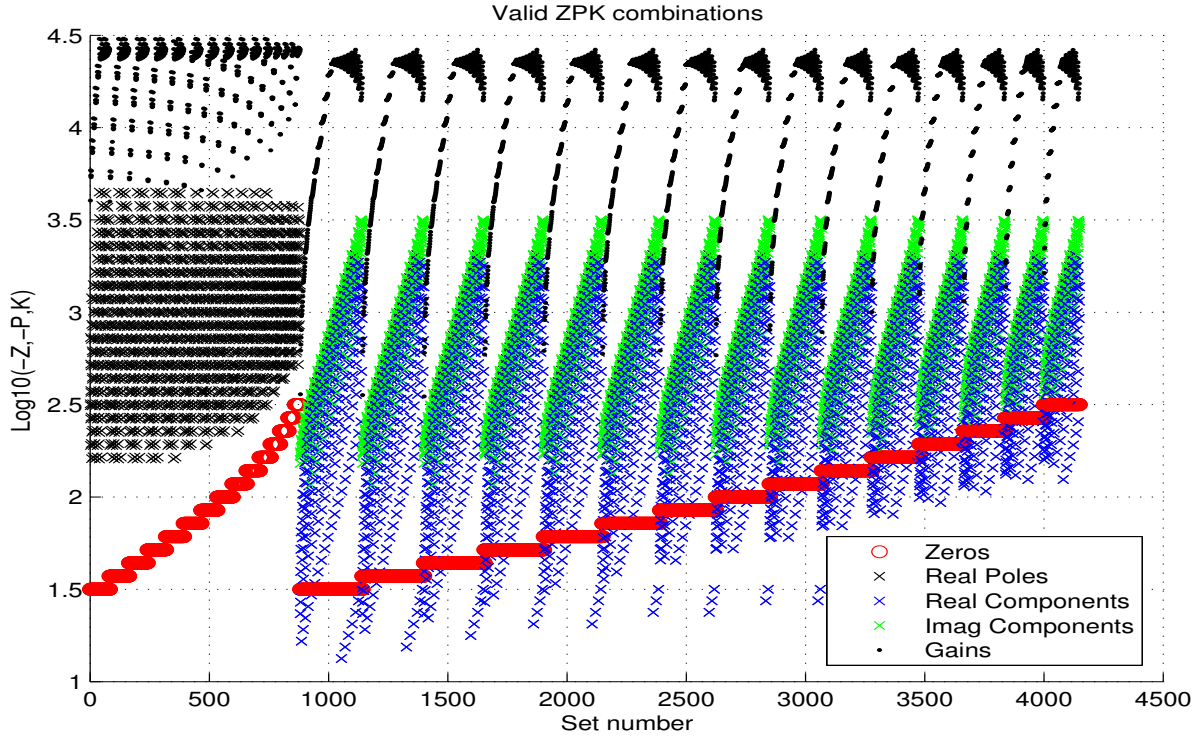


Figure 5.5 Valid pole/zero/gain (P/Z/K) sets for lenient constraints on a unity plant

particular set. The set's pole/zero/gain are still directly readable in a vertical line, but now the slice of the solution surface can be seen. The cost function was the band-limited progressive RMS integral of (5.7). The optimal solution for this set of constraints and cost function is read as the right-most set of values. In this case, it was

$$\{z = -2\pi 10^{1.9}, p_{o\{1,2\}} = -2\pi(10^{2.2} \pm j10^{2.9}), k = 10^{4.35}\}.$$

With these optimal values, the compensator was computed. Figure 5.7 shows the Bode plot of the loop gain for the solution. It simultaneously provided no more than -3dB of gain at both 120 and 5000Hz, nor a phase deviation of more than 60° and has the expected band-pass shape. The disturbance rejection associated with this solution can be seen in Fig. 5.8.

The relatively lenient stability requirements for the optimization resulted in an underdamped, complex compensator response. Although meeting the definition of “optimal” as defined by (5.7), it does not provide much confidence in the cost function as a wideband disturbance rejection metric. In order to see if it was a limitation in the cost function or an artifact of

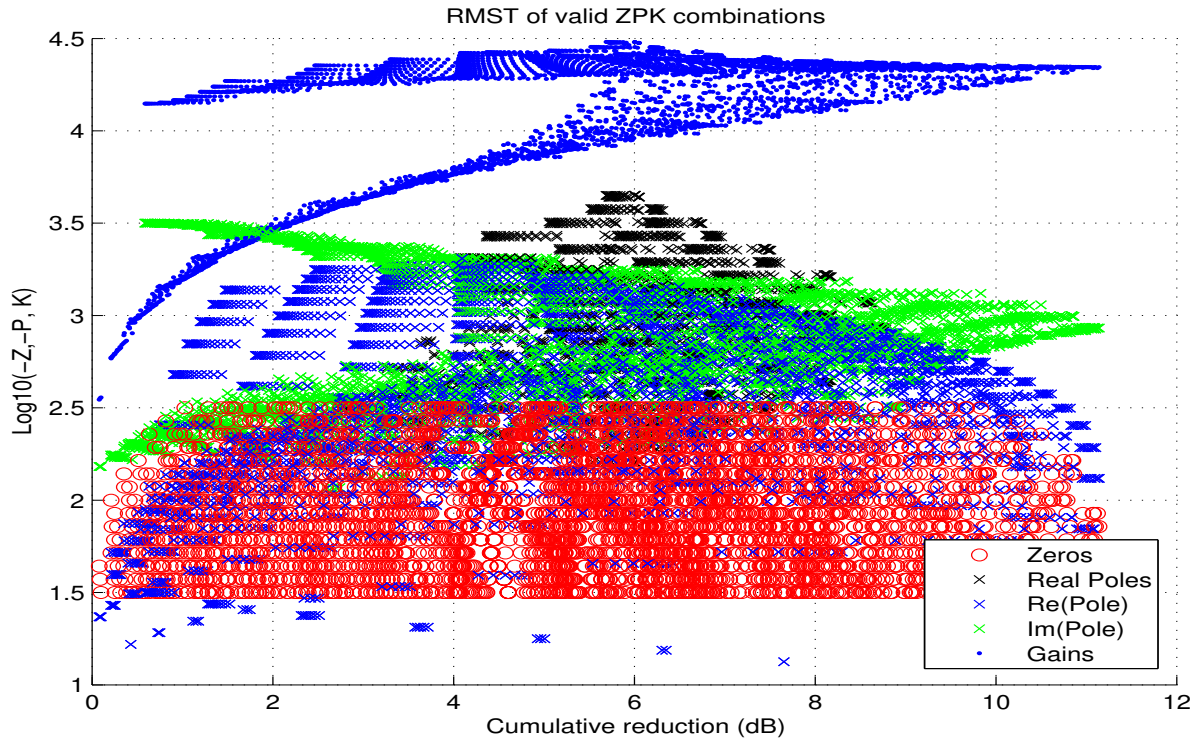


Figure 5.6 Lenient phase P/Z/K mapping by band-limited RMS integral

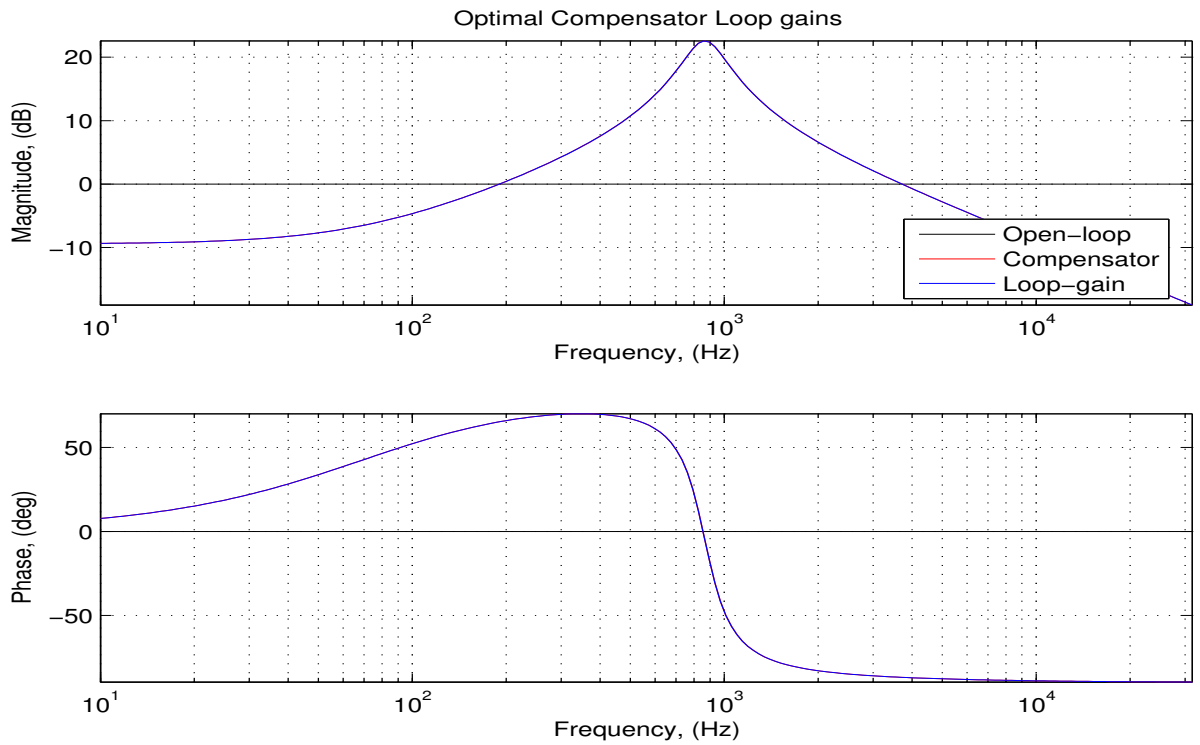


Figure 5.7 Lenient phase optimal loop gain




Figure 5.8 Lenient phase disturbance rejection and performance

the boundary conditions, a new set of stability constraints were chosen. These were the same 3dB gain margin, but this time changing to 30° phase deviation margin at the low frequency.

The mapping of the optimal compensator poles/zeros/gains for this set of constraints is shown in Fig. 5.9. It uses the same band-limited progressive RMS integral as a cost function as that of Fig. 5.6. The difference is clear that the optimal solution now has real poles rather than complex poles. In this case,

$$\{z = -2\pi 10^{2.05}, p_{o1} = -2\pi 10^{2.6}, p_{o2} = -2\pi 10^{2.9}, k = 10^{4.4}\} ..$$

The intuitive reasoning for this is that the more strict phase requirements does not provide enough spectral “room” to allow for significant resonant behavior from the optimal compensator. Indeed, even the best complex solution at poles $\{p_{o\{1,2\}} = -2\pi(10^{3.3} \pm j10^{2.25})\}$ around 6dB on the cost function represent very highly damped poles. Even these cease to be a viable solution once the real poles split.




Figure 5.9 Strict phase P/Z/K ranked by band-limited RMS integral

The Bode plot of the optimal compensator given by the mapping of Fig. 5.9 can be seen in Fig. 5.10 alongside the results for the lenient compensator. Figure 5.11 shows the different disturbance rejection characteristics for the two sets.

The most striking difference between the two optimization results is the difference in the peak loop gain for the lenient stability margins compared to the strict margins. The strict optimization resulted in real poles being optimal, with very smooth magnitude and phase response and the expected band-pass behavior. The lenient optimization results also showed a “band-passed” behavior, but with almost 20dB more peak gain response.

At first glance, the complex (lenient) response appears to be better in all respects than the real-pole solution. The gain was much higher over almost the entire bandwidth of interest. What needs to be emphasized is that the *stability* margins between the two solutions were significantly different. The allowable phase deviations at both the high and low frequencies were much larger in the lenient case. The optimization found a solution along that boundary condition that maximized the disturbance rejection.




Figure 5.10 Lenient vs. strict phase loop gain




Figure 5.11 Lenient vs. strict phase disturbance rejection

From an intuitive sense, the lenient solution already readily exhibits potential implementation problems. Actuator resonance at 120 Hz realistically would go unstable with 60° of phase lead. Likewise, with a 20dB *gain* and 50° phase *lag* by 1kHz, the closed-loop plant would likely be unstable shortly above that frequency as well.

The biggest benefit of comparing the two solutions with the simple unity plant was that the boundary condition tradeoffs were clearly visible. Sufficiently lenient stability margin requirements allowed complex compensator design. With more stringent margins, only real-valued poles could provide the “band-passing” effect without causing too much phase deviation. In Sec. 5.2.2, the model of the actuator dynamics are included and serve to eliminate most of the overly-optimistic compensator solutions from the solution set.

5.2.2 Modeled plant

With the tradeoffs associated with the boundary conditions more clear from the unity plant investigation of Sec. 5.2.1, the simplified model of the plant was added. The optimization routine was performed numerous times with various boundary constraints on the simplified model. From a simplistic standpoint, the expected result of the optimization problem would be inferred from Fig. 5.1. Basically, the lowest possible frequency for the zero and highest for the poles would provide a compensator with the widest bandwidth. With complex poles allowed, the expected result would be similar, but the amount of damping that would provide the greatest broadband performance was not obvious. In general, these predictions were borne out by the results of the optimization, but with a few similar dependencies on boundary conditions as with the results of the unity plant in Sec. 5.2.1.

The results of the optimization with mid-range stability constraints can be seen in Fig. 5.12. The gain and phase deviation margins were set to 3 dB and 60° , respectively. The blue line shows the open-loop plant, and all possible valid compensated loop gains are shown in gray. The darkness of the shading color of the gray candidate loop gains indicate the ranking of that specific solution to the band-limited cost function of (5.1). In addition, the “best” solution according to band-limited, full-band, and peak gain performance metrics are explicitly shown.

A number of important points can be seen in Fig. 5.12. The open-loop plant shows a low-order approximation to the actuator/sensor combination. The low-frequency resonance and high frequency rolloff are clearly seen. The phase leading characteristic of the compensator becomes




Figure 5.12 Candidate compensated loop gains

a detriment to stability at the low frequency. At the high-end, the additional phase rolloff of the compensator also provides a limit. For this set of boundary conditions, the band-limited RMS best also had the highest peak gain. The full-bandwidth RMS result has a significantly lower performance rating in the band-limited sense. The relationship between these two can be seen in Fig. 5.13 and Fig. 5.14.

Figure 5.13 plots the same set of solutions' band-limited RMS progressive frequency integral of disturbance rejection (5.7). The open-loop response has no disturbance rejection, so the frequency-weighted integral asymptotically tends toward 0 dB. The various compensators' aggregate control per Hz is again ranked by gray intensity. The full-bandwidth costs' solution shows a marked reduction in the aggregate disturbance rejection. This is because the full-bandwidth penalizes the out-of-band spillover so much that the in-band performance of its "best" choice suffers. Figure 5.14 shows this cost function. The spillover at low-frequencies and high-frequencies both *magnify* the disturbance in those ranges. This is not as problematic as it may appear, however. The low-frequency controller has significant authority and a highly accurate model at the low frequencies and runs once the high-frequency controller is enabled.




Figure 5.13 Candidate ZPK combinations' bandlimited RMS performance

At the high frequency, any structural damping and passive absorbers become more effective at reducing the disturbance propagation. This was the initial reason for using the band-limited version of the cost function to compute the optimal controller for the final implemented system.

The actual parameters that resulted from the optimization can be seen in Fig. 5.15. Again, these are the unranked pole/zero/gain combinations that were solutions to the stability constraints. The gains were again increased to the maximum possible without exceeding the any stability margins.

Figure 5.16 shows the evolution of all optimized parameters in Fig. 5.15 with respect to the band-limited progressive RMS cost function. In this case, the optimal solution was with parameters

$$\{z = -2\pi 10^{2.2}, p_{o1} = -2\pi 10^{2.9}, p_{o2} = -2\pi 10^{3.7}, k = 10^{4.7}\}.$$

Recall that for this set of boundary conditions, the peak gain performance metric produced the same results as for the band-limited RMS performance metric. The most interesting point to note from the plot is that the maximum gain was *not* produced from a test set with complex




Figure 5.14 Candidate ZPK combinations' full-band RMS performance




Figure 5.15 Valid P/Z/K sets

poles, but rather asymmetrical real poles. The stability constraints were not sufficient to allow a very undamped complex response. This results is similar to that in Sec. 5.2.1.




Figure 5.16 P/Z/K map by bandlimited RMS (3dB GM, 60°PM)

To show the sensitivity of the optimization problem to boundary conditions, another set of conditions was considered. All parameters were unchanged except gain and phase deviation margins at the critical frequency points. The gain deviation margin was increased from 3dB to 6dB, while the phase deviation margin was decreased from 60° to 45°. This set of boundary conditions had a very different optimization result shown in the map of Fig. 5.17.

In this case as well as the previous case, the best-ranked solution for both the band-limited progressive RMS was the same as for the peak gain. The actual values are

$$\{z = -2\pi 10^{2.2}, p_{o\{1,2\}} = -2\pi(10^{3.75} \pm j10^{2.85}), k = 10^{4.45}\}.$$

What was interesting was the different shape of the pole/zero locations as a function of their ranking. Rather than a bifurcation in the real pole frequencies as in Fig. 5.16, a single frequency of $f_{1,2} \approx -10^{3.7}$ was favored. The phase constraints were stringent enough to limit the amount of peakiness of the response, so it was effectively a double-real pole. As an interesting point to




Figure 5.17 P/Z/K map by bandlimited RMS (6dB GM, 45°PM)

note, Fig. 5.18 shows exactly the same solution set, only mapped by peak gain. The optimal solution was the same, but the solution surface was scaled much differently. Rather than being clumped together, the real solutions are spread out in performance just as much as the complex solutions.

The loop-gain response of this configuration can be seen in Fig. 5.19. The small contribution from the imaginary component of the poles can be seen at approximately 3kHz.

What is not shown in the example of Figs. 5.17-5.19 were the numerous other iterations which produced no complex-valued poles. The stability margins on the plant did not allow for many undamped responses as significant as that of Fig. 5.7. This could be attributed to the additional phase effects of the actuator resonance and high-frequency roll-off, rather than the simple, static phase of the trivial unity plant. In any case, even when complex-value poles *were* solutions, they were either non-optimal like those of Fig. 5.16, or so highly-damped to be effectively double real as in Fig. 5.17




Figure 5.18 P/Z/K map by peak gain (6dB GM, 45°PM)




Figure 5.19 Candidate compensated loop gains (6dB GM, 45°PM)

5.3 Summary

This chapter discussed the numerical optimization routine that was applied to the high-frequency controller synthesis for the experimental structure. The simplified model of the system that was only similar to the plant as a very low-order approximation precluded the use of classical loop-shaping techniques such as Nyquist or root-loci manipulation for computing the *best* solution. Rather, a number of frequency-dependent cost function metrics were devised. The optimization routine consisted of solving for all possible solutions to an implementable compensator of one of two forms, then ranking the solutions according to these cost functions. The preferred cost function was the band-limited, progressive RMS integral. It provided the most direct indicator of the controller's performance *in the band of interest* compared to the other controllers to be added to the structure.

The results from the investigation were that the outcome of the optimization was highly dependent on the stability margins chosen as boundary constraints. These stability margins effectively became abstracted design criteria. Very few complex-valued poles were optimal, and those that were ranked "best" were dominated by the real component of their complex-conjugate pole pair. For this reason, the best performance solution to a reasonable set of boundary conditions (Fig. 5.16) was used to give the optimal pole/zero/gain locations for the experimental implementation.

CHAPTER 6

AGGREGATE CONTROLLER ANALYSIS

The synthesis of the three controller types was largely done independently. In this chapter, the analysis of all three controllers together on the same system is investigated. A significant portion of this analysis deals with the intricacies of the combination of continuous-time and discrete-time controllers. Unlike usual controller implementations, the advantage of significant over-sampling was not an option. Aspects of system aliasing and hold characteristics of the discrete-time controller could not necessarily be ignored as small effects. Also, since the continuous-time controller's bandwidth extended above the Nyquist rate of the digital controller, its effect on the aliasing and system identification could not be taken for granted as minimal.

The sampling aspects of both the simple plant, and the plant under continuous-time feedback are investigated in Sec. 6.1. This section reviews the fundamentals of sampled-data systems, and the performance requirements of the necessary anti-aliasing filters. Approximations to the continuous-time plant as represented in a discrete-time framework are justified. These approximations include an uncertainty bound on the aliasing error with and without continuous-time feedback control active.

Section 6.2 discusses the additional error in modeling associated with the multi-rate nature of the chosen control configuration. For consistency in comparing predicted vs. measured performance, the measurements of the system had to be done on the controller computer. This necessitated a higher sampling rate than that for the control. Because of this, both the low-frequency controller and mid-frequency controller were digitally decimated and interpolated. The assumptions made and errors introduced in these operations are discussed in Sec. 6.2. Strictly speaking, this step was not necessary for the mid-frequency control. Indeed, for a final

controller, the additional rate changing would not be desirable, but rather the entire computer would run at a lower sampling rate. Because the bulk of the computations were done at the decimated sample rate however, the processing overhead of this multi-rate sampling was minimal.

Section 6.3 details how the optimal mid-band controller was computed from the analytical model of the system. The transient evolution of this process was deemed unimportant in this investigation. Only the steady-state, optimal solution was used in the analysis. Due to the modeling difficulty of the continuous-time and multi-rate discrete-time aggregate system model, a literal interpretation of the LMS algorithm in the time-domain was used to compute the optimal solution.

The mixed continuous-time, discrete-time aggregate system analysis done in Sec. 6.4 is presented chronologically. The control loops were closed sequentially, so this was the most logical way to approach the analysis. Each loop was closed in turn, and the resulting aggregate system considered. At each stage, justified simplifications were made that allowed for a concise mixed-signal model of the total system. The affect of the different configurations were included in the optimal LMS filter computation from Sec. 6.3.

Finally, Sec. 6.5 looks at the stability of the high-order, low-frequency controller design. The sampled-data aliasing analysis, error propagation and maximum error bounds from the aggregate system model resulted in a total system uncertainty for the low-frequency model. This uncertainty was compared to the worst-case stability bound and the system shown to be stable.

6.1 System Identification

The analysis of the structure and control presented a modeling problem. The biggest problem was finding an acceptable way to accurately model an experimental system using a digital computer to measure the response, but without the usual assumption of a high sample-rate and negligible aliasing. The measurements needed to encompass spectrum *above* the Nyquist rate for the highest-frequency digital controller in order to be able to measure the performance of the continuous-time controller.

As a compromise between *measurement* and *control* computer performance, the sample-rate of the computer was chosen to be higher than the control rate. By including a low-order digital anti-aliasing filter, smoothing filter, and integer decimators and interpolators, the control sample-rate could be reduced while maintaining a higher measurement sampling rate with minimal computational overhead. Appendix A discusses the hardware and software implementation in more detail.

The block-diagram manipulation and techniques presented in [104] were used for this analysis. This technique utilizes the mathematical model of an impulse sampler, and the “*” notation to denote any discrete-time signal or system model. The relevant fundamental assumptions of this sampling are summarized in (6.1)-(6.4) for the characteristic sampled-data system shown in Fig. 6.1

$$r^*(t) = \sum_{k=-\infty}^{\infty} r(t)\delta(t - kT) \quad (6.1)$$

$$R^*(s) = \sum_{k=-\infty}^{\infty} r(kT)e^{-skT} \quad (6.2)$$

$$R^*(s) = \frac{1}{T} \sum_{k=-\infty}^{\infty} R(s - jn\omega_s) \quad (6.3)$$

$$H_z(s) = \frac{1 - e^{-sT}}{s} \quad (6.4)$$

Representing the impulse modulation with its Fourier series, the classical definition of signal spectral aliasing of a signal’s Laplace transform is shown in (6.3).




