Title

Disturbance Rejection and System Identification Using Multi-Channel Adaptive Filtering and Receding-Horizon Control

Author

Tsuchiya, Nolan

Publication Date

2015

Peer reviewed|Thesis/dissertation

University of California

Los Angeles

Disturbance Rejection and System Identification Using Multi-Channel Adaptive Filtering and Receding-Horizon Control

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mechanical Engineering

by

Nolan Eizo Tsuchiya

2015

Abstract of the Dissertation

Disturbance Rejection and System Identification Using Multi-Channel Adaptive Filtering and Receding-Horizon Control

by

Nolan Eizo Tsuchiya

Doctor of Philosophy in Mechanical Engineering

University of California, Los Angeles, 2015

Professor James S. Gibson, Chair

The research presented in this dissertation explores the application of a receding-horizon adaptive controller to increase the disturbance rejection bandwidth far beyond the capabilities of a classical LTI controller. The receding-horizon controller is applied to two laboratory experiments to analyze and test various aspects of the controller. In the *laser beam steering experiment* a complex broadband disturbance is significantly reduced and the laser spot is held tightly on the target. Control penalty parameters are analyzed and frequency weighting is implemented in the single-input-single-output (SISO) experiment.

In the *magnetic bearing experiment* a steel shaft is magnetically levitated and rotated to speeds up to 7200 rpm. The emphasis in this experiment is to test the receding-horizon controller's ability to reject time-varying disturbances and to operate in a multi-input-multi-output (MIMO) configuration. Additionally, the forgetting factor parameter, unique to the recursive least-squares (RLS) algorithm employed by the controller, is examined and analyzed in the context of disturbances with varying statistics. Furthermore, in a collaboration project, a combination of the receding-horizon adaptive controller and

a peak resonator controller is applied in a parallel structure to reduce a complex disturbance. Independence of the controllers is demonstrated and it is shown that the parallel combination is capable of reducing the disturbance to a greater degree than either controller applied individually.

Finally, many advanced controllers including the receding-horizon adaptive controller require an accurate plant model in order to operate effectively. Currently, this plant model must be generated by a one-time, offline system identification procedure in which the system is driven with a white noise sequence with no disturbance present. The collected input/output data is batch-processed at the end of the data acquisition period to generate a plant model which can then be used by the adaptive controller. Because this scenario is not always feasible in practice, a novel system identification procedure in which plant models can be generated in real-time and in the presence of a disturbance is analyzed. The *online system identification* process successfully demonstrates a real-time identification of a complex plant in the presence of a broadband disturbance. The resulting plant model that is generated is comparable in accuracy to the state-of-the-art subspace identification methods.

The dissertation of Nolan Eizo Tsuchiya is approved.

Lieven Vandenberghe

Tetsuya Iwasaki

Tsu-Chin Tsao

James S. Gibson, Committee Chair

University of California, Los Angeles

2015

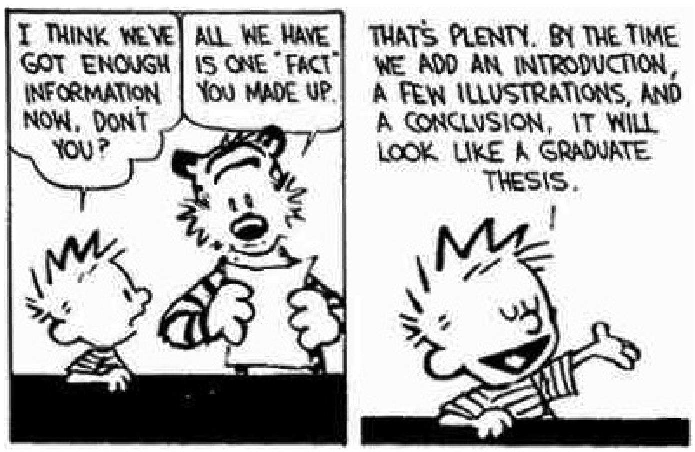

- Bill Watterson, *Calvin and Hobbes* [1]

[1]Used with permission from Universal Uclick / AMU Reprints

Table of Contents

Acknowledgments

I would like to express a sincere appreciation for all who have helped me along my path to the PhD. I am extremely grateful to my committee chair, Professor Steve Gibson, whose academic brilliance is sharply contrasted by his down-to-earth demeanor and comical wit. His continual support and guidance has been instrumental to the completion of my degree. I am also extremely thankful for the financial support for this research provided by the Air Force Office of Scientific Research under grant FA9550-12-1-0375 and by the Air Force Research Laboratory under grants 4-442561-GJ-21264 and 4-442561-GJ-21113. Additionally, I would like to thank my committee members, Professor Tsao, Professor Iwasaki, and Professor Vandenberghe for their guidance on my research. Their input has helped to steer the direction of my research and provide necessary and constructive criticism along the way. I would also like to thank the members of Professor Tsao's Mechatronics lab, who were a constant resource for me on any and all hardware matters. Specifically, Dr. Chris Kang, the resident guru on the magnetic bearing offered invaluable support on my experimental work with the bearing. Many of the results in Chapter 4 and all of the collaborative work in Chapter 5 would have taken significantly longer or would have been impossible were it not for Chris.

During my graduate studies, I also enjoyed a strong support network outside of the laboratory setting. My parents, who continue to be supportive of any path I choose, were a constant blessing for me and have always been vital to my success. My friends and family, who I can always count on for a fun time, have also provided me with a necessary outlet for periodically stepping away from research. Finally, I would like to acknowledge my wife, Karen, whose unwavering love and support over the past 5 years have kept me on track when I would have surely derailed otherwise. I can say with full confidence that I would not have made it this far, academically or otherwise, without her.

Vita

1984	Born, Alameda, California, USA.
2006	B.S. (Mechanical Engineering), University of California, Berkeley.
2009	Professional Engineer (PE) License, Mechanical Engineering, California
2006–2010	• Systems Engineer, Siemens Building Technologies, Hayward, CA