Figure 6.1 Impulse-sampled, zero-order held signal

The experimental setup for identification can be represented exactly using this technique. Figure 6.2 shows the experimental setup within the sampling framework, including all holds (H_z), anti-aliasing filters (B_a), smoothing filters (B_s), and gain stages (A). The analytical input-output characteristics are not affected by the movement of the continuous-time anti-aliasing filter, $B_a(s)$, to the input side as in 6.2(b). All continuous-time filtering can then be described by a single filter, $F(s) = H_z B_s B_a A$ as in 6.2(c). The system was excited directly

from the computer controller at $d^*(t)$, and the response was measured as $e_f^*(t)$. The time-domain data was fit using the NASA SOCIT toolbox [102] as in Chap. 3 and resulted in the input-output frequency characteristic of:

$$\begin{aligned} \frac{E_f^*(s)}{D^*(s)} &= (H_z(s)B_s(s)AB_a(s)P(s))^* \\ &= (FP)^* \end{aligned} \quad (6.5)$$

For the purposes of analysis, however, the combination of (6.5) does not immediately provide a model of P alone, but rather a sampled-version of it.

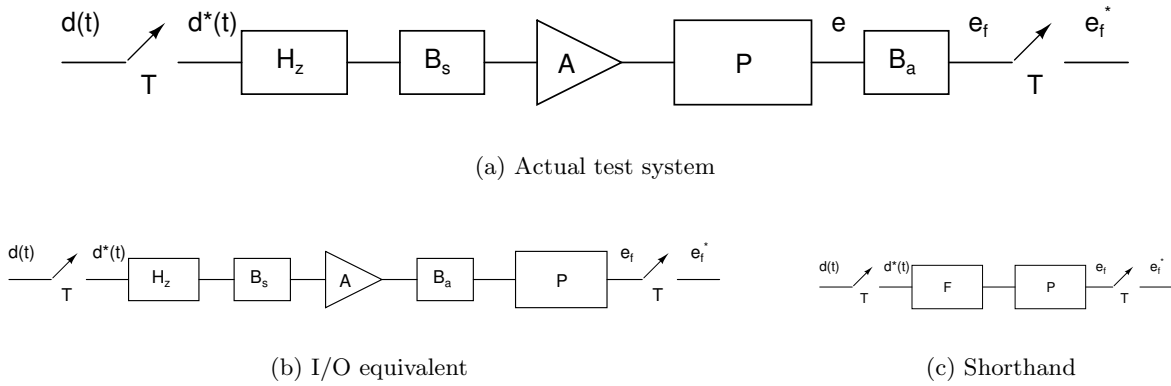


Figure 6.2 Mixed-signal plant ID setup

In order to obtain a model of P by itself, a model of the anti-aliasing filter loop, F , was required. The filters and sample and hold circuits used were well known as 4th order Butterworth filters and zero-order holds. What was not completely known was any slight gain mismatches in the filters, delays in the hold, or additional dynamics within the power amplifier. In order to capture all of these, the plant was removed from the system, and the filters were run through the system identification procedure by themselves. This yielded a model of $(F)^*$. Conceptually, this provides a method which could be used to obtain an approximate model of the plant by itself with

$$\tilde{P}^* = \frac{(FP)^*}{F^*} \quad (6.6)$$

There are a few potential problems with (6.6), however. The first is that the normalization model of F^* appears in the denominator, which effectively inverts the plant and maps zeros to poles. This causes problems when the identified model contains non-minimum-phase zeros as

those identified in Sec. 3.3.3. From inspection, it appears obvious that \tilde{P}^* should be causal and stable, but the numerical computation was not necessarily so. Since both F^* and $(FP)^*$ are obtained from identified models, non-minimum-phase zeros from the identified model mapped to an unstable system when inverted. Also, pole-zero cancellations were not exact so even if they *should* cancel completely, the experimentally-derived models would not and result in numerical problems.

Rather than try to compute the inverse of a plant, an equivalent system was constructed. As discussed in Sec. 3.3, the identification routine used input and output time-domain data directly. Rather than compute (6.6) immediately then, the functional equivalent of Fig. 6.3 was used. The identified model of \tilde{P}^* was computed by filtering the time-domain input data, $d^*(t)$, through the previously-identified model of the filters, F^* , and using this as the input to the identification algorithm. This computed an equivalent of (6.6) without requiring a plant inversion.




Figure 6.3 Mixed-signal plant ID normalized

The second problem with (6.6) is the fundamental limitation that sampled-data systems depend on the periodic time at which they are sampled. In general, for any continuous-time systems, $P(s)$ and $Q(s)$:

$$(P(s)Q(s))^* \neq P^*(s)Q^*(s) \quad (6.7)$$

but rather, by applying (6.3)

$$(PQ)^* = \frac{1}{T} \sum_{k=-\infty}^{\infty} P(j\omega - jk\omega_s)Q(j\omega - jk\omega_s) \quad (6.8)$$

This illustrates the spectral aliasing of the continuous-time system in the discrete-time domain and is further addressed in Sec. 6.1.1.

The individual components of (6.6) were compared to their continuous-time counterparts. Since the filters are well-known, they are easy to illustrate and provide insight into the different

components of the signal aliasing. Figure 6.4 shows the relationship between a continuous-time set of anti-aliasing and smoothing filters and their discrete-time and aliased-terms on a large frequency scale. The $F(\omega)$ plot shows the combination of two 2nd-order Butterworth filters at 4000 Hz and a zero-order-hold at 10kHz. The sampled version of this continuous-time system is also shown with frequency content higher than the typically expressed half-sampling frequency rate.




Figure 6.4 Baseband and impulse-sampled hold and anti-aliasing filter

The baseband frequency aliasing error is of particular interest in the analysis of (6.6). From the definition of frequency sampling in (6.3),

$$\begin{aligned}
 F^*(j\omega) &= \frac{1}{T} \sum_{k=-\infty}^{\infty} F(j\omega - jk\omega_s) \\
 &= F(j\omega) + \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} F(j\omega - jk\omega_s) \\
 &= F[1 + \Delta_F(\omega)]
 \end{aligned} \tag{6.9}$$

where

$$\Delta_F(\omega) = \frac{F^*(\omega) - F(\omega)}{F(\omega)} \quad (6.10)$$

$$= \frac{\sum_{k \neq 0} F(j\omega - jk\omega_s)}{F(j\omega)} \quad (6.11)$$

and

$$|\Delta_F(\omega)| \leq \frac{\sum_{k \neq 0} |F(j\omega - jk\omega_s)|}{|F(j\omega)|}. \quad (6.12)$$

Figure 6.4 also shows this relative aliasing error as $F^* - F$. The minimum error is at $\omega = 0$, and is represented by a notch. Since it is an additive error, the ratio between the two, $\frac{F^*}{F}$ is also a useful plot to compare F^* with F for frequencies lower than the Nyquist rate.

6.1.1 Plant aliasing

With the aliasing error of the anti-aliasing filter addressed with (6.9)-(6.11), the additional component of (6.6) can be added. The same baseband-term subtraction was performed for $(FP)^*$ with (6.8):

$$(FP)^* = (FP)[1 + \Delta_{FP}] \quad (6.13)$$

where

$$\Delta_{FP}(\omega) = \frac{\sum_{k \neq 0} F(j\omega - jk\omega_s)P(j\omega - jk\omega_s)}{F(j\omega)P(j\omega)} \quad (6.14)$$

Using this, (6.6) can be written as:

$$\begin{aligned} \tilde{P}^* &= \frac{(FP)^*}{F^*} = \frac{FP(1 + \Delta_{FP})}{F(1 + \Delta_F)} \\ &= P \left[1 + \frac{\Delta_{FP} - \Delta_F}{1 + \Delta_F} \right] \end{aligned} \quad (6.15)$$

$$= P[1 + \Delta(\omega)] \quad (6.16)$$

The overall goal of the system identification was to determine an unfiltered model of the continuous-time system. The expression for \tilde{P}^* gives a calculable approximation for this, provided the uncertainty can be quantified. With the uncertainty term of Δ_{FP} embedded within (6.16), however, this cannot be assured for arbitrary plant dynamics. This makes sense, since

even with high-order anti-aliasing filters, a continuous-time plant model can be found with sufficient high-frequency gain to introduce arbitrary aliasing errors.

Some of the terms in the $\Delta(\omega)$ expression in (6.16) hide the error relationships in Δ -type relative error terms. To help illustrate the relationship in terms of the sampled and baseband transfer functions, an alternate expression for \tilde{P}^* is shown in (6.17).

$$\begin{aligned}\tilde{P}^* &= P - P + \frac{(FP)^*}{F^*} \\ &= P + \frac{(FP)^* - F^*P}{F^*}\end{aligned}\tag{6.17}$$

For illustration purposes, the 10kHz-sampled model was considered the “true” model, and then “resampled” at 5kHz. Figure 6.5 shows the approximation of (6.17) graphically. This clearly shows the baseband-notch as in the simple filter of Fig. 6.4, but with the effect of the plant shown as well. The notch indicates very little aliasing error at low frequencies, with the actual amount dependent on the plant itself. The aliasing error term becomes significant towards the Nyquist rate of the down-sampled plant, but by an amount determined by the “unmeasurable” out-of-band plant response. This plot shows the need to limit the gain excursions on the aliased high-frequencies in order to put a bound on the aliasing error.

To analytically limit the possibility of arbitrary aliasing errors, a reasonable assumption was made about the continuous-time plant. Rather than assume the very likely case of high-frequency roll-off, a more general limitation on the maximum gain excursions was considered. In particular, we assume that there exists an α constant such that

$$|P(j\omega - jk\omega_s)|_{k \neq 0} \leq \alpha |P(j\omega)|\tag{6.18}$$

for all integers $k \neq 0$. This limitation effectively limits the high-frequency gain of the continuous-time system to be no more than a scalar factor of the baseband plant. It is actually a *worst-case* limitation since it assumes:

- At every frequency in the baseband plant, every aliased frequency has magnitude of α .
- No phase cancellations in the aliased frequencies are present, and thus add the maximum amount of error possible to the modeled baseband plant.

The relationship between the baseband plant and the α bound can be seen in Fig. 6.6. For purposes of illustration, the discrete-time 10kHz sampling model of was used as the “continuous-




Figure 6.5 Absolute aliased plant residue

time” model, but with a “sampling rate” of 5kHz. For the largest peak in the baseband plant, the various corresponding α bounds are shown.

The application of the assumed α bound on the continuous-time plant makes it possible to investigate the aliasing uncertainty in the plant model. Consider the Δ_{FP} term in (6.15):

$$\begin{aligned}
 |\Delta_{FP}| &= \frac{\left| \sum_{k \neq 0} F(j\omega - jk\omega_s)P(j\omega - jk\omega_s) \right|}{|F(j\omega)P(j\omega)|} \\
 &\leq \frac{\sum_{k \neq 0} |F(j\omega - jk\omega_s)P(j\omega - jk\omega_s)|}{|F(j\omega)| |P(j\omega)|} \\
 &\leq \frac{\alpha |P(j\omega)| \sum_{k \neq 0} |F(j\omega - jk\omega_s)|}{|F(j\omega)| |P(j\omega)|} \\
 &\leq \alpha |\Delta_F|
 \end{aligned} \tag{6.19}$$




Figure 6.6 Illustration of different α on a sampled plant

Finally, the uncertainty in (6.16) can be bounded with (6.12) and (6.19):

$$\begin{aligned}
 |\Delta| &= \left| \frac{\Delta_{FP} - \Delta_F}{1 + \Delta_F} \right| \\
 &\leq \left| \frac{\Delta_{FP}}{1 + \Delta_F} \right| + \left| \frac{\Delta_F}{1 + \Delta_F} \right| \\
 &\leq \frac{|\Delta_F| (1 + \alpha)}{|1 + \Delta_F|} \tag{6.20}
 \end{aligned}$$

Thus, the maximum aliasing error in assuming that $P \approx \tilde{P} = \frac{(FP)^*}{F^*} = P(1 + \Delta)$ can be expressed as a simple multiplicative uncertainty. The magnitude of the uncertainty is shown in Fig. 6.7 with varying α bounds. Sample values of α were used just as in the illustrative example of Fig. 6.6, but with the 10kHz sampling rate. Also shown is the error in the anti-aliasing filter, $|\Delta_F|$. What this plot shows is the relationship between the simple aliased filter error, and the error bound under the α assumption of (6.18). For a given α limit however, an anti-aliasing filter can be chosen that reduces the plant's aliased error to less than a specified amount.

Since the error term is multiplicative, expressing the magnitude in relation to the unity factor in (6.16) is a more intuitive quantity. The same curves as in Fig. 6.7 are plotted in this




Figure 6.7 Upper bound on aliasing relative error, $\Delta(\omega)$

manner in Fig. 6.8. This shows that the magnitude error in the plant model can be quantified to be less than a few dB for a given α bound up to a given frequency.

An important note on this uncertainty is recalling that it is a *worst-case* error, bounding high-frequency harmonics below a scalar factor of the corresponding baseband plant model. In actual plants, the relationship is rarely as severe as this. The largest effect of these errors would be seen where the plant has a baseband zero and corresponding aliased resonances. Since the overall goal of the control algorithm was to reduce the resonant peaks, however, the baseband zeros are not as important to model accurately as the resonant peaks. A less conservative α -bound could be chosen in this case.

6.1.2 Plant aliasing with feedback

A comparable amount of care must be taken when considering the analysis of the mixed-signal plant with continuous-time feedback. Figure 6.9 shows the system with the addition of the continuous-time feedback controller. The actual test system with anti-aliasing filters as




Figure 6.8 Upper bound on continuous-time plant model error due to aliasing

present is Fig. 6.9(a), but for analysis, the input-output equivalent of Fig. 6.9(b) was used where $F = H_z B_s B_a A$ as before.

Solving the resulting block diagram to obtain a mixed-signal transfer function relationship yields:

$$\frac{E_f^*(s)}{D^*(s)} = \left(\frac{FP}{1 - PAC} \right)^* \quad (6.21)$$


To make this analysis similar to that of Sec. 6.1.1, there is a desire to isolate the continuous-time portion of the measured discrete-time transfer function. Considering the pre-filtered ID routine as before, this expression can be written as:

$$\tilde{P}_a^* = \frac{\left(\frac{FP}{1 - PAC} \right)^*}{F^*} \quad (6.22)$$


$$= \frac{(FP_a)^*}{F^*} \quad (6.23)$$

where

$$P_a = \frac{P}{1 - PAC} \quad (6.24)$$



(a) Actual test system



(b) I/O equivalent

Figure 6.9 Mixed-signal plant ID setup with analog feedback

This agrees with the intuitive concept that the analog controller simply modifies the plant completely in the continuous-time domain. For consistency, a bound on this expression using the same α bound on the open-loop plant was desired. An attempt to do the same manipulation in (6.17) yields:

$$\tilde{P}_a^* = \frac{P}{1 - PAC} + \frac{\left(\frac{FP}{1 - PAC}\right)^* (1 - PAC) - F^*P}{F^*(1 - PAC)} \quad (6.25)$$

Unfortunately the error term of (6.25) does not provide the same insight into the sampling relationship as did (6.17). It now contains additional terms that are dependent on not only the plant, but also the specific feedback control addition. A relative error expression can still provide a bound on the error analogous to (6.15)

$$\tilde{P}_a^* = P_a \left(1 + \frac{\Delta_{FP_a} - \Delta_F}{1 + \Delta_F} \right) \quad (6.26)$$

$$= P_a (1 + \Delta_a(\omega)) \quad (6.27)$$

A similar bound on the Δ_a magnitude can be expressed in terms of $|\Delta_{FP_a}|$, once its relationship to the α -bound can be established. Writing

$$\begin{aligned} |\Delta_{FP_a}| &= \left| \frac{\sum_{k \neq 0} \frac{F(j\omega - jk\omega_s)P(j\omega - jk\omega_s)}{1 - P(j\omega - jk\omega_s)AC(j\omega - jk\omega_s)}}{\frac{F(j\omega)P(j\omega)}{1 - P(j\omega)AC(j\omega)}} \right| \\ &\leq \frac{\alpha_a |P_a(j\omega)| \sum_{k \neq 0} |F(j\omega - jk\omega_s)|}{|F(j\omega)P_a(j\omega)|} \\ &\leq \alpha_a \frac{\sum_{k \neq 0} |F(j\omega - jk\omega_s)|}{|F(j\omega)|} \end{aligned} \quad (6.28)$$

$$\leq \alpha_a |\Delta_F| \quad (6.29)$$

yields a bound on $|\Delta_{FP_a}|$ if there exists a bound α_a such that

$$\left| \frac{P(j\omega - jk\omega_s)}{1 - P(j\omega - jk\omega_s)AC(j\omega - jk\omega_s)} \right|_{k \neq 0} \leq \alpha_a \frac{|P(j\omega)|}{|1 - P(j\omega)AC(j\omega)|}. \quad (6.30)$$

The inequality of (6.30) represents the same type of bound as in the open-loop plant in (6.18), except for the possibility of increased gain due to the continuous-time controller, $C(j\omega)$. The left-hand side of (6.30) represents gain of all of the non-baseband frequencies. The goal of the analog controller was to reduce the closed-loop plant gain, so the only increase in these high-frequency gains would be due to high-frequency spillover from the analog controller. The performance characteristics of the feedback controller were investigated in Chap. 5, where the worst-case phase-margin was a design parameter. This phase margin provides a limit to how small the denominator of (6.30) can become. Additionally, the numerator is limited in size by the α -bound from the previous section. Thus, assuming a complex-valued PAC , an acceptable α_a bound can be written as

$$\begin{aligned} \alpha_a &\leq \frac{\alpha}{|1 - P(j\omega_{pm})AC(j\omega_{pm})|} \\ &\leq \frac{\alpha}{\sqrt{(1 - \cos(\phi_{pm}))^2 + \sin^2(\phi_{pm})}} \end{aligned} \quad (6.31)$$

A graphical interpretation of this limit is shown in Fig. 6.10. Since the same type of aliasing error magnitude bound is present as in the open-loop plant, the closed-loop plant error can be written just as (6.16), but with the updated α .

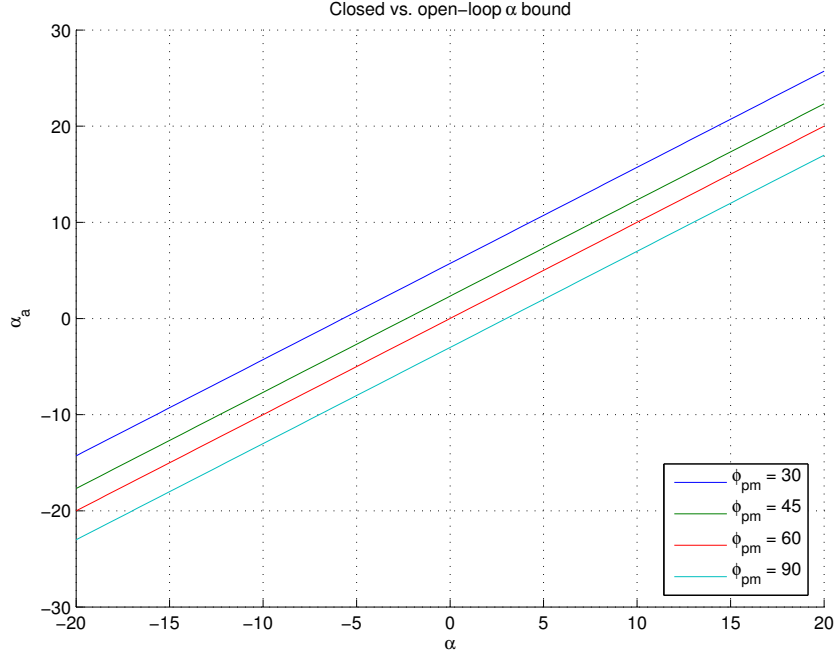


Figure 6.10 Relationship of α to α_a at different phase margins

Finally, the uncertainty in (6.27) can be bounded with (6.12) and (6.29) to form (6.32) where the α_a closed-loop limit can be expressed in terms of the open-loop limit through (6.31).

$$\begin{aligned}
 |\Delta_a| &= \left| \frac{\Delta_F P_a - \Delta_F}{1 + \Delta_F} \right| \\
 &\leq \left| \frac{\Delta_F P_a}{1 + \Delta_F} \right| + \left| \frac{\Delta_F}{1 + \Delta_F} \right| \\
 &\leq \frac{|\Delta_F| (1 + \alpha_a)}{|1 + \Delta_F|} \tag{6.32}
 \end{aligned}$$

So, by making a rather modest scalar bound on the maximum out-of-band magnitude as compared to the in-band magnitude, the aliasing error can be bounded. In addition, this sampling error bound can be adjusted to include any out-of-band spillover introduced by a continuous-time controller. Thus, the approximation of $\tilde{P}_a^* = P_a (1 + \Delta_a(\omega))$ can be assured within a specified error bound. This final form was included in the error analysis of the closed-loop controller discussed in Sec. 6.5.

6.2 Discrete-time Rate Changing

In addition to the aliasing introduced by the discrete-time sampling of the continuous-time system, an additional source of potential aliasing errors was introduced by the decimation of the low-frequency controller. Most control literature does not address such sampling rate changes, but the signal processing literature utilizes it frequently. The primary reference used for this analysis is [115], which uses the complex exponential frequency-domain representation of the Fourier transform. This is a bit different from the pole/zero notation typically used in the control literature, and generally assumes that a frequency response is the desired result.

This section looks at the modeling errors introduced with the decimation and interpolation of the controller implementation. In the end, the operations were very similar to that of the sampling of a continuous-time system as investigated in Sec. 6.1. The aliasing introduced nonlinear errors because of frequency re-mapping. The errors associated with the implementation of direct discrete-to-discrete remapping by Matlab's `d2d` function were quantified and provided justification for its use in the analysis. Section 6.2.1 quantifies the errors associated with decimation, and Sec. 6.2.2 those attributed to interpolation.

6.2.1 Decimation

Decimation of the high-frequency sampled data was the first operation performed for the low-frequency implementation. The effects of this on the analytical model were quantified by considering the digital anti-aliasing filters and errors associated with the frequency remapping.

From the definition of the discrete-time Fourier transform, a discrete time signal $x[n] = x_c(nT)$ has the frequency response

$$X(e^{j\omega}) = \frac{1}{T} \sum_{k=-\infty}^{\infty} X_c \left(j \frac{\omega}{T} - j \frac{2\pi k}{T} \right). \quad (6.33)$$

The same continuous-time system, $x_c(t)$, sampled at a integer lower sampling rate, $T' = MT$ has frequency response

$$X_d(e^{j\omega}) = \frac{1}{MT} \sum_{r=-\infty}^{\infty} X_c \left(j \frac{\omega}{MT} - j \frac{2\pi r}{MT} \right). \quad (6.34)$$

With a bit of algebra, and utilizing the definition of the discrete-time Fourier transform, it can be shown that the relation between $X(e^{j\omega})$ and the decimated $X_d(e^{j\omega})$ is (6.35).

$$X_d(e^{j\omega}) = \frac{1}{M} \sum_{i=0}^{M-1} X \left(e^{j(\omega/M - 2\pi i/M)} \right) \quad (6.35)$$

Ignoring the decimation aliasing terms of (6.36), the approximation of the decimation transfer function (6.37) results.

$$= \frac{1}{M} X(e^{j\omega/M}) + \frac{1}{M} \sum_{i=1}^{M-1} X \left(e^{j(\omega/M - 2\pi i/M)} \right) \quad (6.36)$$

$$\approx \frac{1}{M} X \left((e^{j\omega})^{1/M} \right). \quad (6.37)$$

For a high-order system such as the one under examination in this work, the classical representation of a discrete-time system as a transfer function becomes numerically sensitive. From numerous investigations with Matlab on the identified system model, rational expressions of high-order polynomials caused numerical errors with over a few tens of poles and zeros. For this reason, a pole/zero/gain model was used where possible to reduce the system order. This had the effect of reducing the high-order rational expressions into equivalent “1st-order” sections and decreasing the numerical sensitivity. Such a system can be represented as (6.38) with N_p poles $\{p_{i_p}\}$, and N_z zeros, $\{z_{i_z}\}$. Applying the decimation equation (6.36) to (6.38), the full aliased decimated system’s frequency response can be represented by (6.39).

$$X(z) = K \frac{\prod_{i_z=1}^{N_z} (z - z_{i_z})}{\prod_{i_p=1}^{N_p} (z - p_{i_p})} \quad (6.38)$$

$$X_d(e^{j\omega}) = \frac{K}{M} \frac{\prod_{i_z=1}^{N_z} (e^{j\omega/M} - z_{i_z})}{\prod_{i_p=1}^{N_p} (e^{j\omega/M} - p_{i_p})} + \frac{K}{M} \sum_{i=1}^{M-1} \left(\frac{\prod_{i_z=1}^{N_z} (e^{j(\omega - j2\pi i)/M} - z_{i_z})}{\prod_{i_p=1}^{N_p} (e^{j(\omega - j2\pi i)/M} - p_{i_p})} \right) \quad (6.39)$$

Although the first term of (6.39) represents the fundamental non-aliased decimated plant, it cannot be easily separated from the aliased components in the second term. The rational expression requires tedious algebra and numerically sensitive computations to compute the modification of the poles and zeros of the fundamental plant by the aliasing terms exactly. To minimize these modifications to the baseband plant, a discrete-time anti-aliasing filter would be used in practice.

In addition, the decimation of the discrete-time plant does not lend itself to a unique model transformation from one sampling rate to the other. The transformation is not unique because it must assume some type of hold between the sampled data points and tends to be modeled as a conversion from discrete-time to continuous-time and back [118]. This type of mapping can generally map the poles of a system uniquely between continuous-time and discrete time, but not the zeros [119]. Matlab’s **d2d** command provides one possible solution to this with the logarithm of the state transition matrix but it has also been done by fitting time-domain response data [119].

The accuracy of the approach used by Matlab is of particular interest. The frequency-domain response of a decimated plant provides one means to ascertain the accuracy of the model. The basic decimator block is shown in Fig. 6.2.1, and assumes the presence of an anti-aliasing filter. The frequency response of the different models of Fig. 6.11(b) produced by Matlab is shown in Fig. 6.12. The baseband plant is shown for reference, as well as the baseband plant with the digital anti-aliasing filter. The decimation rate was chosen as a factor of 8 from the 10kHz sampling rate of the analysis ID model. The “ Σ FRD” contains the decimated plant’s frequency-response data, with all of the aliased terms of (6.36) added in. Also shown is the version of the decimated plant’s frequency response as computed by Matlab. It differs from the ideal decimated frequency response in both magnitude and phase. This can be explained by the assumption made by the **freqresp** function used to compute the frequency response. This function assumes an implicit ZOH on the output in order to obtain the continuous-time frequency response. When the frequency response of a ZOH is removed from the result, the plot is shown in magenta. With that removed, it is almost identical to the **d2d** function response with digital anti-aliasing filter.




Figure 6.11 Decimation analysis setup

The additional phase delay of the hold was a concern from the analysis point of view, but it is only an artifact of the tools used to present the frequency response. The actual discrete-time




Figure 6.12 Decimation aliasing errors

model was investigated further to ensure that it did not contain any additional unmodeled phase lag. The conclusion from the decimation analysis was that although the problem is inherently lossy, the errors in the decimated model can be kept to a quantifiable minimum.

6.2.2 Interpolation

A similar analysis was performed for the interpolation errors in the control design. Referencing [115] again, the system's frequency response $X(e^{j\omega})$ up-sampled by L has the frequency response in (6.40).

$$\begin{aligned}
 X_i(e^{j\omega}) &= \sum_{n=-\infty}^{\infty} \left(\sum_{k=-\infty}^{\infty} x[k] \delta[n - kL] \right) e^{j\omega n} \\
 &= X(e^{j\omega L})
 \end{aligned} \tag{6.40}$$

This in effect compresses the frequency response of the discrete-time system. As a result, the baseband frequency responses get aliased L times at higher frequencies and must be suppressed with a smoothing filter.

Figure 6.13 shows the resulting block diagram from the total plant manipulation. Again, it was useful to compare the true frequency response characteristics of the interpolator to that produced by the solution presented by Matlab’s **d2d** function. Figure 6.14 shows the cascade of the functions. Again, the original plant model is shown for reference. The original plant with the frequency response of both anti-aliasing filters is also shown as a reference. The “ Σ FRD” contains the actual frequency response data performed as calculated by the cascade of (6.36) and (6.40). The final plot is the cascade of the **d2d** from the original rate, to the decimated rate, and back, with the filters as arranged in Fig. 6.13. The most interesting point to note is that the Matlab frequency response was definitely performing a model reduction compared to the true frequency response. This was not mentioned in [103], but was likely to be practically required anyway. Interpolation of an N th-order, decimated system by a factor of L would inherently contain at least $L \times N$ states. The necessary smoothing filter would effectively eliminate all but N , however. Numerical problems plagued the analysis of this system throughout, so keeping unmeasurable states was certainly not desirable.



Figure 6.13 Decimation/interpolation of plant

The final piece to consider in this analysis is how much of an error is present in the aliasing of the decimated plant. Figure 6.15 shows the same data as Figs. 6.12 and 6.14, but relative to the true plant’s response. All are normalized to the cascade of the original plant to the two (identical) anti-aliasing filters. As expected, the relative error is very small, at less than a few dB almost up to the decimated Nyquist rate.

The error analysis results summarized in Fig. 6.15 meant that for analysis purposes, the **d2d** function within Matlab did not introduce significant errors as compared to the true decimated frequency response. This was important, because as explained in [118], there is no single, unified way to compare systems of different sampling rates. Likewise, Fig. 6.14 shows a negligible increase in modeling uncertainty due to interpolation. The aggregate decimation and interpolation present errors to the model in addition to adding the required digital anti-aliasing and smoothing filters shown in Fig. 6.15. These additional errors still show negligible contributed errors up to the control frequency of interest. Even without additional errors, the




Figure 6.14 Interpolation errors

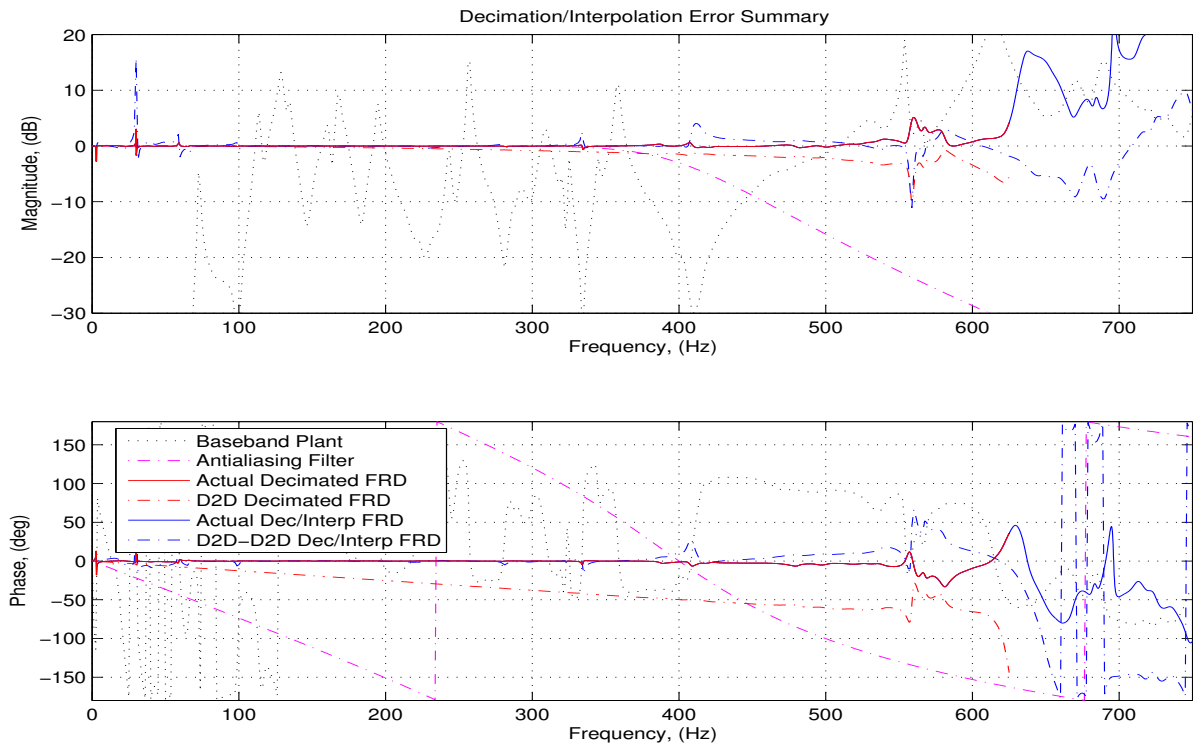


Figure 6.15 Decimation and interpolation error summary

filters and holds themselves must still be present in the loop-gain analysis in order to truly model the experimental implementation. This was an important distinction to make. Although the approximated “equivalent” system representation for the decimated/interpolated system introduced little error, the *required* additions of filters and holds still had to be included. The small effect of the error to the loop gain is included in the uncertainty analysis of Sec. 6.5.

6.3 LMS Optimal Solution Computation

At this point in the analysis, it was necessary to consider the inclusion of the feed-forward controller. The transient performance of the FxLMS algorithm used was not directly under investigation. What *was* important was the optimal solution to the LMS algorithm given all of the modeling difficulties of the high-order, hybrid system. This section describes the generalized technique used to literally implement the optimal LMS solution utilizing a time-domain simulation of the system. The time-domain simulation included time delays, holds, and decimation/interpolation that were difficult to model analytically. This time-domain simulation was used to estimate the parameters necessary to compute the optimal filter coefficients.

The formulation of the LMS algorithm has taken many different permutations [84] [89] [94]. These have been briefly introduced in Sec. 2.4, but will be explained in more detail here for the case of the filtered-x, LMS. In particular, the optimal solution of the algorithm given a deterministic system model will be computed.

The general estimation problem that LMS tries to solve is shown in Fig. 6.16. Phrased this way, the “desired” signal is the contribution of the disturbance through the primary disturbance path. The optimization problem normally considers the more general case of a stochastic noise source along with the signal. For the reasons mentioned in Sec. 6.4.1 however, it is useful here even with deterministic models. The optimal solution was calculated using a deterministic simulation of the exact FxLMS algorithm in order to replicate the experimental setup. The model would eventually include all filters, samplers, holds, and the other controllers that were present in the experimental setup, but initially only the simple case is considered here.

The framework of the discrete-time form of the Wiener-Hopf equation is discussed in more detail in [89], but the results require the construction of time-domain sequences and matrices since all the models used in this investigation were IIR (discrete-time state-space, specifically).




Figure 6.16 General estimation filter

The optimization problem assumes an FIR filter weighting past values of the reference and has the form of (6.41) for a minimum mean-squared error signal.

$$0 = r_{fd}(k) - \sum_{i=0}^N c_i^o r_{ff}(k-i) \quad 0 \leq k \leq N-1 \quad (6.41)$$

The optimal filter coefficients for an N -length filter are given by c_i^o , where r_{fd} is the digital cross-correlation between the reference and the desired signals, and R_{ff} is the digital autocorrelation matrix of the reference signal. Recalling the definition of correlation of two discrete-time signals is (6.42), the definition of the autocorrelation matrix is shown in (6.43).

$$r_{xy}(m) = E[x(n)d(n+m)] \quad (6.42)$$

$$R_{ff} = \begin{bmatrix} r_{ff}(0) & r_{ff}(1) & r_{ff}(2) & \dots & r_{ff}(N-1) \\ r_{ff}(1) & r_{ff}(0) & r_{ff}(1) & \dots & r_{ff}(N-2) \\ \vdots & \vdots & \vdots & & \vdots \\ r_{ff}(N-1) & r_{ff}(N-2) & r_{ff}(N-3) & \dots & r_{ff}(0) \end{bmatrix} \quad (6.43)$$

Solving (6.41) for the filter yields (6.44) with the optimal FIR filter given as vector \vec{c}^o .

$$\vec{c}^o = R_{ff}^{-1} \vec{r}_{fd} \quad (6.44)$$

Using the identified model of the system, it was excited in simulation with a white-noise reference. The white-noise ensured adequate excitation of the model to provide the necessary system response. The algorithm was the unmodified filtered-x LMS so the “reference” signal is filtered by the secondary path model, P . The “desired” signal is filtered by the disturbance path model, P_d , as in Fig. 6.17. These two time-domain signals were easily modified to contain any additional filtering, shaping or delay elements as would eventually be added in Sec. 6.4.2-6.4.3.

The autocorrelation matrix and cross-correlation vector were computed with this filtered time data using (6.42)-(6.43). Using (6.44), the optimal filter coefficients were computed for the




Figure 6.17 Estimation filter with actuator dynamics

given system model. As required by [89], the white-noise reference signal provided spectrally rich content to ensure the invertibility of R_{ff}^{-1} .

Another advantage of this method is that it allows for realistic sampling and decimation of the plant for the feed-forward controller. The baseline plant model used for the analysis was sampled at 10kHz. Once the control algorithm was running, however, the sample-rate would be lowered to reduce computation. The time-domain generation of the optimal filter coefficients allows this exact same process to be utilized in computing the optimal filter coefficients. The time-domain response of the analysis model was computed on the higher order, high-sample-rate model which better approximates the real system's continuous-time dynamics.

Figure 6.18 shows an example of the continuous-time frequency response used to compute a FxLMS optimal filter using a representative model of the structure. Since this was merely a representative example, including all of the filters, decimators, and interpolators to the computation would have only complicated the setup and reduced the clarity of the example. Rather, the samplerate of the controller was set at 5kHz, with the high-order identified model at the full 10kHz sample-rate. The actuator path response (FxPmodel) shows the effective model of the actuator path as was seen by the FxLMS algorithm. It obviously shows high-order dynamics just as the desired disturbance path model (FxPdmodel), and certainly significant dynamics through the target sampling rate of 5kHz.

Figure 6.19(a) shows a pseudo-color rendition of the autocorrelation matrix and 6.19(b) its inverse. While not necessarily completely intuitive, they provide a visual means to examine the “difficulty” of the filter fit and the ringing in the structural impulse response. Models requiring longer filter lengths can be seen in the complexity of the R_{ff}^{-1} matrix. As can be seen from Fig. 6.19(c), the optimal solution has very aggressive high-frequency response, and is not completely






Figure 6.18 No filter FxLMS input frequency response

captured with a 256-point impulse response, although it *is* reduced by a factor of 10 by the end of the filter length.

The computation of the autocorrelation matrix required averaging of the time-domain simulation data, so an attempt was made to quantify the effect of changing the number of components in the average. The true autocorrelation matrix (6.43) was numerically computed by approximating the expected value of (6.42) as the mean over a number of non-overlapping sequences of time-domain simulations data. Although R_{ff} is only *approximated* by this numerical computation, a relatively low number of averages were required to provide adequate convergence. Lengths ranging from 2-1600 were tried, but anything over 25 averages had little effect. Figure 6.19(d) shows the effect of varying the length of the time-domain simulation used to obtain the optimal filter. It shows very little difference in the result between 5 different numbers of averages ranging from 25 to 1600. This is expected, since the underlying model is deterministic only, but it reassures that there weren't numerical problems in the computation of the average. Only the length 25 was used for subsequent estimations of the autocorrelation matrix.




(a) R_{ff}



(b) R_{ff}^{-1}

131



(c) $\vec{c}^{\vec{o}}$



(d) Filter response (25-1600 averages)

Figure 6.19 Optimal filter computations

Finally, Fig. 6.20 shows the frequency response of the optimal unweighted FxLMS filter in both 5KHz sample-rate and continuous time. The black line represents the frequency response of the FIR filter generated by the optimal filter computation. Due to the addition of holds and smoothing filters, the actual frequency response of the “continuous-time” implementation of the FIR filter is shown in blue. Also, this allows for a more consistent comparison with the error-weighting strategies employed in Sec. 6.4.2-6.4.3. The output of the feed-forward controller was subject to a zero-order hold at 5kHz and 4kHz smoothing filters. The resulting aliased components around 2.5 kHz were reduced considerably (20dB) by 2.6kHz.




Figure 6.20 No filter FxLMS output filters

The primary benefit of this method of optimal filter computation is that it allows for accurate capturing of the errors caused by sampling. By separating the *model* samplerate from the *controller* samplerate, variations on the controller parameters could be easily adjusted. The issues of different types of holds, anti-aliasing filters, and the small residual aliasing terms could be included in the analysis. It also allowed the integration of the high-frequency analog controller’s affect to be included in the model that would otherwise be out of band for the

model used in the analysis. A large portion of the usable bandwidth of the analog controller is completely above the Nyquist rate for the feed-forward controller.

The performance of the unweighted FxLMS optimal filter can be seen in 6.21. Although it produces respectable reduction throughout the entire bandwidth, it clearly does not have enough degrees of freedom at only 256-points to adequately model the approximately 250th-order state-space model of the plant. Because of the synthesis of the controller with three independent frequency bands, error-weighting was desired to shift the optimal filter into the desired bandwidth. This would allow the limited degrees of freedom of the filter to more adequately fit the bandwidth of interest. This is considered in Sec. 6.4.2-6.4.3




Figure 6.21 Broadband FxLMS performance

6.4 Total System Modeling

The analysis of the combination of the controllers required a careful study of all of the individual components of the different control types. As mentioned in Sec. 6.1, the usual assumptions of high sampling rates, low phase-rolloff, and minimal controller interaction were

not necessarily valid. The system identification error modeling of Sec. 6.1.1 and 6.1.2 was extended to incorporate the feed-forward and low-frequency feedback controllers.

6.4.1 Model extension for feed-forward

The feed-forward controller application implied that at least some of the source of the vibration came from a measurable disturbance. While technically possible to run a feed-forward controller in the continuous-time domain, practically speaking it requires a digital controller for any system with nontrivial dynamics. From the digital controller's perspective then, the analog control becomes a transparent modification to the continuous-time plant. For an individual actuator/error sensor pair, the disturbance can be modeled as an uncoupled system injecting disturbance energy at the error sensor. When the analog feedback controller is added, the resulting system (Fig. 6.22(a)) becomes coupled only through the error summing junction. This cross-coupling can be easily shown in (6.45), or conceptually by moving the feedback through the summing junction in the block diagram in Fig. 6.22(b).

$$\begin{aligned} E &= (PA)D + (P_w)W + (PAC)E \\ &= \left(\frac{PA}{1 - PAC} \right) D + \left(\frac{P_w}{1 - PAC} \right) W \end{aligned} \quad (6.45)$$

By moving the feedback through the summing junction, the collocated control's effect on the disturbance path becomes more clear. The goal of the controller is to reduce the effect of the disturbance in the error signal. The only way that this is done is by increasing the loop gain PAC . The addition of the $\frac{1}{1-PAC}$ term, provides decreased disturbance throughput when the loop gain is high.

In order to apply the discrete-time feed-forward control for the mid-frequency, this continuous-time plant must be sampled. At this point, it was necessary to invoke the hybrid continuous/discrete analysis similar to that used in Sec. 6.1.1. The block diagram of the feed-forward system in combination with the analog controller is shown in Fig. 6.23. The C_{ff}^* term represents the discrete-time controller obtained using the LMS algorithm, so it is present in the sampled-data portion of the signal path. Also represented in the discrete-time domain is the d^* signal, which represents a control input to be used for the low-frequency feedback control. The continuous-time filtering elements are grouped into the $F(s)$ transfer function as in Sec. 6.1.1. The $F(s)$ equation groups together the zero-order hold, smoothing filter, analog amplifier

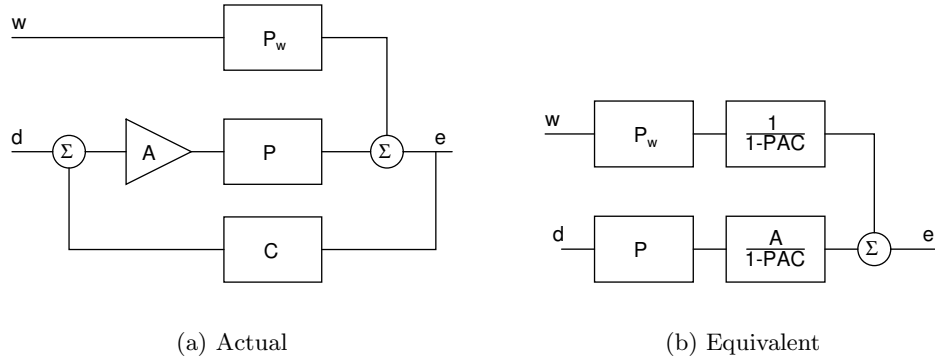


Figure 6.22 Analog actuator and disturbance model

gain, and anti-aliasing filter into a single continuous-time transfer function for the secondary (control) path, (6.46), and the primary disturbance path, (6.47).

$$F(s) = H_z B_s A B_a \quad (6.46)$$

$$F_w(s) = H_z B_s B_a \quad (6.47)$$

In reality, the anti-aliasing filter, B_a , exists on the error signal after the summing junction but before the sampler to e_f^* . Since this filter is still in the continuous-time domain and the models are linear, it can be brought through the summing junction.




Figure 6.23 Hybrid continuous-discrete feed-forward model.

Writing the output error in terms of the mixed-signal input disturbance, w , and actuator control signal, d , yields (6.48). Sampling the system of (6.48) using the “*”-operator distribution properties in [104] results in an accurate representation of the frequency response in the discrete-time domain in (6.49).

$$E_f = [F_w P_w + C_{ff}^* F P] \left(\frac{1}{1 - PAC} \right) W^* + [F P] \left(\frac{1}{1 - PAC} \right) D^* \quad (6.48)$$

$$E_f^* = \left[\frac{F_w P_w + C_{ff}^* F P}{1 - PAC} \right]^* W^* + \left[\frac{F P}{1 - PAC} \right]^* D^* \quad (6.49)$$

It is obvious from (6.49) that the desired disturbance-minimizing transfer function would be $C_{ff}^{*o} = \frac{-F_w P_w}{F P}$. Unfortunately, there are a few problems with this expression. First, the simple expression for direct disturbance cancellation does not provide any means to weight the frequency response. Model reduction with frequency weighting terms could potentially be done, but this assumes that a numerically stable full-order model could be obtained. Such methods were attempted by re-identifying the system *through* the $(FP)^*$ path in a manner similar to Fig. 6.3. Unfortunately, the system order necessary to obtain sufficient accuracy within the band of interest led to numerical computation problems with this approach.

Second, the intended operation of the feed-forward control is through the LMS algorithm. Although it has very attractive adaptive performance, computation, and convergence qualities it yields a finite-impulse-response (zero-only) controller. It therefore cannot exactly equal the recursive model $\frac{-F_w P_w}{F P}$. Moreover, the plant has such a lightly-damped characteristic that it has an extremely long impulse response. This requires very long filter lengths for an FIR filter to adequately approximate and was the motivation for the frequency-separation control system topology to begin with.

Third, it assumes that P^{-1} can be obtained. The identified model of P was not minimum-phase, so its inverse was not stable. Attempting to isolate the minimum-phase and non-minimum-phase components proved to be too numerically sensitive as well. The required order of the plant model was in excess of 50, and attempting any computations requiring P^{-1} either did not converge, or produced unstable results. One technique attempted to reduce the non-minimum phase model was to employ the zero conjugate-reciprocal mapping presented in [115]. Both the P and P_w SISO models were derived from the same identified state-space model, so they had identical poles. The ratio of $\frac{P_w}{P}$ can then be expressed as (6.50).

$$\frac{P_w}{P} = \frac{\prod_{i=1}^{N_{z_w}} (z - z_{w_i})}{\prod_{i=1}^{N_z} (z - z_i)} \quad (6.50)$$

By transforming the zeros of P by the conjugate reciprocal mapping in [115], the magnitude response could be maintained while preserving P^{-1} 's stability.

Figure 6.24 shows the frequency response of the resulting minimum-phase model of P . The magnitude response is correct, but the phase is not matched. This result was typical of applying signal-processing techniques for use in real-time control. Often phase or causality requirements are “adjusted” with shifting to allow for a correct magnitude response, but the result is unusable

for control. For use in the feed-forward controller, this signal-processing tool could not provide adequate performance with such mismatching phase. For this reason, the analysis of the feed-forward was left to the more literal implementation of the optimal solution presented in [89] in Section 6.3.




Figure 6.24 Minimum-phase representation of P

6.4.2 Filtered-error LMS

The basic FxLMS algorithm did not provide enough design parameters for the controller synthesis. In particular, the lack of frequency-weighting in the optimal solution of the standard FxLMS algorithm was mentioned in 6.4.1. Figure 6.25 shows the most intuitive form of error-weighting for the feed-forward controller. It assumes a shaping filter on the error signal, but does not directly appear in the signal flow path. The signal e_f^* represents the anti-alias-filtered error signal, with the smoothing, zero-order hold, and anti-aliasing filter represented by the F and F_w blocks. Also note that the analog controller was included into the actuator and

disturbance paths, which are denoted by (6.51) and (6.52), respectively.

$$P_a = \frac{P}{1 - PAC} \quad (6.51)$$

$$P_{w_a} = \frac{P_w}{1 - PAC} \quad (6.52)$$

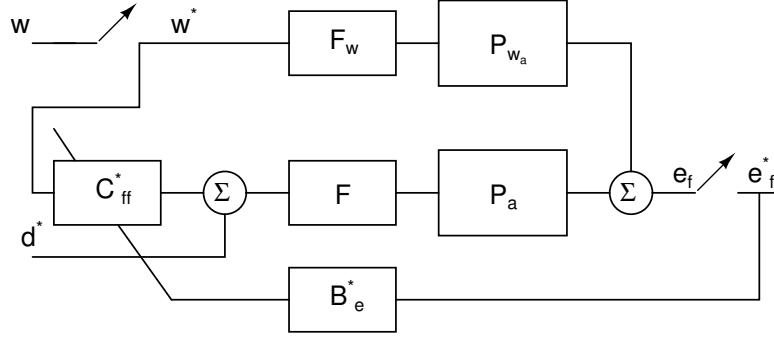


Figure 6.25 FxLMS error-weighting

Unfortunately, the representation of Fig. 6.25 does not fit into the standard Wiener-Hopf formulation [89]. By applying the block-diagram manipulation rules in [104], the block-diagram equivalent system of Fig. 6.26 could be formulated. By adding and removing the filter block in the error signal path, the error signal of (6.53) was modified.

$$e_f^* = [(F_w P_{w_a})^* + C_{ff}^* (F P_a)^*] w^* + [(F P_a)^*] d^* \quad (6.53)$$

The advantage of this equivalent representation is that the secondary disturbance path now includes the error-shaping filter, B_e^* .

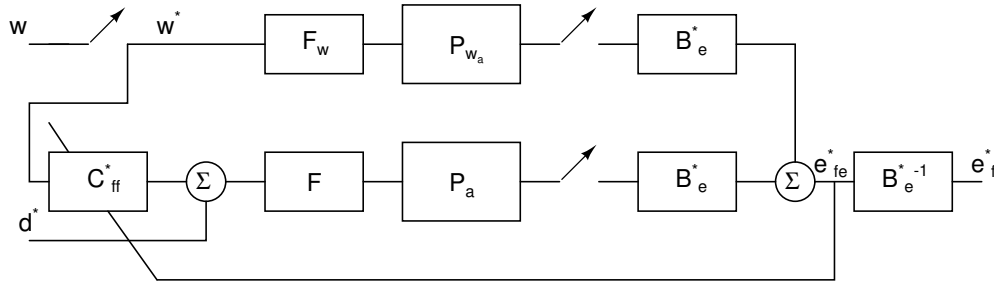


Figure 6.26 FxLMS error-weighting equivalent

Equation (6.54) shows the addition of the error-shaping filter in both paths.

$$e_{fe}^* = [(F_w P_{w_a})^* B_e^* + C_{ff}^* (F P_a)^* B_e^*] w^* + [(F P_a)^* B_e^*] d^* \quad (6.54)$$

Both ends of the analog plant system (FP_a and $F_wP_{w_a}$) are sampled-data systems. Because they sample at the same time and same rate however, the error-filtering block can be moved from one side to the other as in Fig. 6.27. In this configuration, both the disturbance path and actuator path of the standard controller (Fig. 6.17) can be represented as a discrete-time system as far as the feed-forward controller is concerned. This allows for the standard FxLMS algorithm to be applied to the optimal filter calculation of Sec. 6.3.

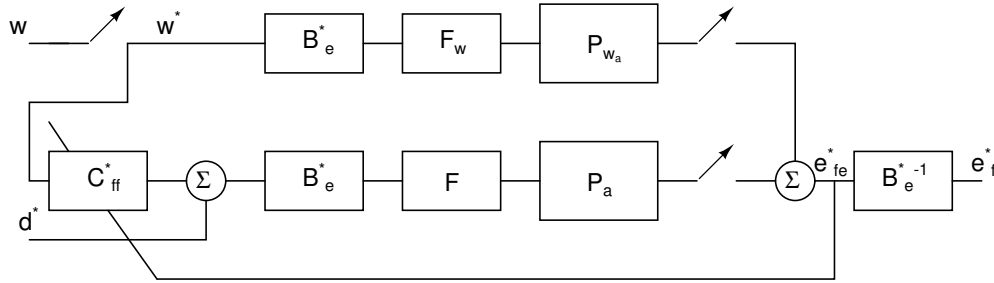


Figure 6.27 FxLMS error-weighting equivalent consolidated

The “disturbance path” contains the the error filter, and the actual disturbance path added disturbance rejection provided by the analog loop. It also contains the models for the required smoothing and anti-aliasing filters. The “secondary path” contains the same error weighting, smoothing, holding, and actuator paths. The resulting paths can be seen in Fig. 6.28 for an example weighting filter. The B_e^* weighting filter consists of a bandpass between 400-1000Hz, with a 5kHz sampling rate.

Although the autocorrelation matrix inversion is a computation that results from the direct solution to the FxLMS algorithm, it proved to be numerically sensitive with the addition of the error-weighting. The two filtered plant responses in 6.28 were used to generate the time-domain signal just as in the unweighted FxLMS case. Unfortunately, the dynamic range involved in the bandpass-filtered plants prevented the correct coefficients from being calculated. To solve this numerical problem, a small amount of conditioning was added to the R_f matrix before inversion by way of $R'_{ff} = R_{ff} + I * (10^{-6} \|R_{ff}\|)$. This had the effect of reducing the condition of the matrix by a factor of over 10^8 . Running this through the modified FxLMS routine yields Fig. 6.29 as an output filter.

The output filter generated by the error-weighting LMS algorithm of Fig. 6.29 does not look appreciably different from the unweighted FxLMS of Fig. 6.20 with one exception. The




Figure 6.28 Error-weighting FxLMS input frequency response




Figure 6.29 Error-weighting FxLMS output filters

out-of-band regions of the filter no longer attempt to match the unweighted plant dynamics of Fig. 6.20. While not immediately obvious by inspection, the fit of the optimal filter *within* the passband is better in the error-weighted case.

The reason for band-passing the error for the LMS controller was to target the algorithm to the mid-frequency regime. While the optimal filter shown in Fig. 6.29 shows a good match within the passband, the actual performance of this filtering strategy shown in Fig. 6.30 shows a few problems. Figure 6.30 shows that the mid-frequency performance of the error-weighted FxLMS scheme is improved over the unweighted performance of 6.21. The resonant peaks in the 400-1kHz range were reduced 10-15 dB for the unfiltered FxLMS. For the filtered error weighing, the peaks were reduced 15-25 dB within the same frequency range. This increase in performance shows that the optimal filter was better able to utilize the finite number of degrees of freedom with the FIR filter to the desired frequency region.




Figure 6.30 Broadband FxLMS performance with error weighting

Unfortunately, this filtered-error solution exhibits out-of-band artifacts. There is some significant spillover at the low frequencies, and continued performance above the designed fre-

quency range. A more predictable error-weighted strategy was implemented with the addition of a filtered reference signal.

6.4.3 Filtered-error, filtered-reference LMS

The intuitive filtered-error method of frequency weighting the FxLMS algorithm produced the undesirable optimal control spillover for out-of-band frequencies. As is often done in applied FxLMS [96], a reference filter was added to the system to aid in reducing out-of-band spillover. After the same block diagram manipulation that resulted in Fig. 6.26 to move the error filter to the left side of the model, the filtered-error, filtered-reference LMS weighting scheme can be represented as Fig. 6.31. The signal e_f^* signal is the original filtered error, with the “filtering” being the continuous-time anti-aliased filter contained in the F and F_w . The e_{fe}^* is the same filtered error signal, but with the addition of the discrete-time error filter B_e^* . In the configuration of Fig. 6.31, both B_e^* filters are taken to be identical, although in general they need not be. This addition of the filtered reference to the optimal filter computation did not remove the spillover effect in the computation of the optimal filter. In the final *implemented* configuration (6.55) however, the additional filter reduced the spillover in the optimal filter, because it stayed in the secondary path (6.56) as a residual B_e^* .

$$e_{fer}^* = \left[(F_w P_{wa})^* B_e^* + C_{ff}^* (F P_a)^* B_e^{*2} \right] w^* + [(F P_a)^* B_e^*] d^* \quad (6.55)$$

$$e_f^* = \left[(F_w P_{wa})^* + C_{ff}^* (F P_a)^* B_e^* \right] w^* + [(F P_a)^*] d^* \quad (6.56)$$

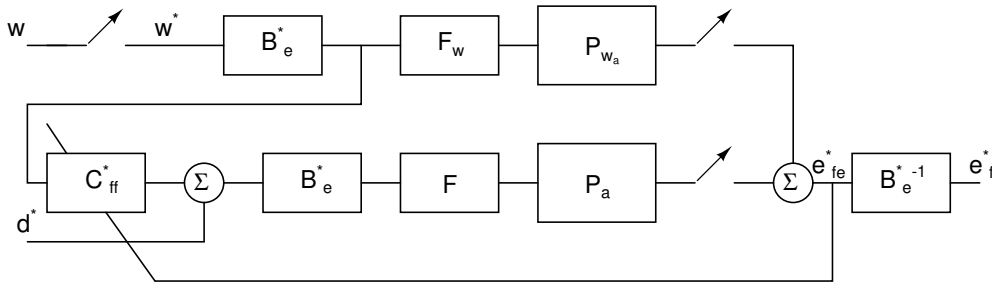


Figure 6.31 FxLMS error-weighting equivalent with reference filtering

Using the same bandpass error filter as in the previous case, the optimal 256-point filter was computed in Fig. 6.32. The spillover is still present, as is the aliasing at 2.5kHz. The

final continuous-time filter applied to the high-frequency model is again shown in blue. The additional filtering term, B_e^* attenuates the low-frequency spillover resulting from the optimal filter computation penalizing only the bandpass-filtered error.




Figure 6.32 Error-weighted, reference filtered output filters

The simulated performance of this scheme can be seen in Fig. 6.33. Again, the reduction in the peaks in the target 400-1kHz bandwidth is increased from the unfiltered case. The spillover is essentially eliminated, although the in-band performance does decrease slightly. This can be explained by looking at the output of the optimal filter between the filtered-error and filtered-error with filtered reference. The black frequency response curves of Fig. 6.32 shows an increase in the optimal filter’s control authority in the transition bands (200-400 Hz and 1100-1300 Hz) over that in the filtered error case in Fig. 6.29. This increased gain of the optimal filter is largely eliminated by the *implemented* filter, or the blue curve in Fig. 6.32. So, while the filtered-reference has the desired effect of greatly reducing spillover and targeting the optimal filter to the desired frequency range, it does so at the expense of a bit of in-band performance. Still, the reduction within the narrow band is greater than with the unfiltered case.




Figure 6.33 Broadband FxLMS performance with error weighting and filtered reference

The inclusion of the filtered reference in the optimal design can be rationalized intuitively as well. The function of the filter in the error path shapes the *error*, and thus affects the convergence and optimal result of the LMS algorithm. When a filter is added to the reference path, it shapes the *reference* spectrum and can filter out the spillover produced by the filtered-error LMS optimal filter. It does this at the expense of remaining in the feed-forward path during filter operation. Thus, the optimal filter contains the inverse of the error filter, only to be removed by the reference filter at run-time. The optimal filter “wastes” some of its degrees of freedom adapting to what will ultimately be removed anyway.

6.4.4 Model extension for low-frequency feedback

One of the more tedious modeling problems for the controller was adding in the low-frequency feedback control to the analog and feed-forward system. For the case of analyzing only the high-frequency and feed-forward control, the analog control simply modified the apparent plant as far as the discrete-time controller was concerned. This moved the plant’s natural frequencies in a manner investigated in Sec. 4.2. With the mixed feed-forward *and*

feedback digital control, a more in-depth analysis was necessary. The mixed feed-forward and feedback had to be included, as well as the additional decimation and interpolation factor of the low-frequency feedback controller. The actual three-controller layout can be seen in Fig. 6.34. This includes the filtered-error, filtered-reference weighting discussed in Sec. 6.4.3 to optimize the mid-frequency feed-forward control. Just as in Sec. 6.4, the block-diagram manipulation and sampled-data systems techniques of [104] were used.

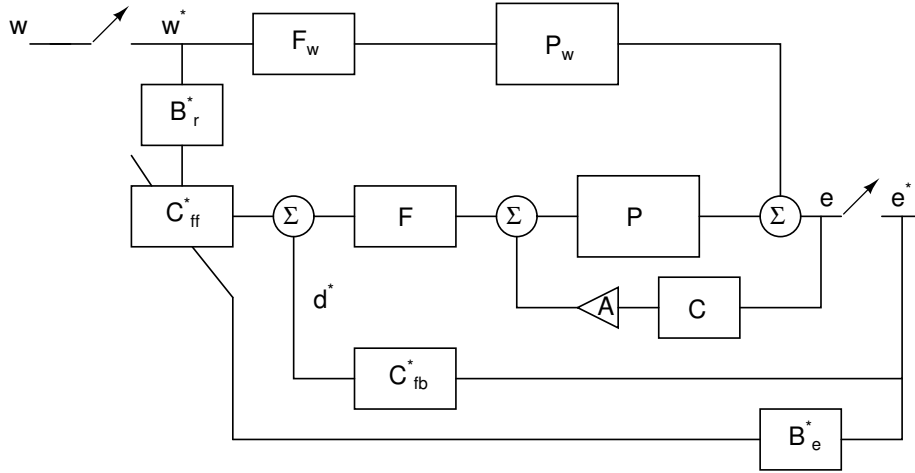


Figure 6.34 Three controller combination configuration

The consolidation of the continuous-time controller can be seen in Fig. 6.35. The summing junction of the primary (disturbance) and secondary (control) path can be separated into two modified subsystems. The same subsystems as defined in (6.51) and (6.52) were used to show the effect of the continuous-time feedback controller in both paths. For clarity of illustrating it as a feedback path, the equivalent of the system is shown equivalently *as* feedback in the disturbance path, with (6.57)

$$P_{w_a} = \frac{P_w}{1 - \left(P_w \frac{P}{P_w}\right) AC} \quad (6.57)$$

With the continuous-time feedback path adequately modeled in Fig. 6.35, the sampler and error filters could be moved through the summing junction. In general, this is not possible with sampled-data systems, but in this case it was. The two paths are simultaneously sampled, and the sampling operator is linear [104]. The two signals are simultaneously-sampled, so the time-varying effect can be ignored. Figure 6.36 shows an analytically equivalent system, P_a^* and $P_{w_a}^*$. From the standpoint of both the feed-forward and low-frequency feedback controller, the

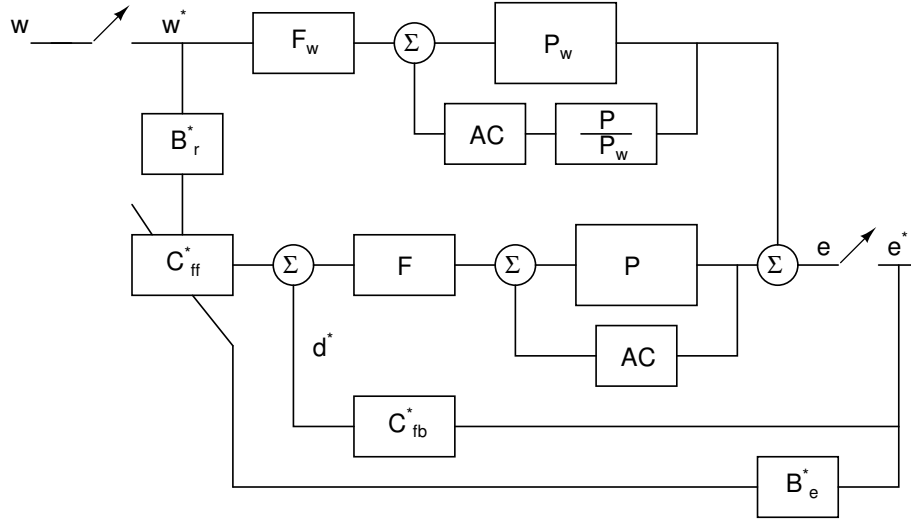


Figure 6.35 Three controller combination with high-frequency controller equivalent

continuous-time system could now be viewed as an “equivalent” sampled-data system. Since the anti-aliasing, hold and smoothing filters are contained within F and F_w , it made sense to group them together with the continuous-time system. The sampled-data system models of the two paths are summarized in (6.58)-(6.59) and Fig. 6.37

$$P_a^* = \left(\frac{FP}{1 - PAC} \right)^* \quad (6.58)$$

$$P_{w_a}^* = \left(\frac{F_w P_w}{1 - PAC} \right)^* \quad (6.59)$$

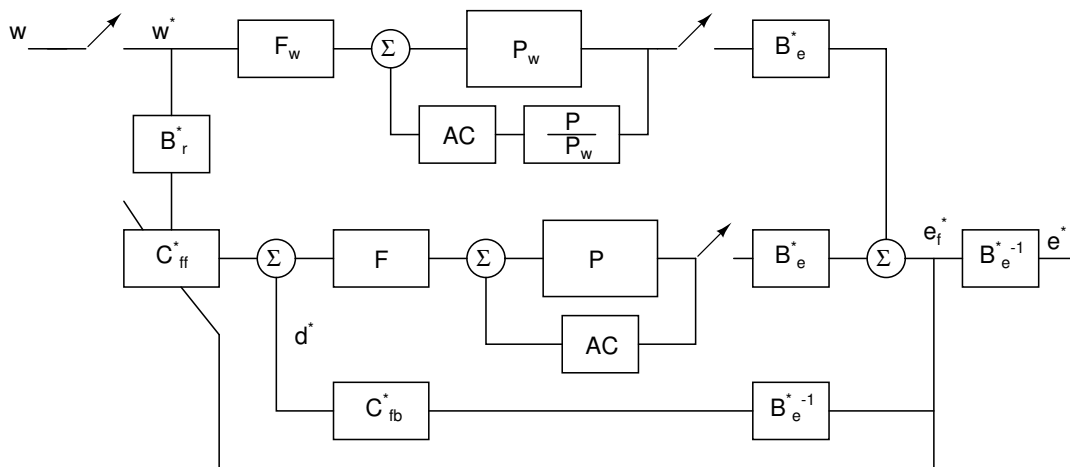


Figure 6.36 Three controller combination equivalent with in-path error filter




Figure 6.37 Optimal filter computations

Equation (6.60) shows the literal translation of Fig. 6.36.

$$E^* = \left[F_w \frac{P_w}{1 - PAC} + B_r^* C_{ff}^* \frac{FP}{1 - PAC} \right]^* W^* + \left[\frac{FP}{1 - PAC} C_{fb}^* \right]^* E^* \quad (6.60)$$

Combining feedback terms yields the disturbance to error relationship of (6.61).

$$\frac{E^*}{W^*} = \frac{\left[F_w \frac{P_w}{1 - PAC} + B_r^* C_{ff}^* \frac{FP}{1 - PAC} \right]^*}{1 - \left[\frac{FP}{1 - PAC} C_{fb}^* \right]^*} \quad (6.61)$$

Applying the simplifications of (6.58)-(6.59) yields the transfer function relationship of (6.62), since the “*” operator can be distributed through addition.

$$\frac{E^*}{W^*} = \frac{P_{w_a}^* + B_r^* C_{ff}^* P_a^*}{1 - C_{fb}^* P_a^*} \quad (6.62)$$

Finally, the “feedback form” of the two paths can be seen in (6.63), and graphically in Fig. 6.38.

$$\frac{E^*}{W^*} = \frac{P_{w_a}^*}{1 - C_{fb}^* P_a^*} + B_r^* C_{ff}^* \frac{P_a^*}{1 - C_{fb}^* P_a^*} \quad (6.63)$$

Pulling the B_e^* error filter inside the loop for Figs. 6.36-6.38 would not be strictly necessary, except that one more simplification was made. Because of the simultaneous sampling at the input and output of the system, the error filter can be moved to before the plant. Also, this investigation assumed that the reference filter, B_r^* and the error filter, B_e^* were the same. Consequently, the final system of Fig. 6.39 could be obtained.

What Fig. 6.39 ultimately shows is that the sequential loop closure removes the dynamic interaction between the controllers. All of the continuous-time controller dynamics are contained within the P_a^* and $P_{w_a}^*$ blocks. As assumed from the beginning of the investigation in Sec. 3.1 the feed-forward path gets a reference signal from an uncoupled source. With this assumption, the effects of the actuator are not seen by the reference signal, so there is no high-frequency

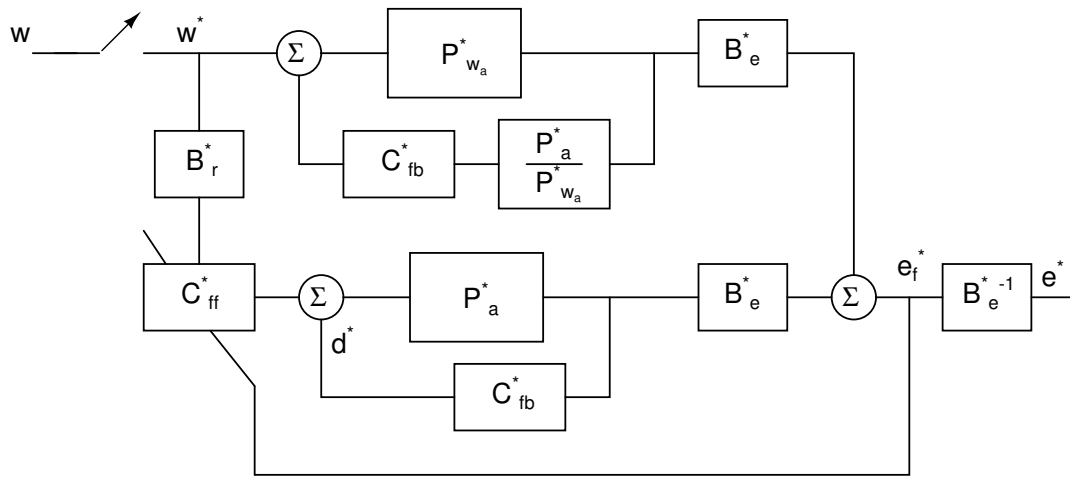


Figure 6.38 Three controller combination equivalent

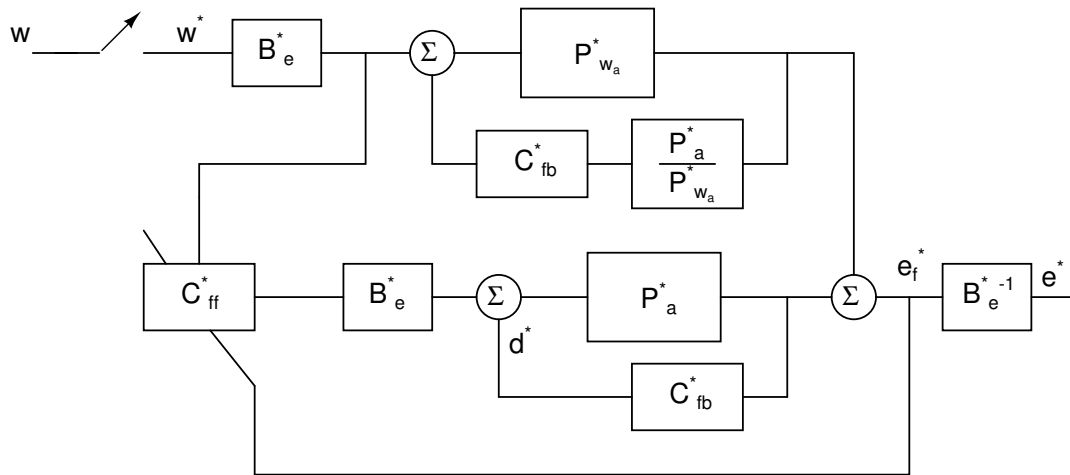


Figure 6.39 Three controller combination equivalent with $B_r^* = B_e^*$

interaction from the feed-forward control. Finally, the low-frequency actuator path stability was one of the *defining* characteristics of its synthesis in Sec. 4.3. As it turns out, this loop’s stability was also the characteristic for the disturbance path as well. The effects of the errors in the loop are investigated in Sec. 6.5.

6.5 Stability Analysis

The block diagram of Fig. 6.39 clearly shows all relevant parts of the control scheme in a “standard form.” The continuous-time feedback, anti-aliasing filtering, and sampling operations are contained within the P_a^* and $P_{w_a}^*$ blocks. The only potential for feedback between the two coupled plants, P_w and P are through the adaptation path of the feed-forward controller and its coupling to the reference signal. Because the reference signal is assumed to be decoupled from the structure, however, this means that there is no feedback path coupling the two inputs to the structure. This fundamental assumption is what allows the independent stability analysis of the two loops. The stability and design of the high-frequency analog loop gain of *PAC* was investigated in Sec. 4.2. This, coupled with the sequential-loop closure meant that the stability of the low-frequency loop could be investigated relatively independently.

The feedback portions of the actuator loop in Fig. 6.36 can be seen in Fig. 6.40. The ‘**’-notation represents the low-frequency controller at it’s decimated, low-frequency sampling rate. The F_L^* blocks represent the required digital anti-aliasing and smoothing filters following the decimation and interpolation for the low-frequency controller. The C_L^{**} block represents the low-frequency controller feedback compensator, implemented at the decimated sampling rate.

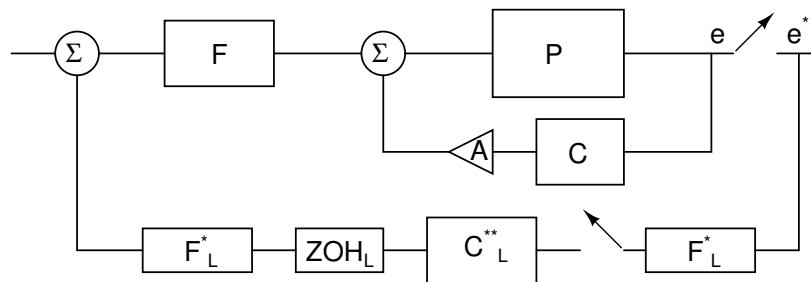


Figure 6.40 Actuator path loop gain

Although the decimated/interpolated blocks are simple to express, they present difficulty in ensuring the closed-loop system stability. For classical digital control design, the sampling

rate and associated filters are generally taken to be sufficiently high so as to be largely ignored. Within the control field, they are also generally taken to be at the *same* sampling rate. The multi-loop/multi-rate control proposed throughout this project can not necessarily assume either.

The sampling rate adjustment and filtering requirements were initially addressed in Sec. 6.2. In order to ensure the stability of the entire system however, the effects of the various gain bounds and sampling aliasing were quantified. Unfortunately, different components of the single control loop gain of Fig. 6.40 have error mechanisms from different sources. The sequential-loop closure accounted for the interaction terms from previously enabled controllers, but not the effects of sampling. To consider this, an uncertainty argument from the robust control field was used.

Figure 6.41 shows the generalized form of multiplicative uncertainty in a feedback plant. Using the techniques summarized in [120], this Δ addition to the loop gain of a linear feedback system results in the total loop gain of $GK(I + \Delta)$. This form of loop perturbation leads to the consideration of the complimentary sensitivity function, $T = (KG)(I - KG)^{-1}$. So long as $\bar{\sigma}(\Delta) < \frac{1}{\bar{\sigma}(T)}$, the closed-loop, perturbed system remains stable. This is a conservative bound that depends only on the magnitude of the uncertainty. Phase can be assumed to be worst-case.

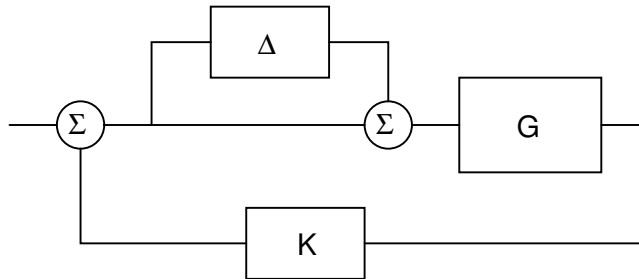


Figure 6.41 Multiplicative uncertainty

In general, all components of the system can be multi-variable and fully coupled. For the purposes of examining the feedback of Fig. 6.40 however, a SISO analysis was adequate. Because a SISO case was sufficient for the loop analysis, the maximum singular values degenerate into the frequency response of the respective parts.

The error mechanisms for the mid-frequency system identification were initially addressed in Sec. 6.1.2 for the case when the continuous-time controller loop was closed. The result of

that analysis was an α_a bound that limited the uncertainty in the plant model due to sampling and aliasing errors and is summarized in (6.64).

$$\Delta_{P_a^*} \leq \alpha_a \frac{\sum_{k \neq 0} |F(j\omega - jk\omega_s)|}{|F(j\omega)|} \quad \text{from (6.28)} \quad (6.64)$$

The low-frequency, decimated error mechanism can be written in multiplicative error form in (6.65).

$$\Delta_{C_{fb}^*} = \frac{\tilde{C}_{fb}^*}{C_{fb}^*} - 1 \quad (6.65)$$

The equivalent discrete-time frequency response of the low-frequency controller in mid-frequency sampling rate is expressed as C_{fb}^* . The *actual* implemented gain with digital anti-aliasing filters and hold equivalents is represented by \tilde{C}_{fb}^* in (6.66).

$$\tilde{C}_{fb}^* = [F_L^* C_L^{**}]^{**} ZOH_L F_L^* \quad (6.66)$$

The error mechanisms of Fig. 6.40 are shown as uncertainty in Fig. 6.42. Individually, they are represented in Fig. 6.42(a), with the combined error in Fig. 6.42(b). The combination of the two errors is summarized with $\Delta_{P_a^* C_{fb}^*}$ in (6.68).

$$1 + \Delta_{P_a^* C_{fb}^*} = (1 + \Delta_{P_a^*}) (1 + \Delta_{C_{fb}^*}) \quad (6.67)$$

$$= 1 + \left(\Delta_{P_a^*} + \Delta_{C_{fb}^*} + \Delta_{P_a^*} \Delta_{C_{fb}^*} \right) \quad (6.68)$$

$$\leq 1 + \left(|\Delta_{P_a^*}| + |\Delta_{C_{fb}^*}| + |\Delta_{P_a^*} \Delta_{C_{fb}^*}| \right) \quad (6.69)$$

Using the analytical model of the plant and controllers, the frequency response of the various error mechanisms can be seen in Fig. 6.43. The plot assumes the mid-range α_a parameter of 0 dB, which results in the Δ_{P_a} in the black line. The other contribution of $\Delta_{C_{fb}}$ including decimation/interpolation errors can be seen by the blue line. As expected, the higher-order terms contribute little additional error to the final $\Delta_{P_a C_{fb}}$ frequency response shown with the red line.

Finally, the total uncertainty in the loop was compared with the complimentary sensitivity plot of the nominal loop gain in Fig. 6.44. The digital anti-aliasing filter dominates the response above the low-frequency decimated sampling rate, as suspected. At all frequencies however, the

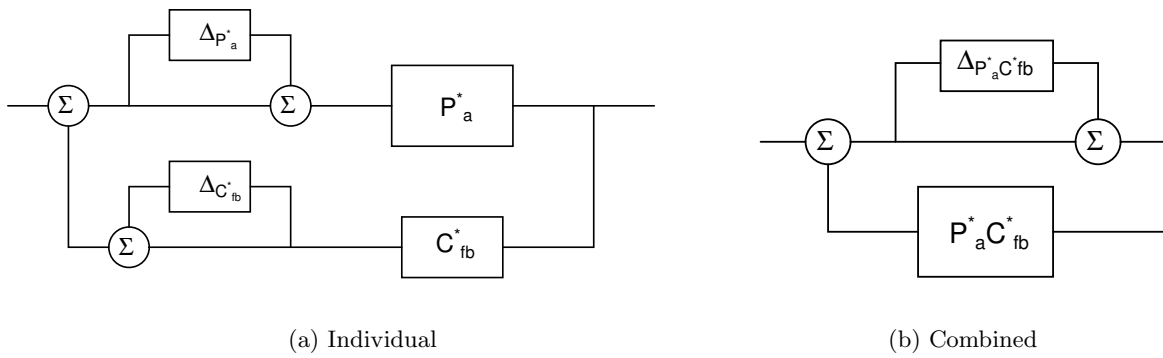


Figure 6.42 Actuator loop gain uncertainty




Figure 6.43 Actuator loop error mechanisms




Figure 6.44 Actuator loop sensitivity bounds

aggregate aliased, hold, and filtered uncertainty, $\Delta_{P_a C_{fb}}$, remains less than $1/T$. This ensures that the loop remains stable.

What this analysis shows is that the error mechanisms of all of the continuous-time and discrete-time controllers do not cause any stability concerns. The sequential loop closure and uncoupled reference signal in the feed-forward controller already minimized interaction between the “baseband” equivalents of the plants. With all of the implementation details of uncertainty in sampling, decimation, and interpolation, it was unclear as to whether this assumption still held. The uncertainty analysis here combines the different error mechanisms to show that it *did* hold throughout and the errors cause to not cause instability of the final system. This uncertainty analysis was extended to include experimentally-measured plant variation in Sec. 7.4.3.

6.5.1 Disturbance path stability equivalent

Section 6.5 addresses the *actuator* path of the complete system. In this section, feedback portions of the *disturbance* path are considered. Although drawn in feedback form in Fig.

6.39 for clarity, the disturbance path actually has a different structure. Figure 6.45 shows the disturbance path with all of the loops closed. Rewriting the low-frequency feedback component results in the forward path of Fig. 6.46. Finally, the effect of the continuous-time controller can be decoupled as in Fig. 6.47.




Figure 6.45 Disturbance “feedback” loop from total model




Figure 6.46 Disturbance “feedback” loop partially unwrapped

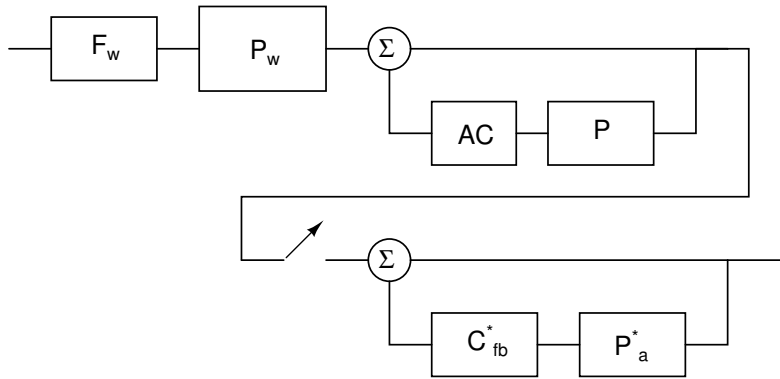


Figure 6.47 Disturbance “feedback” loop unwrapped

These loop manipulations preserve the sampling operator’s effect on all of the feedback loops. Fortunately, both of these feedback components have already been investigated for stability. The continuous-time loop is the same as was analyzed in Sec. 4.2, and the low-frequency loop in Sec. 6.5. Thus, for both the actuator path, and the disturbance path, the stability of the inner feedback loops was assured.

6.6 Summary

This chapter presented the analysis of the complete, hybrid continuous discrete model of the controller implementation. It included a thorough quantification of the errors associated with the sampling operation both with and without continuous-time feedback in place. A bound limiting the amount of this error was included for both cases. In addition, the errors associated with modeling the decimation and interpolation operation as implemented were analyzed and found to be relatively minimal.

Finally, a complete model of the system was used to combine the modeling errors into a single, unified uncertainty for the low-frequency controller. This uncertainty implies stability under these errors for this controller. The stability of the high-frequency and mid-frequency control were considered in Sec. 4.2 and 4.4, respectively.

CHAPTER 7

EXPERIMENTAL RESULTS

This chapter discusses the process of implementing the controller strategy on the experimental structure and the measured results obtained. Due to the complexity of the controller topology, sequential loop closures, and different possible design synthesis tradeoffs, the layout of this chapter is primarily chronological. The individual loops are closed, the results discussed, and then summarized. The hardware experimental controller setup and computational performance is described in detail in Appendix A-B, so this chapter focuses on only the results.

The high-frequency, continuous-time controller loop was closed first, and its measured implementation and performance is discussed in Sec. 7.1. Once that loop was closed, the low-frequency controller's implementation could be pursued. This process, the design tradeoffs due to actual experimental implementation, and the resulting performance are detailed in Sec. 7.2. Section 7.3 investigates the experimental implementation of the mid-frequency, adaptive feed-forward controller. The details of the experimentally-required design constraints are also considered. Finally, Sec. 7.4 discusses the aggregate performance of all the controllers. This includes their contribution to one another, and the overall contribution to the structural control. In addition, performance degradation due to unmodeled system dynamics and plant uncertainty are also experimentally quantified.

7.1 High-frequency Results

The high-frequency controller was constructed as in the schematic of Fig. B.1 with details in Appendix B. The frequency response was optimally computed in Sec. 5.2.2 and amounted to a very mild band-pass filter with high-frequency rolloff. Figure 7.1 shows this optimal high-frequency compensator as compared to the experimentally-measured controller. The loop was

measured by exciting the analog circuitry with the white noise from control computer and measuring the response after the output of the power amplifier, but before driving the control actuator. The control computer was sampling at the higher-frequency “measurement” rate of 10kHz, although this is still somewhat of a factor in measuring this continuous-time controller.

There were a number of reasons for using the controller computer for measurement collection. The most significant reason was because all of the implementation-specific hardware (anti-aliasing filters, smoothing filters, power amplifiers, high-frequency controller dynamics, discrete-time controller delay, etc) made it very difficult to match an external measurement system with the controller system. For these reasons, the controller computer was used to measure the system’s response for ID and subsequent controller synthesis. If the final measurements were done on a completely different acquisition system, the open/closed loop comparisons would require significant post-processing to directly compare them. The different transfer functions of the two systems would have to be normalized out. Since the controller computer was already in place and used for system identification, it was straight-forward to increase the sampling rate to measure more of the continuous-time controller’s response. Additionally, the computer’s effective “control” sampling speed was decimated from the actual sampling speed. This provided additional measurement bandwidth with minimal computational increase. Appendix A discusses the implementation in further detail.

The effective bandwidth of the continuous-time controller was higher than was measurable by even this increased rate, but all of the transition-region was captured. This transition region includes the mid-frequency controller’s bandpass filter (up to 1kHz), any out-of-band, feed-forward spillover, and ultimately the anti-aliasing filters and transition *through* the 2.5kHz Nyquist rate. The anti-aliasing filter and zero-order hold can be seen as a substantial phase rolloff at the higher frequencies. In order to compare the optimally-computed compensator (black line) to the experimentally measured results, the anti-aliasing filter and holds were included with the optimal compensator frequency response in Fig. B.1. These were not actually present in the loop however, as the continuous-time controller used the error signal *before* it was sent through the anti-aliasing filter, and its output was not sent through the smoothing filter. The comparison of the optimal and experimental plots are generally in agreement. Some aliasing can be seen as a measurement artifact above 3kHz, but the majority of the measurement bandwidth agrees with the designed controller.




Figure 7.1 Measured continuous-time controller compensator

Three different experimental loop gains were tested to see the increasing performance at the expense of stability margin. The synthesis section on the simplified model (Sec. 4.2.3) predicts spillover near the actuator resonance frequency. This provided one means to see the impending instability with increasing loop gain. The experimentally-measured results for the actuator path can be seen in Fig. 7.2. Decreasing the gain from the optimal controller had the expected result of decreasing the attenuation at higher frequencies, while also decreasing the spillover at actuator resonance.

Although difficult to measure, there was an additional bit of information in the experimental implementation that lends further credibility to the high-frequency optimization procedure of Chap. 5. While iterating on gains to obtain the best intermediate signal scaling between the controller and power amplifier, the analog gain was often adjusted. When set to a marginally-stable high gain, the instability on the structure could be audibly heard. Interestingly, a low-frequency “rumble” and high-frequency “squeal” grew in amplitude at the same rate and intensity as one another. This would continue unchecked until the high-frequency controller gain was reduced (or the amplifier clipped if allowed to grow too long). This simultaneous growth




Figure 7.2 Measured continuous-time controller performance (actuator path)

implies that the optimal controller design was successful in maximizing the gain between the high and low instability frequencies. If one frequency had gone unstable first, it would have implied there was more spectral “room” available at the other frequency and the controller was not optimal.

The disturbance path attenuation can be seen in Fig. 7.3. Again, the increasing gain toward the optimal controller resulted in increased performance in the high-frequency band. The actuator spillover is largely missing from the disturbance path, although can still be slightly seen if examined closely.

The high-frequency controller loop was closed throughout the rest of the experimental measurement. As discussed in Sec. 4.2, this modified the plant as far as the discrete-time controllers were concerned. All subsequent system identification was done with it in place and its stability was ensured from the stability margin in the design of Chapter. 5.

7.2 Low-frequency Results

The low-frequency controller implementation was the second loop to close for the total controller. Since the high-frequency continuous-time controller was already closed, all of its dynamics were already present in the synthesis design model. The system was excited with random noise at the decimated sampling rate. All of the anti-aliasing filters (both continuous-time and discrete-time) were present in the loop for the measurement. The actuator and the disturbance input were excited in turn, and the resulting time-domain data recorded.

The time-domain data was identified with the SOCIT toolbox as in Chap. 3. With the addition of the decimators, holds, and anti-aliasing filters however, the system ID resulting from the direct time-domain, decimated data was a more accurate representation of the the effective plant. It captured all of the relevant dynamics that would be seen by the controller, without requiring analytical model order reduction. This resulted in a lower order, and simultaneously more accurate model than with the large full sample-rate identified system used for analysis.

Another reason for this was discussed in Sec. 3.3. The SOCIT toolbox does not provide accurate models when given fewer degrees of freedom than states of the identified system. If the plant has 100 identifiable states, merely allowing for 100 states does not result in as accurate of a model as allowing for more states in the identification and then reducing.




Figure 7.3 Measured continuous-time controller performance (disturbance path)

The quality of the system identification can be seen in Figs. 7.4-7.5. Both figures show the raw time-domain data spectral response shown as an FFT transfer function from excitation to output acceleration in black. The blue shows the full-model order (198 state) 2-input, 1-output state-space model. At almost all frequencies, the baseline model order was sufficient to almost completely describe the measured response. The reduction in accuracy of the reduced-order model can be seen by the various dashed red plots. As can be seen from the response, 100 states was largely sufficient to model all of the dynamics, but reducing this to 50 resulted in significant deviations from the measured response. The green line shows the result of a 98-state model directly identified by the SOCIT toolbox. Upon close examination, this fit was not as accurate as a higher-order fit reduced down to 98 states.

7.2.1 Controller variations

With the identified analytical model of Sec. 7.2, the synthesis computation for the low-frequency controller could be performed. The high-order identified model was reduced to 100 states because it sufficiently captured the system dynamics, but allowed for much less computation of the optimal low-frequency controller. As discussed in Sec. 4.3, the H_2 design process allowed for a number of different design parameters to tailor the result. The frequency-weighting term was accomplished by a simple second-order filter, but the “aggressiveness” of the design was adjusted by simulating different sensor noise. For the purposes of this investigation, it was sufficient to leave this noise unweighted in frequency since the net goal was to decrease the effect of a broadband disturbance.

Four different values of sensor noise were used in design of the experimental measurement controller. The absolute magnitude of this sensor noise parameter are only relevant within the context of the identified plant model. This included any sets of continuous-time controller gains and external power amplifier gains. Comparison singular value plots similar to the synthesis design example of Fig. 4.16 were used to ensure that the specific gains used in the experimental model were appropriate. The four different resulting compensators can be seen in Fig. 7.6 for the varying noise floor thresholds. Each of the noise floor levels was decremented by a factor of 2. The exact gain change at all frequencies depended on the outcome of the H_2 -optimal controller, but nominally increased the gain of the compensator by an amount similar to the 6dB change in sensor noise. If the noise level was decreased much beyond $V = 0.25$, the






Figure 7.4 Low-frequency ID model order comparison (actuator path)



191

Figure 7.5 Low-frequency ID model order comparison (disturbance path)

resulting design from the H_2 optimization routine was identified as unstable. These unstable compensators were not implemented experimentally and provided a limit to the performance of the controller.




Figure 7.6 Low-frequency controller gains varying design noise level

The loop-gain (compensator with actuator) with the four different compensators was calculated and shown in Fig. 7.7. Compared to open-loop, the compensated gain shows an increase in loop gain at some frequencies, with a significantly reduced loop gain at others. One of these large reductions in gain was at the actuator resonant frequency around 120Hz. This makes sense to normalize out the effect of the actuator. The actuator was a relatively low mass in comparison to the plate it was attached to. Thus, its effect on the disturbance to error path would be minimal.

7.2.2 Performance

The measured magnitude performance of the four different “gains” can be seen in Fig. 7.8. Fig. 7.8(a) shows the closed-loop actuator gain with the compensator in the loop. Since the H_2 optimal controller targeted the *disturbance* path rather than the actuator path, some increases

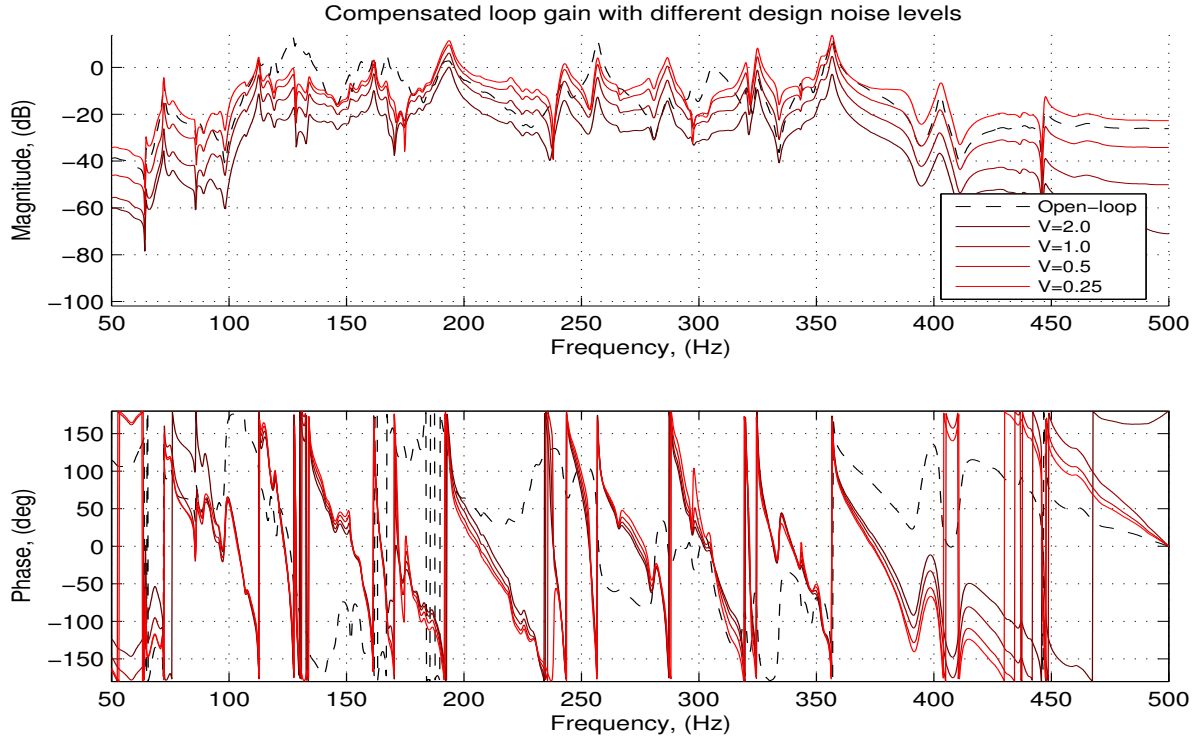



Figure 7.7 Low-frequency actuator loop gain varying design noise level


can be seen in the actuator’s closed-loop response. The 200Hz frequency in particular shows how the zero at the disturbance path results in little penalty for actuator overshoot there.

Figure 7.8(b) shows the magnitude of the measured modeled disturbance rejection. The resonant frequencies’ contribution are significantly reduced. With the increasing assertiveness of the optimal controller with the smaller noise levels ($V = 0.25$), the largest resonant peaks are reduced over 10 dB.

The most aggressive controller was also compared to the simulation results in Figs. 7.9 and 7.10. The measured results for the actuator loop are shown in Fig. 7.9. At almost all frequencies, the measured response was as predicted. At 110, 260, and 360Hz, the closed-loop actuator was reduced as expected from the open-loop gain of Fig. 7.7. At the actuator frequency of 120 Hz, there was a slight deviation. This can be explained by the small difference in the reduced-order identified model to the measured frequency response. Fortunately, this unmodeled overshoot did not show appreciably in the disturbance performance, nor cause stability problems. Other than modeling errors, the experimental results of the actuator loop were almost exactly as predicted from the optimal control synthesis.



(a) Actuator path



(b) Disturbance path

Figure 7.8 Low-frequency measured performance with varying aggressiveness




Figure 7.9 Low-frequency measured performance (actuator path)



The measured results for the disturbance path are shown in Fig. 7.10. Again, the expected results were computed using the full-order model, but the design implemented with a reduced-order model. The model order reduction does not cause any appreciable error at 120Hz. Only around 160Hz can the experimental model be seen to deviate from the expected results. Again, this small deviation between measured and predicted performance can be explained by the errors in the reduced-order model.

7.2.3 Implementation artifacts

The experimental results for the full-order optimal controller were almost exactly as predicted from the synthesis computation. The small discrepancies can be explained by analytical finite model order as compared to the experimental measurement. To investigate this phenomenon in detail, the model order was decreased even further. This alleviated some of the computational load, but in excess caused stability problems.

Figure 7.11 shows the effect of reducing the model order from the designed order comparable to the model ID of 100 states, to only 60. This model reduction was done with a balanced realization, minimal-contribution state removal. At all but the frequencies around 120Hz, 60 states adequately captures the compensator dynamics. The resulting experimental measurement for the 60-state model order reduction case is shown in Fig. 7.12 looks different, however. The model mismatch causes significantly different results than expected at actuator resonance. For the 50-order model, the resulting experimental model was unstable so no experimental measurements could be taken.

The model order reduction can be seen to be an important aspect of the design. As previously mentioned, the SOCIT identification toolbox did not work well with too few degrees of freedom to choose its model. Much higher quality identifications were possible when the identification model order allowed to be higher than necessary and then the resulting model reduced. This model order reduction should be done *before* the computation of the optimal controller to minimize low-gain model mismatches such as in Fig. 7.11. In reality however, the computer hardware utilized for this investigation generally had sufficient processing power to implement the low-frequency feedback loop at its full design model order of approximately 100 states. Since the implementation was done using a direct state-space model, the numerical stability problems associated with transfer functions of such a high order model were minimized.



170

Figure 7.10 Low-frequency measured performance (disturbance path)




Figure 7.11 Low-frequency reduced order compensator




Figure 7.12 Low-frequency reduced order compensator measured performance

7.3 Mid-frequency Results

The mid-frequency computer control was implemented after both the high-frequency continuous-time controller and the low-frequency high-order discrete-time controller were activated. Just as in the low-frequency case, all anti-aliasing filters and holds were present in the loop when the system was excited with white noise at the mid-frequency sampling rate. The resulting time-domain data was both processed into a direct spectral transfer function, as well as run through the SOCIT toolbox to produce an analytical model.

Figure 7.13 shows the comparison between the direct frequency response data and the model-fit frequency response. In this case however, the 100 states that were sufficient to capture almost all of the system dynamics for the low-frequency controller in Sec. 7.2 were only sufficient to capture the dominant plant dynamics. The intricacies of the closely-spaced, lightly-damped modes were not completely resolved.

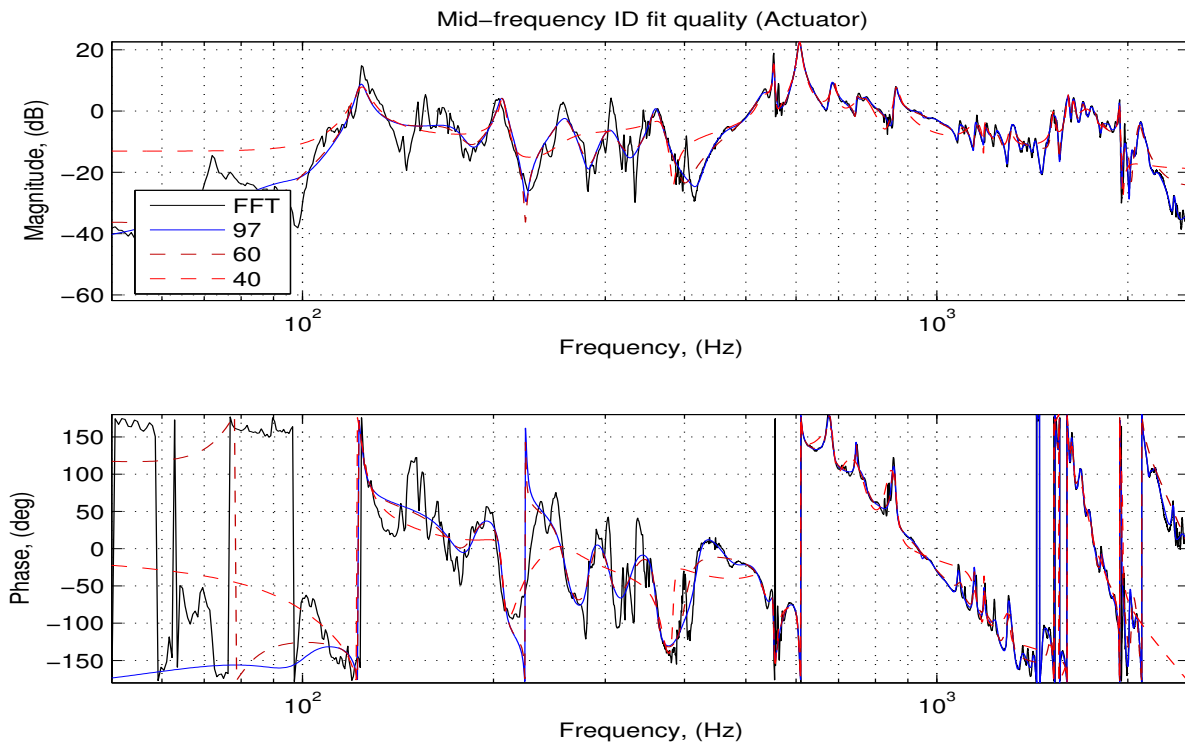


Figure 7.13 Mid-frequency ID model order comparison (actuator path)

Fortunately, as discussed in Sec. 4.4, the model order requirements for the adaptive FxLMS controller were not as demanding as for the feedback controller. So long as the phase estimate

of the actuator path is accurate to within $\pm \frac{\pi}{2}$, the adaptation algorithm will remain stable. The penalty for model mismatches of this order are only visible as increased adaptation time. From Fig. 7.13 it can be seen that a model order as low as 40 states captured the system dynamics to within this relatively generous phase constraint. For the experimental implementation in this section, a model order of only 40-states was used in the adaptive controller's estimate of the actuator path's dynamics.

One additional consideration of the mid-frequency controller was that of the filtered-error, filtered-reference investigated in Sec. 6.4.3. In order to consider this, two different error filter frequency ranges were implemented. The error filter was experimentally changed from between 400-1000Hz to 100-1000Hz. Because the actuator had very small gain below resonance, this effectively became the equivalent of the unfiltered, broadband FxLMS. The experimental results comparing these two filters can be seen in Fig. 7.14. As expected, the additional degrees of freedom made possible by band-passing the error to between 400-1000 Hz (red line) provided some increased performance in-band. The 100-1000Hz plot shows slightly smaller reduction within the 400-1000Hz bandwidth, but with some fairly large reductions below 400Hz.

The explanation for this can be seen with in the structural response itself. Within the mid-frequency region, the number of significant modal peaks in the disturbance path is relatively small and not too lightly damped. Because of this, the 256-states in the adaptive controller were mostly sufficient to capture the dynamics of the disturbance path. Additional degrees of freedom available due to not penalizing the disturbance out of band only resulted in small additional improvements.

It is also worthwhile to put this into the context of coherence-limited results. One of the limits to control performance is the limit in measurement coherence [121] as $P_m(\omega) = 1 - |C(\omega)|$. Another one is that gradient-descent algorithms rely on minimizing the squared error, and thus tend to reduce disturbances down to a constant level (the noise floor). Finally, the frequency granularity possible with a finite number of filter taps provides another constraint on maximum possible filter performance.

Figure 7.15 shows the filtered and unfiltered performance normalized to the mid-frequency controller open-loop. It also shows $1 - |C(\omega)|$ for the measurement data taken. In the mid-band frequencies (400-1kHz), both filtered and unfiltered result in a reduction down to the highest common constraint. Comparing Fig. 7.15 to 7.14, the limits of of coherence (e.g. 500-600Hz),




Figure 7.14 Mid-frequency disturbance rejection unfiltered vs. filtered

relatively-fixed performance level (e.g. 400-500Hz), and frequency granularity (e.g. 700-800Hz) can be identified.




Figure 7.15 Mid-frequency disturbance rejection vs coherence limit

Finally, the experimental results compare favorably to the expected results from the optimal controller investigated in Sec. 6.3. The band-passed (400-1kHz) results look very similar to those in Fig. 6.33 (pp. 144), with the additional constraint of measurement coherence. The band-passed (100-1kHz) results look very similar to those of the unfiltered FxLMS results in Fig. 6.21 (pp. 133). The 256-coefficient filter does not provide enough degrees of freedom or frequency granularity to provide reduction to the noise floor in the low-frequencies. It does however, provide enough degrees of freedom to obtain most of the in-band performance obtainable with the filtered-error, filtered-reference from 400-1000Hz. For these reasons, unless noted the additional low-frequency performance of the 100-1000Hz filter was used for the remainder of the experimental results.

7.3.1 Transient performance

One other way to compare the differences between the two different FxLMS filtered variants is to see the transient performance of both. Figure 7.16 shows the filtered error signal for the two different versions of the controller. Starting with zeroed filter coefficients, the controller filters were allowed to adapt. The raw time-domain error signal was windowed with a 5 second time average RMS and plotted against the entire filter convergence time. After 5 minutes, the convergence of the filter had been assured. The first-order decay characteristic of the gradient-descent LMS algorithm can be clearly seen, and at least 3 time constants of even the slowest controller variation have elapsed.




Figure 7.16 Mid-frequency transient filtered error

The two solid lines show the disturbance-path filtered error reduction for the 400-1kHz and 100-1kHz filter bandwidth LMS controllers. These are shown with their respective filters on the time-domain data, which illustrates what the controller “sees” as it adapts. The dashed line shows the time-domain error of the 100-1kHz FxLMS controller, but bandpass-filtered from 400-1kHz rather than with its respective filter. This is useful to further illustrate how the additional

degrees of freedom in the 400-1kHz controller can be used to gain additional performance *within its band*. For this particular structure, the disturbance rejection was ultimately better with the 100-1kHz filter, even though it had less performance in the 400-1kHz band.

7.4 Aggregate Results

The three different controllers depended upon each other due to the sequential loop closure implementation. This section describes the performance measurement for the three systems together, and illustrates how each one sequentially improves the total system performance.

Figure 7.17 shows the full-bandwidth measured spectrum of the actuator loop with all three controllers active, as well as the open-loop. All of these controllers were at the most aggressive possible while still maintaining system stability. The high-frequency was set to 0dB from the optimal, the low-frequency was set to minimum noise in the H_2 design ($V = 0.25$), and the mid-frequency controller used the 100-1kHz filter. The high-frequency contribution is unchanged from the pure high-frequency controller in Fig. 7.2. The low-frequency controller contribution likewise looks similar to Fig. 7.8.

The disturbance path can be seen in Fig. 7.18. It shows the significant added disturbance rejection in the mid frequency regime as well as some additional reduction in the low-frequency regime.

The frequency-response performance of all of the controllers as shown in Fig. 7.18 is an intuitive way to see all three controllers' effect on the resulting disturbance rejection, but it does not lend itself to a critical interpretation of the aggregate broadband performance. For that, the progressive RS integral shows the performance more clearly.

7.4.1 Broadband performance measurement

The extremely high-order, lightly-damped nature of the structure makes the Bode representation of the performance difficult to interpret. Rather, a variation on the cost function equation, (5.1), in Sec. 5.1 was used. By removing the mean portion of the RMS definition, using just the root-square (RS), and leaving a frequency-dependent term, (7.1) has a direct




Figure 7.17 Full-bandwidth measured performance (actuator path)




Figure 7.18 Full-bandwidth measured performance (disturbance path)

relationship to the transmitted disturbance power.

$$P(\omega) \triangleq \sqrt{\int_0^\omega \left| \frac{E(\Omega)}{D(\Omega)} \right|^2 d\Omega} \quad (7.1)$$

In this progressive RS definition, the $E(\Omega)$ term is the spectrum of the error signal on the structure, and the $D(\Omega)$ term is the spectrum of the disturbance function. The entire expression integrates the energy from DC up until the maximum frequency value of interest, which allows for specific modal contributions to be seen in the non-decreasing resultant plot. The difference between (7.1) and (5.1) is that (5.1) was assuming a low-order model of the system, so there were no resonant peaks in the output. For the actual measured system performance, only resonances with appreciable gain contributed to the overall energy contribution. These are seen as areas of the resulting plot with steep, positive slope.

Figure 7.19 shows the same information as in the Bode representation of Fig. 7.18, but in the progressive RS way. Individual contributions at frequencies such as 110Hz, 190Hz, 360Hz, etc can be seen as increasing the distance between the two integration lines. It should also be noted that this is a plot of *aggregate* broadband performance, but with specific spectral contributions easily identifiable. Because of this, the higher the frequency on the plot, the more difficult it is for a controller to increase the distance between the open and closed-loop plots. To mitigate this decreased sensitivity at high frequencies, a progressive RMS plot (normalized to the aggregate integrated bandwidth) was also investigated. Unfortunately, the resulting interpretation did not show individual controller contributions as clearly as the RS plot, so was not included in the results.

The progressive RS method of examining the results, the spectral contribution from each individual controller can be viewed together in one plot. Figure 7.20 shows each controller activated, in turn. The high-frequency broadband performance is dominated by the spillover at 120Hz and 350Hz. Those two frequencies increase the amount of contribution from the disturbance over even the open-loop. Above 1kHz however, the additional reduction by the addition of the continuous-time control can be clearly seen. Comparing to the Bode representation in Fig. 7.3, each significant gain excursion that is reduced causes the aggregate distance between the open-loop integration and high-frequency-closed-loop line to increase.

The next loop that was closed was the low-frequency feedback loop. Comparing the green line on Fig. 7.20 to the Bode representation of Fig. 7.10, significant reductions were added




Figure 7.19 Full-bandwidth measured progressive RS performance (disturbance path)




Figure 7.20 Full-bandwidth measured progressive RS sequential performance (disturbance path)

to the aggregate performance at 110Hz, 190Hz, and 350Hz. Since the high-frequency loop was already closed for this controller, the effective open-loop contained the spillover. Not only was this spillover removed, but the structural contributions were also removed as predicted.

Finally, the contribution from the mid-frequency feed-forward controller can be seen for both instances of the error filter. The minimal broadband improvement between 400-1kHz is obvious here with the red line. The significant reduction from including the lower-frequencies can be seen, even though they are not adequately captured with the 256-point filter. The very small reduction in performance in the 400-1kHz band can be seen, but the benefit at lower frequencies significantly outweighs this.

7.4.2 Unmodeled disturbance

The preceding results target the goal of broadband reduction of a known disturbance injected within the structure. This model was entirely unnecessary for the high-frequency, collocated control because it was assumed that the structural variation was too high to robustly use the information within the feedback loop. At least one instance of the disturbance path was used in the design to tailor the H_2 low-frequency controller, but its effect was only to minimize a cost function in the design. The resulting controller's *stability* was only a function of the variability of the collocated actuator/sensor pair.

The purpose of using feedback was to not only minimize the known disturbance, but to also help reduce any unknown disturbances. In order to provide some measurement of this effectiveness, another set of measurements was taken. This one was with another disturbance actuator on the structure, located at a completely different position on the structure from the modeled disturbance. The low-frequency controller was designed using the original disturbance path. The measured performance for this unmodelled disturbance path can be seen in Fig. 7.21. In order to gauge the unmodelled performance with different low-frequency design “aggressiveness,” the performance was measured using the same compensators as those used in Sec. 7.2.1. The performance results for the unmodeled disturbance path in Fig. 7.21 can be directly compared to those of the modeled disturbance path in Fig. 7.8(b). The unmodelled disturbance shows reduction where the two disturbance paths shared common modes such as 260Hz, 325Hz, and 360Hz. Where the modelled disturbance path had a zero however, the controller did not penalize structural coupling at that frequency. A good example of this is at 210Hz. As the

aggressiveness of the optimal controller is increased, the overshoot in the unmodelled path also increases at these frequencies.




Figure 7.21 Low-frequency measured performance with varying aggressiveness

Figure 7.22 shows the full-bandwidth performance for the unmodeled disturbance. As expected, the feed-forward control has little effect, since it was measuring the disturbance path from the other disturbance actuator. Both the feedback controllers had some performance at most of the modal frequencies. The notable exception to this was the 210Hz spillover.

The sequential effect of all controllers (including the spillover penalty) can be seen in the progressive RS plot of Fig. 7.23 for the unmodeled disturbance. Although there are some reductions from the low-frequency controller, the entire progressive broadband performance is dominated by the overshoot at 210Hz. By the time the high-frequency controller becomes active above 600Hz, it starts to provide additional performance. Because the high-frequency controller is only low-order, it essentially only moves the structural poles for which significant gain excursions exist. Zeros are unaffected, as are poles with small amplitude.

The most important result of looking at the unmodeled disturbance path is to realize that a single instance of modeled disturbance may not be sufficient to adequately weight the H_2




Figure 7.22 Full-bandwidth measured performance (unmodeled disturbance path)




Figure 7.23 Full-bandwidth measured progressive RS sequential performance (unmodeled disturbance path)

design filters. The zeros of any specific disturbance path can introduce large spillover in other, unmodeled disturbance paths. Still, for the common modes that are present in both paths, the feedback controllers provided disturbance rejection for even the unmodeled disturbances.

7.4.3 Unmodeled structural variation

The hierarchal control design relied heavily on proper model identification for the low-frequency synthesis. To a lesser extent, the mid-frequency controller also relied on high-order model identification. In order to verify system performance and stability, the flexible plates on the structure were randomly perturbed to measure the systems' response. For all of the perturbations, the nominal ID was used for both the low-frequency controller and the control to error path ID. The adaptive controller was the only one allowed adapt to the modified plant.

Figure 7.24 shows the actuator path Bode plot with approximately 10 different plate variations. The mean of the open-loop actuator path is shown in black. The mean of the closed-loop actuator path is shown in blue, with the individual instances of the closed-loop performance shown in gray. The gray show a relatively widely varying closed-loop gain at a few of the lower frequencies. This did not affect the stability of the system, but a further analysis was warranted.

The robustness of the low-frequency controller to unmodeled structural variation was important to quantify. The argument for the complicated, hierarchal control strategy was that the structural unknowns were relatively small at the low-frequency, becoming increasingly important at higher frequencies. A mean/variance approach in addition to the loop sensitivity bounds of Sec. 6.5 were used as an indicator to stability under model errors. Although only done for a subset of possible plant variations, the approach could be extended to verify stability under a wider array of conditions.

The same Δ_P stability margin equation (6.65) was computed using measured experimental actuator loop gains and the mean loop gain. The total unknown quantities were computed from this experimentally-measured data just as in Fig. 6.44 in Sec. 6.5. This result is shown in Fig. 7.25. Since the uncertainty in the experimentally-measured path is less than the reciprocal of the complimentary sensitivity function, the closed-loop system is ensured to be stable. This compares favorably to the predicted uncertainty limitation in Fig. 6.44, although the measured uncertainty was much as 20dB higher than the modeling errors analytical limit.




Figure 7.24 Full-bandwidth measured performance with varying structure (actuator path)




Figure 7.25 Low-frequency robustness to structural variation

The robustness of the mid-frequency controller is ensured because of the generous phase margins allowed by the FxLMS algorithm. As discussed in Sec. 4.2, phase excursion at modal frequencies is ensured to be less than $\pm \frac{\pi}{2}$ for any structure with finite damping. Thus, as a worst-case, the collocated structure of the actuator/error sensor pair would allow a model that only fit to the mean phase rolloff. Adaptation time would suffer, but the steady-state performance would be equivalent. For the plant variations in this experimental setup, such extreme measures to ensure stability were not necessary. The same 40-state model was sufficiently accurate for all tests.

The actuator-path variation and Δ_P computation meant that the combination of all controllers would remain stable, but it does not indicate anything about the disturbance-rejection performance. This performance under structural variability for the can be seen in Fig. 7.26. The increase in actuator-path loop dynamics at frequencies such as 120Hz, 160Hz, and 170Hz can be seen somewhat in individual instances on the Bode plot. When looked at from a progressive RS standpoint in Fig. 7.27 however, the result is little change in aggregate performance.




Figure 7.26 Full-bandwidth measured performance with varying structure (disturbance path)




Figure 7.27 Full-bandwidth measured progressive RS performance with varying structure (disturbance path)

Finally, the unmodeled disturbance path was also measured under the variation in the experimental structure. Figure 7.28. Since this path was unmodeled to begin with, a change in its performance was an artifact of the other paths. The same spillover at 210Hz is present, as is the high-frequency performance above 600Hz.




Figure 7.28 Full-bandwidth measured performance with varying structure (unmodeled disturbance path)

7.5 Results

The results of the experimental test were very positive. The sequential-loop closure made for a robust implementation that did not rely on highly-accurate models from previous stages to close subsequent loops. The high-frequency controller produced the expected loop gain and performance, complete with predicted low-frequency spillover at actuator resonance. The low-frequency controller closed around this path was capable of reducing this spillover, and providing a very effective reduction in the modeled disturbance path. Unfortunately, this design choice leads to the tradeoff between modal characterization of the structure, and the zeros imposed

by individual actuator placements. Where the modeled-disturbance had zeros in the transfer function to error, the H_2 optimal controller would often inject spillover energy that can be seen at other disturbance locations.

Finally, the mid-frequency adaptive controller performed similar to what was predicted. The extension of its bandwidth to low frequencies was only partially predictable from the optimal filter computations of Sec. 6.3. The other limitations to performance were a function of the noise floor and the coherence of the measurement data.

The overall result was most clearly demonstrated in Fig. 7.19-7.20. The total performance is shown most clearly in Fig. 7.19 in both a Bode and broadband representation. The progression of the loop-closure sequence in Fig. 7.20 clearly shows how each controller added increases the performance in the designed frequency range.

CHAPTER 8

CONCLUSIONS AND CONTINUED RESEARCH

This chapter discusses the results of this dissertation research on wideband active vibration control on an uncertain resonant structure. The proposed research contributions outlined in Chapter 1 are described in Sec. 8.1. Each is briefly summarized, along with the conclusions drawn from the design and implementation. Suggestions for extensions to this work are discussed in Sec. 8.2

8.1 Contributions

The work presented in this dissertation primarily describes the synthesis of an *implementable* control strategy for wideband active vibration control. Every step in the synthesis process was determined by implementational constraints. Most of these were governed by the structure itself including experimentally-verified measurements of the extremely high order, modal density, low structural damping, structural variability, and non-invasive actuators/sensors performance. Other implementation constraints were given by the practical computational complexity that was available by the real-time digital control computer hardware. One could make the argument that the computational processing is irrelevant since it is always increasing, but there will *always* be a practical limit to the computational (and numerical) complexity that will be implementable in real-time. The contributions of this work included a synthesis procedure that considered all of these constraints. A combination of control techniques was used to extend the limitations of each to include a very wide aggregate bandwidth. The resulting controller configuration provided wideband disturbance rejection with minimal parasitic interaction between the different controllers.

An optimization procedure was developed to provide a means of choosing the configuration of the low-order feedback controller. This utilized a simplified, yet worst-case stability model of the experimental system including the finite actuator bandwidth. A number of cost functions were considered to rank the performance of candidate compensators. In the end, a band-limited summation of the aggregate disturbance rejection was chosen as the cost function most representative of the design goals. It was also discovered that boundary conditions on the optimization problem were the primary driving force on its result. These effectively provided an abstraction to the controller synthesis just as the frequency-weighting parameters of the H_2 design did in Sec. 4.3.

An analysis of the errors introduced by multiple levels of sampling was conducted. A modest bound on the continuous-time plant's maximum gain above the sampling rate was introduced to limit the maximum sampling error. This was extended to include contributions from the highest frequency controller, which was operated *above* the sampling rate of the digital controllers. Also, the errors introduced by decimation and interpolation to reduce computational requirements of the high-order control were quantified and found to numerically justify the assumption of minimal impact on the resulting closed-loop performance and stability. The careful analysis of all of these error sources justified the operation of the digital controllers well beyond the usual "5-10 times the highest frequency of interest" rule of thumb for digital control sample rate choice.

Finally, the synthesis and optimization procedure was experimentally validated on a test structure with an extremely high order. Over 1000 states were identified in the 50-5000Hz bandwidth alone. The progressive effect of closing each control loop provided additional performance over the previous configuration. The clearest quantification of this result is in Fig. 7.20 on page 182. Subsequent controller loop closure even experimentally demonstrated spillover reduction introduced by previous controllers.

The effects of unmodeled disturbance rejection were also experimentally measured. The high-order feedback clearly showed a design tradeoff between maximizing *modeled* performance and minimizing *unmodeled* spillover in Fig. 7.21. Even so, contribution of structural modes with common influences to both disturbances were reduced even with an unmodeled disturbance excitation.

8.2 Further Research

A research topic as multi-faceted as the one pursued in this work has a very broad choice of expansion possibilities. Just a few of the more incremental extensions to the proposed control strategy are mentioned here.

The analysis presented here was for a single channel of error minimization and control actuators. A multivariable extension to the work has already been investigated with encouraging results. By duplicating the control in all three axes (XYZ) as in Fig. 8.1, simultaneous reductions in all three channels were experimentally observed. An extension to include a fully-coupled, multivariable implementation of the H_2 low-frequency controller was also included in these experimental tests. The analysis of the multivariable version of this has been started, although has a number of modeling difficulties. In particular, the inclusion of all error sources including sampling, decimation, and interpolation makes it difficult to ensure stability analytically. This area has good potential for further study.

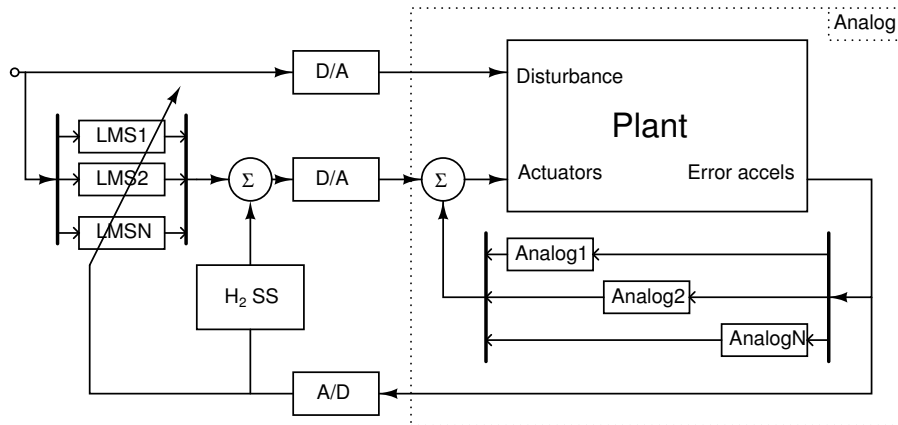


Figure 8.1 Possible multivariable extension

Another area that would be very interesting to pursue would be to remove the assumption of decoupled disturbances. Directly measuring a disturbance, rather than relying on a structurally decoupled, yet correlated, disturbance would widen the range of applications. By removing this assumption, adaptation path for the feed-forward control would no longer be free of high-frequency dynamics. Again, the nonlinear sampling errors and multivariable stability requirements would have to be investigated under this more general assumption to ensure stability with this implementation extension.

APPENDIX A

DIGITAL CONTROLLER

This section describes the implementation details of the digital controller portion of the experimental controller. It details the hardware, custom software, and off-the-shelf software decisions that were important to ensure adequate hard real-time performance.

A.1 Hardware

The digital computer used for the low and mid-frequency controller and data acquisition was a combination of standard, off-the-shelf personal computers. A 750 MHz Pentium 3 and a 2.4 GHz Pentium 4 machine were used in combination. The actual data acquisition hardware was a National Instruments PCI-6070E PCI card with 1.25 MS/s of 12-bit analog to digital conversion split between up to 16 single-ended channels. The output card was a National Instruments PCI-6713 PCI card, with eight, 12-bit digital to analog converters capable of outputting at 1MS/s on each. The software employed was a combination of Mathwork's xPC-target extension to Matlab's Real-Time Workshop (RTW) environment and custom C-code. The advantage of this arrangement was that it allowed for the control algorithm to be written in a high-level, yet highly optimizable programming language. At the same time, the Real-Time Workshop attribute of the implementation meant that this code could be very easily extended in an object-oriented way. The interface to the control code was "written" as a Simulink block diagram.

The code and interface were written and compiled on the 750MHz machine, called the *host* computer. This was then uploaded to the 2.4GHz machine (*target* computer) which was running a custom real-time operating system kernel from a diskette generated from the xPC Target utilities on the host. This combination allowed the real-time kernel on the target computer to have full access to all hardware without the overhead of running a full operating system.

Data collected on the target machine during controller execution was made available to the host machine during or after the run, depending on the operational mode.

A.2 Control Code

The biggest advantage of running the xPC Target environment was the ability to modularize the control code. Once a library of code was written, it could be easily extended to higher channel counts. Standard components such as drivers for the hardware, discrete-time transfer functions and state-space blocks were already included in the toolbox. Of course there was no “LMS” block in the pre-written toolbox, so this adaptive code had to be written. The disadvantage of xPC target environment for this was having to deal with the additional overhead of the object-oriented nature of the Simulink/RTW environment.

Although many other software features were available in the RTW environment, the fundamentally important ones to writing the control code were functionalized as the `UpdateStates()` function and the `CalculateOutputs()` function. These are analogous to the $\dot{x} = Ax + Bu$ and $y = Cx + Du$ equations in state-space framework, respectively. States which represent the simulation/controller’s internal states of a block are updated in the code once per sampling period by executing the `UpdateStates()` function. The output of the block is then calculated once per sampling period during the `CalculateOutputs()` function.

The problem with this was that in order to make the toolbox general-purpose and object-oriented, each state in a block that was to be persistent at the next time had to be passed as functional outputs from the `UpdateStates()` function into inputs to the `CalculateOutputs()` function. There was no provision for globally-shared states except for this message-passing format. For implementation of the LMS algorithm, this additional message-passing overhead was much too inefficient for real-time computations in excess of 10-20 states. To get around this limitation, a global, memory-mapped state variable structure was used to ensure unconditional access to the state variables from either function. Once this abstraction was made, computations of the LMS algorithm were possible up to 1024-points at a 10kHz sample rate. The disadvantage of this approach was that a *true* global variable had to be used. The hierarchy did not provide an obvious mechanism for limiting the scope of a persistent piece of memory to a specific instance of a block. Because of this, the custom code had to *also* be duplicated with unique global

memory variable names. Simply copying the block representing the code would not generate a new global variable name, but rather would access the same variable as the original. Careful header definitions greatly simplified this renaming requirement so it was not much of a concern at compile/run time.

Aside from the memory-mapped requirement, the algorithm of (4.17) and (4.18) were more or less computed directly. The final implementation of FxLMS as an LMS block can be seen in Fig. A.1. The control block was written with with the filter length (N), gradient gain (μ), and leakage parameter (λ) as input parameters to the S-function. The maximum length of the filter was adjustable at compile-time, but was run-time adjustable. A ring-buffer data structure for the filter coefficients was used internally to reduce the amount of memory copying that was required during every UpdateState iteration.

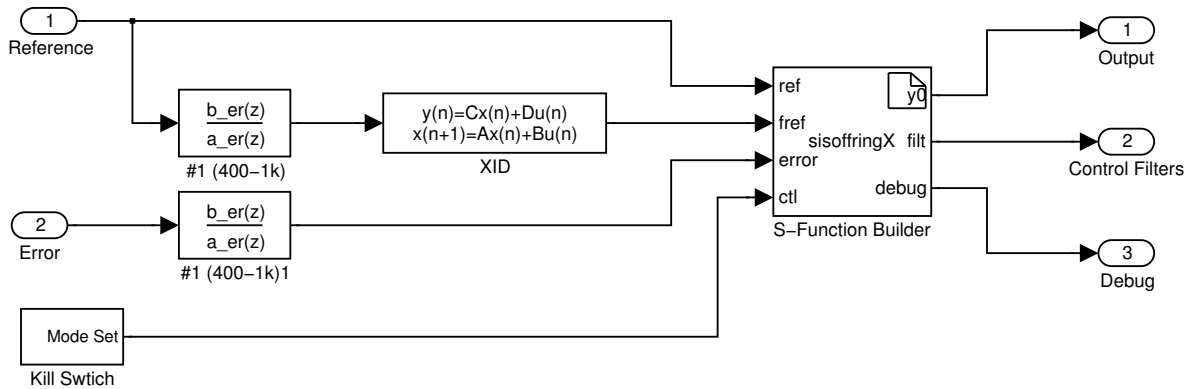


Figure A.1 LMS code block as Simulink object

Other implementation features of the control code were externally-accessible reference, filtered-reference, and error inputs. This allowed easy adjustment of the error weighting filters, and models of the actuator path. The output (y_0) was the direct output of the LMS filter and was suitable for directly adding to the digital controller summation with the low-frequency controller. The filter coefficients were made available out the #2 output in an interleaved fashion. Because of the overhead previously mentioned, it was infeasible to output all of the filter coefficients at all time steps. Rather, the control filter output was a 2x1 vector that included one of the filter coefficients and its index. The index (and corresponding filter coefficient) was incremented at the next time step. Since the evolution of the filters was slow by design of LMS, this was not viewed as a significant drawback. At only 256-point filters at 10kHz sample rate,

the entire filter was updated 40 times per second. Realistically, this output had only debugging value since the output filters were never directly used in the analysis, but they were available if desired.

Finally, a few other implementation features were added to the code. A hard mode switch was included to control the operation of the LMS block. The internals of this mode switch block can be seen in Fig. A.2. A single digital I/O pin was configured with a physical switch. This switch allowed cycling the LMS block through four distinct modes of operation: normal, disabled, reset, or freeze as described below:

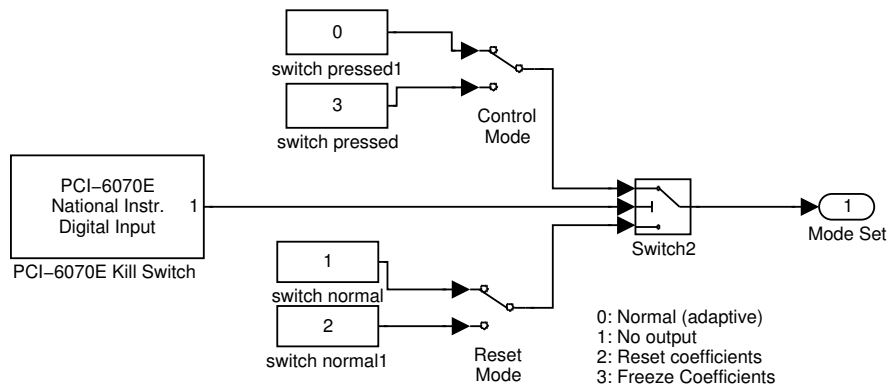


Figure A.2 Hard real-time mode switch as Simulink object

- **Normal:** Normal LMS operation using entered parameters for N , μ , and λ .
- **Disable:** Disabled the output of the LMS block and paused filter evolution.
- **Reset Coefficients:** Reset filter coefficients. This was by default zero, but could be changed to any pre-configured set of filter coefficients desired.
- **Freeze:** Stopped the evolution of the filter coefficients, but continue to provide filtered output. This mode was useful to compare different controller configurations with a *static* amount of adaption time allowed.

Although the single-bit switch only allowed for external selection of the LMS mode between one of two options, the manual switching in the Simulink diagram allowed for selection of these two modes. The primary reason for the hard switch was to allow modification of the real-time system when operating in free-run mode. The free-run mode was a reduced-flexibility

communication mode between the target and host computer that removed the host computer’s ability to change the target’s parameters at run-time. By removing this ability, the target computer had additional processing resources and reduced latency as compared to the normal, polled mode. When operating with a very high sampling rate or model order complexity, this mode was sometimes required to ensure hard real-time limits to latency.

One other bit of implementational detail that strays from the obvious was the decimation/interpolation blocks necessary for the low-frequency control. Figure A.3 shows the implementation of the decimated low-frequency control as done in Simulink, but does not fully describe the design choices required. There were two additional parameters in the rate-transition block that were important to have set correctly. One was “Ensure data integrity during data transfer,” and the other was “Ensure deterministic data transfer (maximum delay).” The hard-real-time nature of the RTW meant that data in memory was not necessarily guaranteed to be updated before or after the rate transition choice was made. This trade-off between deterministic behavior (double-buffering data and semaphore-control) and performance (copy read data immediately) was by default chosen in favor of determinism, and both options were checked. Unfortunately, this maximized the amount of delay in the rate transition block and guaranteed one entire *decimated* sample delay in the throughput. Un-setting both of these options reduce the decimation delay to what would provide the least delay. Otherwise, the discrete-time anti-aliasing, smoothing, and state-space blocks were implemented as normal Simulink/RTW tools.




Figure A.3 Decimated H_2 block as Simulink object

The final configuration of the controller (simplified somewhat for clarity) is shown in Fig. A.4. The custom code block for the LMS algorithm is all contained in the central block. All accelerometer and excitation logging was done within their respective blocks. Where necessary for hardware-in-the-loop identification discussed in Sec. 7.2 and 7.3, these blocks also contained copies of the relevant filtering and rate-changing blocks. The debugging scopes were an imple-

mentation aid to view the progress of the filter adaption and instantaneous RMS error on the target computer *while running* the control.

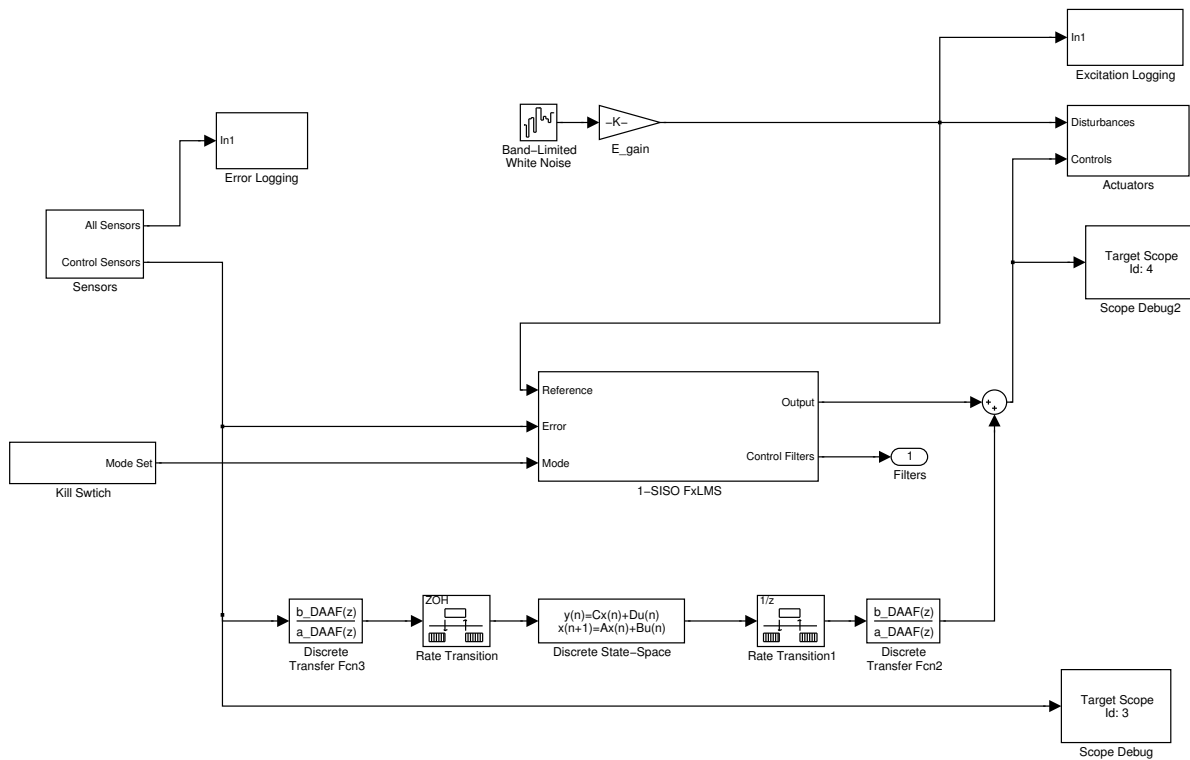


Figure A.4 Complete controller as Simulink object

APPENDIX B

ANALOG IMPLEMENTATION

The high-frequency synthesis of Sec. 4.2 and optimization of Chap. 5 provided the desired control-style transfer function of the high-frequency controller. What they did not provide were the translation between s -domain transfer functions and real, continuous-time analog signal processing. For this, an op-amp circuit was constructed that used the impedance relationships between frequency independent impedances ($Z_R = R$) and frequency-dependent impedances ($Z_C = \frac{1}{j\omega C}$).

Various undergraduate texts cover such op-amp circuits, but one of the more succinct in terms of *implementation* is [122]. The schematic of the final implemented circuit is shown in Fig. B.1.

The circuit follows very clearly with the desired transfer function consisting of two poles, one zero, and an aggregate gain. Although the design could have been done with only two op-amp stages, this setup had the advantage of providing some isolation between setting the different parameters. If everything was computed with a minimum of stages, changing any one value would alter other values as well. For testing purposes in particular, it was useful to allow the analog or digital control loop to be opened individually without affecting the other's gain.

The first stage provides both a zero and a pole and has the s -domain transfer function of (B.1).

$$G_1 = \left[\frac{-C_1}{C_{f_1}} \right] \frac{s + \frac{1}{R_1 C_1}}{s + \frac{1}{R_{f_1} C_{f_1}}} \quad (\text{B.1})$$




Figure B.1 Analog compensator schematic

The second stage needed only a single pole, so the input parallel capacitance was omitted in (B.2)

$$G_2 = \left[\frac{-1}{R_2 C_{f_2}} \right] \frac{1}{s + \frac{1}{R_{f_2} C_{f_2}}} \quad (\text{B.2})$$

The third (B.3) stage provides a simple, variable gain stage that isolates the high-frequency feedback from the summing junction (B.4).

$$G_3 = \frac{-R_{f_3}}{R_3} \quad (\text{B.3})$$

$$G_4 = \left[\frac{-R_{f_4}}{R_{4_a}} \right] v_{a_1} + \left[\frac{-R_{f_4}}{R_{4_b}} \right] v_b \quad (\text{B.4})$$

The summing junction adds in the signal provided from the low and mid-frequency controller implemented in the digital computer.

As a final result, the concatenation of all stages for the analog controller is (B.5), and for the digital controller (B.7)

$$G_a = \left[\frac{C_1 R_{f_3} R_{f_4}}{C_{f_1} R_2 C_{f_2} R_3 R_{4_a}} \right] \frac{\left(s + \frac{1}{R_1 C_1} \right)}{\left(s + \frac{1}{R_{f_1} C_{f_1}} \right) \left(s + \frac{1}{R_{f_2} C_{f_2}} \right)} \quad (\text{B.5})$$

$$= [121 \times 10^3] \frac{s + 2\pi 156}{(s + 2\pi 885)(s + 2\pi 4820)} \quad (\text{B.6})$$

$$G_b = \frac{-R_{f_4}}{R_{4_b}} \quad (\text{B.7})$$

The component values that were chosen had to simultaneously provide the desired aggregate gain, individual pole/zero corner frequencies, not emphasize non-ideal characteristics in the op-amp, and be standardized values where possible. They had to maximize propagation signal to noise ratio by ensuring that no intermediate stage had too little gain. Each stage was also limited to the power supply rails of $\pm 15\text{V}$. Table B.1 summarizes the final component choices.

Component	Value
R_1	6.8 k Ω
R_{f_1}	1.2 k Ω
R_2	3.3 k Ω
R_{f_2}	2.2 k Ω
R_3	100 k Ω
R_{f_3}	0-600 k Ω
R_{4_a}, R_{4_b}	12 k Ω
R_{f_4}, R_{x_4}	12 k Ω
$R_{x_1}, R_{x_2}, R_{x_3}$	10 k Ω
C_1, C_{f_1}	150 nF
C_{f_2}	15 nF

Table B.1 Analog component values

In the final configuration, the 0dB, -6dB, and -12dB gain settings in the experimental results settings of Fig. 7.1 were with R_{f_3} equal to 336k Ω , 168k Ω and 84k Ω , respectively. It should also be noted that the power amplifier was downstream from this circuit, and its gain also affected the aggregate loop gain. It was set to a fixed value of approximately 10, but left unchanged throughout the entire investigation.

REFERENCES

- [1] H. Nyquist, "Regeneration theory," Bell Systems Technical Journal, Tech. Rep., 1932.
- [2] H. S. Black, "Stabilizing feedback amplifiers," Bell Systems Technical Journal, Tech. Rep., 1934.
- [3] H. W. Bode, "Feedback amplifier design," Bell Systems Technical Journal, Tech. Rep., 1940.
- [4] R. E. Kalman, "Contributions to the theory of optimal control," *Bol. Soc. Math. Mexicana*, 1960.
- [5] P. Dorato, "A historical review of robust control," *IEEE Control Systems Magazine*, pp. 44–47, April 1987.
- [6] J. C. Doyle and G. Stein, "Multivariable feedback design: Concepts for a classical/modern synthesis," in *IEEE Transactions on Automatic Control*, vol. AC-26, February 1981, pp. 4–16.
- [7] A. A. Rodriguez and D. R. Carter, "Hierarchical HAC_∞/LAC vibration suppression for a flexible space telescope: SPICE," in *Proceedings of the American Control Conference*, 1994, pp. 686–960.
- [8] R. Gran and M. Rossi, "A survey of the large structures control problem," in *Proceedings of the 18th Conference on Decision & Control*, vol. 2, December 1979, pp. 1002–1007.
- [9] M. Balas, "Trends in large space structure control theory: Fondest hopes, wildest dreams," in *IEEE Transactions on Automatic Control*, vol. AC-27, June 1982, pp. 522–535.
- [10] L. Meirovitch, *Analytical Methods in Vibrations*. Macmillan Publishing Co., 1967.
- [11] C. M. Harris and C. E. Crede, Eds., *Shock and Vibration Handbook*. McGraw-Hill Book Company, 1976.
- [12] R. Bishop and D. Johnson, *The Mechanics of Vibration*. Syndics of the Cambridge University Press, 1979.
- [13] X. Li, R. K. Agarwal, and S.-P. Shue, "Optimal control and H_∞ filter for control of Timoshenko beam vibrations using piezoelectric material," in *Proceedings of the 37th IEEE Conference on Decision & Control*, 1998, pp. 1566–1571.
- [14] R. Clark, W. Saunders, and G. Gibbs, *Adaptive Structures Dynamics and Control*. John Wiley & Sons, Inc., 1998.
- [15] P. Gardonio and M. J. Brennan, "On the origins and development of mobility and impedance methods in structural dynamics," *Journal of Sound and Vibration*, 2002.

- [16] R. H. Lyon and R. G. DeJong.
- [17] R. S. Langley, P. J. Shorter, and V. Cotoni, "Predicting the response statistics of uncertain structures using extended versions of SEA," in *Inter-Noise Environmental Noise Control*, August 2005.
- [18] R. S. Langley, R. H. Lande, P. J. Shorter, and V. Cotoni, "Hybrid deterministic statistical modelling of built-up structures," in *Twelfth International Congress on Sound and Vibration*, July 2005.
- [19] R. S. Langley, V. Cotoni, and M. R. . F. Kidner, "Response statistics and variance for a single SEA subsystem: Theory and experiment," in *Vibro-Acoustics Users Conference Belgium*, 2003.
- [20] P. J. Shorter and R. S. Langley, "Vibro-acoustic analysis of complex systems," *Journal of Sound and Vibration*, 2004.
- [21] R. S. Langley, "Natural frequency statistics and universality," in *Twelfth International Congress on Sound and Vibration*, July 2005.
- [22] M. J. Balas, "Direct velocity feedback control of large space structures," *Journal of Guidance, Control, and Dynamics*, vol. 2, pp. 252–253, May–June 1979.
- [23] E. T. Falangas, J. Dworak, and S. Koshigoe, "Controlling plate vibrations using piezoelectric actuators," in *IEEE Control Systems Magazine*, vol. 14, Aug 1994, pp. 34–41.
- [24] R. L. Clark, "Accounting for out-of-bandwidth modes in the assumed modes approach: Implications on collocated output feedback control," in *Transactions of the ASME*, vol. 119, September 1997, pp. 390–395.
- [25] R. L. Clark, "Moving away from collocated control," *Journal of Sound and Vibration*, vol. 190, pp. 129–136, 1996.
- [26] S. M. Joshi, *Lecture Notes in Control and Information Sciences*. Springer-Verlag, 1989.
- [27] E. Sim and S. W. Lee, "Active vibration control of flexible structures with acceleration feedback," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 2, pp. 413–415, 1993.
- [28] J. M. Hyde and W. P. Seering, "Using input commands pre-shaping to suppress multiple mode vibration," in *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, 1991, pp. 2604–2609.
- [29] R. Gueler, A. von Flotow, and D. Vos, "Passive damping for robust feedback control of flexible structures," *Journal of Guidance, Control, and Dynamics*, vol. 16, July–August 1993.
- [30] H. Pota, S. Moheimani, and M. Smith, "Resonant controllers for flexible structures," in *Proceedings of the 38th Conference on Decision & Control*, December 1999, pp. 631–636.
- [31] J. Ronald L. Spangler and S. R. Hall, "Broadband active structural damping using positive real compensation and piezoelectric simultaneous sensing and actuation," *Smart Materials & Structures*, vol. 3, pp. 448–458, 1994.
- [32] A. J. Bronowicki, B. S. Abhyankar, and S. F. Griffin, "Active vibration control of large optical space structures," *Smart Materials & Structures*, vol. 8, no. 6, pp. 740–752, 1999.

- [33] A. J. Bronowicki, "A layered vibration control strategy for space telescopes," in *Smart Structures and Materials 2003: Smart Structures and Integrated Systems*, M. Shao, Ed., 2003.
- [34] R. Alkhatib and M. F. Golnaraghi, "Active structural vibration control: A review," in *The Shock and Vibration Digest*, vol. 35, Save Publications, September 2003, pp. 367–383.
- [35] A. Arbel and N. K. Gupta, "Robust collocated control for large flexible space structures," *Journal of Guidance, Control, and Dynamics*, vol. 4, pp. 480–486, Sept-Oct 1981.
- [36] M. Balas, "Feedback control of flexible systems," in *IEEE Transactions on Automatic Control*, vol. AC-23, August 1978, pp. 673–679.
- [37] G. Balas and J. C. Doyle, "Robustness and performance trade-offs in control design for flexible structures," in *IEEE Transactions on Control Systems Technology*, vol. 2, December 1994, pp. 352–361.
- [38] J. C. Doyle, K. Glover, P. P. Khargonekar, and B. A. Francis, "State-space solutions to standard H_2 and H_∞ control problems," in *IEEE Transactions on Automatic Control*, vol. 34, August 1989, pp. 831–847.
- [39] M. Hatayama, F. Matsuno, and Y. Sakawa, "Modeling and robust control of flexible solar array paddles," in *Proceedings of the 1993 Conference on Industrial Electronics*, vol. 3, November 1993, pp. 2039–2044.
- [40] IEEE, "Special issue on the LQG problem," in *IEEE Transactions on Automatic Control*, vol. AC-16, December 1971.
- [41] E. Tse, "On the optimal control of stochastic linear systems," in *IEEE Transactions on Automatic Control*, vol. AC-16, December 1971, pp. 776–785.
- [42] M. Athans, "The role and use of the stochastic linear-quadratic-gaussian problem in control system design," in *IEEE Transactions on Automatic Control*, vol. AC-16, December 1971, pp. 529–552.
- [43] D. C. Youla, J. Joseph J. Bongiorno, and H. A. Jabr, "Modern Wiener-Hopf design of optimal controllers," in *IEEE Transactions on Automatic Control*, vol. AC-21, 1976.
- [44] G. Stein and M. Athans, "The LQR/LTR procedure for multivariable feedback control design," in *IEEE Transactions on Automatic Control*, vol. AC-32, February 1987, pp. 105–114.
- [45] H. Kwakernaak, " H_2 -optimization theory and applications to robust control design," *Annual Reviews in Control*, vol. 26, pp. 45–56, 2002.
- [46] J. Doyle, "Robust and optimal control," in *Proceedings of the 35th Conference on Decision & Control*, December 1996, pp. 1595–1598.
- [47] R. Smith, Cheng-Chu, and J. Fanson, "The design of H_∞ controllers for an experimental non-collocated flexible structure problem," *IEEE Transactions on Control Systems Technology*, vol. 2, June 1994.
- [48] M. Karkoub, G. J. Balas, and K. Tamma, "Collocated and noncollocated control design via μ -synthesis for flexible manipulators," in *Proceedings of the American Control Conference*, June 1995, pp. 3321–3325.

- [49] R. Fraanje, M. Verhaegen, N. Doelman, and A. Berkhoff, “MIMO H_2 optimal and robust feedback controller estimation for a vibrating plate using subspace model identification,” Sound and Vibration UT-TNO, Robust Active Control, Tech. Rep., 2003.
- [50] S. Joshi and P. Maghami, “Robust dissipative compensators for flexible spacecraft control,” in *IEEE Transactions on Aerospace and Electronic Systems*, vol. 28, 1992, pp. 768–774.
- [51] D. S. Bernstein and W. M. Haddad, “LQG control with an H_∞ performance bound: A riccati equation approach,” *IEEE Transactions on Automatic Control*, vol. 34, March 1989.
- [52] F. Paganini, “Frequency domain conditions for robust H_2 performance,” *IEEE Transactions on Automatic Control*, vol. 44, January 1999.
- [53] G. J. Balas, A. K. Packard, M. G. Safonov, and R. Y. Chiang, “Next generation of tools for robust control,” in *Proceedings of the 2004 American Control Conference*, 2004, pp. 5612–5615.
- [54] H. R. Pota and I. R. Petersen, “Robust control of an acoustic cavity,” in *Proceedings of the 2004 IEEE International Conference on Control Applications*, 2004, pp. 771–775.
- [55] I. R. Petersen and H. R. Pota, “Experiments in feedback control of an acoustic duct,” in *Proceedings of the 2000 IEEE International Conference on Control Applications*, 2000, pp. 261–266.
- [56] G. Rafaely and S. J. Elliott, “ H_2/H_∞ active control of sound in a headset: Design and implementation,” in *IEEE Transactions on Control Systems Technology*, vol. 7, January 1999, pp. 79–84.
- [57] S.-H. Wang and E. J. Davison, “On the stabilization of decentralized control systems,” in *IEEE Transactions on Automatic Control*, vol. AC-18, October 1973, pp. 473–478.
- [58] R. Saeks, “On the decentralized control of interconnected dynamical systems,” in *IEEE Transactions on Automatic Control*, vol. AC-24, April 1979, pp. 269–271.
- [59] A. Linnemann, “Decentralized control of dynamically interconnected systems,” in *IEEE Transactions on Automatic Control*, vol. AC-29, November 1984, pp. 1052–1054.
- [60] G. West-Vokovich, E. Davidson, and P. Hughes, “The decentralized control of large flexible space structures,” in *IEEE Transactions on Automatic Control*, vol. 29, October 1984, pp. 866–879.
- [61] P. Dix, U. Ozguner, and R. W. Gordon, “Decentralized control experiments on a truss structure,” in *Proceedings of the 29th IEEE Conference on Decision & Control*, 1990, pp. 2001–2006.
- [62] J. P. Corfmat and A. S. Morse, “Decentralized control of linear multivariable systems,” in *Automatica*, vol. 12, 1976, pp. 479–495.
- [63] J. N. Aubrun, “Theory of the control of structures by low-authority controllers,” *Journal of Guidance and Control*, vol. 3, pp. 444–451, Sept-Oct 1980.
- [64] T. W. C. Williams and P. J. Antsaklis, “Limitations of vibration suppression in flexible space structures,” in *Proceedings of the 28th Conference on Decision and Control*, Dec 1989, pp. 2218–2222.

- [65] M. G. Safonov, R. Y. Chiang, and H. Flashner, " H_∞ control synthesis for a large space structure," *Journal of Guidance and Control*, vol. 14, pp. 513–520, May-June 1991.
- [66] D. C. Hyland, J. L. Junkins, and R. W. Longman, "Active control technology for large space structures," *Journal of Guidance, Control, and Dynamics*, vol. 16, pp. 801–821, Sept-Oct 1993.
- [67] D. J. Phillips, J. A. King, L. D. Davis, and D. C. Hyland, "High-performance active control of dynamically complex space structures," *Journal of Guidance, Control, and Dynamics*, vol. 19, pp. 1071–1080, 1996.
- [68] R. J. Benhabib, R. P. Iwens, and R. L. Jackson, "Stability of large space structure control systems using positivity concepts," *Journal of Guidance, Control, and Dynamics*, vol. 4, Sept-Oct 1981.
- [69] M. McLaren and G. Slater, "Robust multivariable control of large space structures using positivity," *Journal of Guidance, Control, and Dynamics*, vol. 10, July-August 1987.
- [70] S. M. Joshi, P. G. Maghami, and A. G. Kelkar, "Design of dynamic dissipative compensators for flexible space structures," in *IEEE Transactions on Aerospace and Electronic Systems*, vol. 31, October 1995, pp. 1314–1324.
- [71] R. Lozano-Leal and S. M. Joshi, "On the design of dissipative LQG-type controllers," in *Proceedings of the 27th IEEE Conference on Decision & Control*, 1988, pp. 1645–1646.
- [72] R. Leal and S. Joshi, "On the design of dissipative LQG-type controllers," in *Proceedings of the 27th Conference on Decision & Control*, December 1988, pp. 1645–1646.
- [73] A. Bhaya and C. Desoer, "On the design of large flexible space structures (LFSS)," *IEEE Transactions on Automatic Control*, vol. AC-30, pp. 1118–1120, November 1985.
- [74] S. Gupta and S. M. Joshi, "State space characterization and robust stabilization of dissipative LTI systems," in *Proceedings of the American Control Conference*, June 1995, pp. 3616–3619.
- [75] W. M. Haddad and D. S. Bernstein, "Robust stabilization with positive real uncertainty: Beyond the small gain theorem," in *Proceedings of the 29th Conference on Decision and Control*, December 1990.
- [76] P. M. Clarkson and P. R. White, "Simplified analysis of the LMS adaptive filter using a transfer function approximation," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-35, July 1987, pp. 987–993.
- [77] S. J. Elliott, I. M. Stothers, and P. A. Nelson, "A multiple error LMS algorithm and its application to the active control of sound and vibration," in *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-35, October 1987, pp. 1423–1434.
- [78] M. Rupp and A. H. Sayed, "A time-domain feedback analysis of filtered-error adaptive gradient algorithms," in *IEEE Transactions Signal Processing*, vol. 44, June 1996, pp. 1428–1439.
- [79] P. Darlington and G. Xu, "Equivalent transfer functions of minimum output variance mean-square estimators," in *IEEE Transactions on Signal Processing*, vol. 39, July 1991, pp. 1674–1677.

- [80] R. Venugopal and D. S. Bernstein, "Adaptive disturbance rejection using AR-MARKOV/toeplitz models," *IEEE Transactions on Control Systems Technology*, vol. 8, pp. 257–269, March 2000.
- [81] P. R. Kumar and M. Campi, "Adaptive control." University of Illinois at Urbana-Champaign ECE 417 Adaptive Control Course Notes, August 1999.
- [82] A. H. Sayed and T. Kailath, "A state-space approach to adaptive RLS filtering," *IEEE Signal Processing Magazine*, pp. 18–60, July 1994.
- [83] S. J. Elliott and P. A. Nelson, "Active noise control," *IEEE Signal Processing Magazine*, pp. 12–35, 1993.
- [84] B. Widrow and S. Stearns, *Adaptive Signal Processing*. Prentice Hall, 1-ed., March 15 1985.
- [85] C. R. Fuller and A. H. von Flotow, "Active control of sound and vibration," *IEEE Control Systems Magazine*, vol. 15, pp. 9–19, December 1995.
- [86] E. Bjarnason, "Analysis of the Filtered-X LMS algorithm," in *IEEE Transactions on Speech and Audio Processing*, vol. 3, November 1995, pp. 504–514.
- [87] C. C. Boucher, S. J. Elliott, and P. A. Nelson, "Effect of errors in the plant model on the performance of algorithms for adaptive feedforward control," *IEE Proceedings of Radar and Signal Processing*, vol. 138, pp. 313–319, Aug 1991.
- [88] S. Dasgupta, J. S. Garnett, and J. C. Richard Johnson, "Convergence of an adaptive filter with signed filtered error," in *IEEE Transactions on Signal Processing*, vol. 42, April 1994, pp. 946–950.
- [89] P. A. Nelson and S. J. Elliott, *Active Control of Sound*. Academic Press, 1992.
- [90] V. H. Nascimento and A. H. Sayed, "An unbiased and cost-effective leaky-LMS filter," in *IEEE Transactions on Signal Processing*, vol. 47, December 1999, pp. 3261–3276.
- [91] V. H. Nascimento and A. H. Sayed, "An unbiased and cost-effective leaky-LMS filter," in *Conference Record of the Thirtieth Asilomar Conference on Signals, Systems and Computers*, November 1996, pp. 1078–1082.
- [92] K. A. Mayyas and T. Aboulnasr, "Leaky LMS: A detailed analysis," in *IEEE International Symposium on Circuits and Systems*, vol. 2, 1995, pp. 1255–1258.
- [93] R. Sharm, W. A. Sethares, and J. A. Bucklew, "Asymptotic analysis of stochastic gradient-based adaptive filtering algorithms with general cost functions," in *IEEE Transactions on Signal Processing*, vol. 44, September 1996, pp. 2186–2194.
- [94] M. Rupp and A. H. Sayed, "Robust FxLMS algorithms with improved convergence performance," in *IEEE Transactions on Speech and Audio Processing*, vol. 6, January 1998, pp. 78–85.
- [95] D. P. Das and G. Panda, "Active mitigation of nonlinear noise processes using a novel filtered-s LMS algorithm," in *IEEE Transactions on Speech and Audio Processing*, vol. 12, May 2004, pp. 313–322.
- [96] S. J. Elliott, "Filtered reference and filtered error LMS algorithms for adaptive feedforward control," *Mechanical Systems and Signal Processing*, pp. 769–781, 1998.

- [97] S. H. Oh and Y. Park, "Active noise control algorithm using IIR-based filter," *Journal of Sound and Vibration*, vol. 231, pp. 1296–1412, 2000.
- [98] C. Mosquera, J. A. Gomez, F. Perez, and M. Sobreira, "Adaptive IIR filters for active noise control," in *Sixth international congress on sound and vibration*, July 1999.
- [99] C. M. Papenfuss, M. R. F. Kidner, C. R. Fuller, and W. T. Bauman, "Experimental implementation of wideband active vibration control," in *International Symposium on Active Noise and Vibration Control*, 2004.
- [100] M. Phan, L. G. Horta, J.-N. Juang, and R. W. Longman, "Improvement of observer/kalman filter identification (OKID) by residual whitening," in *AIAA Guidance, Navigation and Control Conference*, August 1992, pp. 344–353.
- [101] L. G. Horta, J.-N. Juang, and C.-W. Chen, "Frequency domain identification toolbox," National Aeronautics and Space Administration Technical Memorandum 109039, Tech. Rep., September 1996.
- [102] J.-N. Juang, "State-space system realization with input-and-output-data correlation," National Aeronautics and Space Administration, Tech. Rep., April 1997.
- [103] T. Mathworks, *Control System Toolbox for use with Matlab*. The Mathworks, Inc., June 2005.
- [104] G. F. Franklin, J. D. Powell, and M. L. Workman, *Digital Control of Dynamic Systems*. Addison Wesley, 1998.
- [105] T. Chen, K. Huang, and C. Liaw, "High-frequency switching-mode power amplifier for shaker armature excitation," in *IEEE Proceedings-Electric Power Applications*, vol. 144, Nov 1997, pp. 415–422.
- [106] F. Fahy, *Foundations of Engineering Acoustics*. Academic Press, 2001.
- [107] D. A. Bies and C. H. Hansen, Eds., *Engineering Noise Control*. Unwin Hyman Ltd, 1988.
- [108] C. Fuller, *Encyclopedia of Acoustics*. John Wiley & Sons, Inc, 1997.
- [109] A. Preumont, *Vibration Control of Active Structures*. Kluwer Academic Publishers, 2002.
- [110] G. D. Margin, "On the control of flexible mechanical systems," Ph.D. dissertation, Stanford University, Stanford University, 1978.
- [111] F. Matsuno and T. Ohno, "Exponential stabilization of vibration for a large space structure with distributed and lumped flexibility," in *Proceedings of the 38th Conference on Decision & Control*, 1999, pp. 657–662.
- [112] S. J. Elliott, P. Gardonio, T. C. Sors, and M. J. Brennan, "Active vibroacoustic control with multiple local feedback loops," *Journal of the Acoustical Society of America*, pp. 908–915, February 2002.
- [113] T. Mathworks, *Robust Control Toolbox for use with Matlab*. The Mathworks, Inc., June 2001.
- [114] M. G. Safonov, A. J. Laub, and G. L. Hartmann, "Feedback properties of multivariable systems: The role and use of the return difference matrix," in *IEEE Transactions on Automatic Control*, vol. AC-26, February 1981, pp. 47–65.

- [115] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*. Englewood Cliffs, New Jersey 06632: Prentice Hall, 1989.
- [116] J. Hong, J. C. Akers, R. Venugopal, M.-N. Lee, A. G. Sparks, P. D. Washabaugh, and D. S. Bernstein, "Modeling, identification and feedback control of noise in an acoustic duct," *IEEE Transactions on Control Systems Technology*, vol. 4, pp. 283–291, May 1996.
- [117] J. Hong and D. S. Bernstein, "Bode integral constraints, collocation, and spillover in active noise and vibration control," *IEEE Transactions on Control Systems Technology*, vol. 6, pp. 111–120, January 1998.
- [118] D. Babic, V. Lehtinen, and M. Renfors, "Discrete-time modeling of polynomial-based interpolation filters in rational sampling rate conversion," in *Proceedings of the 2003 International Symposium on Circuits and Systems*, vol. 4, May 2003, pp. IV–321–IV–324.
- [119] I. Kollar, G. Franklin, and R. Pintelon, "On the equivalence of z-domain and s-domain models in system identification," in *IEEE Instrumentation and Measurement Technology Conference*, June 1996, pp. 14–19.
- [120] K. J. Astrom, R. L. Dailey, E. J. Davison, G. F. Franklin, H. Khalil, F. L. Lewis, L. Ljung, M. K. Maasten, R. A. Shoureshi, and D. D. Sworder, *Modern Control Systems*. Institute of Electrical and Electronics Engineers, 1995.
- [121] S. V. Narasimhan, S. Veena, H. Loksha, and S. S. Shankarling, "Algorithms for active noise control and their performance," in *International Conference on Smart Materials Structures and Systems*, July 2005, pp. 77–85.
- [122] P. Horowitz and W. Hill, *The Art of Electronics*. Cambridge University Press, 2-ed., 1989.

VITA

Cory Marvin C. Papenfuss was born November 28, 1975 in Dubuque, Iowa to Marvin and Linda Papenfuss. Cory has two siblings, Tracey and Lucas. Cory lived with his family in Dubuque, Iowa until 1992, except for the academic year 1986-1987, when the entire family lived in Fairbanks, Alaska. In 1992, the family moved to Juneau, Alaska. After graduating from Juneau Douglas High School in 1994, he began his undergraduate studies in electrical engineering at the University of Alaska-Fairbanks. During the summer of 1997 and through the 1997-1998 academic year, he worked as an undergraduate researcher in the Alaska Space Grant Program's Student Rocket Project in Fairbanks designing and constructing instrumentation and power conversion circuitry.

After earning a bachelor's degree in electrical engineering in 1998, he started his graduate education at the University of Illinois at Urbana-Champaign. While studying there, he became interested in the multi-disciplinary aspect of engineering control systems and power electronics. He was the Electrical Engineering Design Team Leader for the University of Illinois FutureCar project, which was a prototype hybrid diesel-electric vehicle. He earned a master's degree in electrical engineering in 2001 with thesis entitled, *Control and Instrumentation for Power Management in a Hybrid Electric Vehicle*.

He continued his graduate education at Virginia Polytechnic Institute and State University in 2001 with further research in applied control systems. After working on a number of projects in the Center for Power Electronic Systems, he began studying active vibration control at the Vibrations and Acoustics Lab in the Mechanical Engineering Department in 2002. He finished his research and dissertation and graduated with a Ph.D. in electrical engineering in 2006.

Other noteworthy accomplishments by the author include a private pilot certificate earned in 2002, followed with an instrument rating in 2003. In addition to many hours of flying in the

Lower 48 states, his flying experience includes a cross-country flight from Blacksburg, Virginia to Juneau, Alaska in a single-engine aircraft during the summer of 2005.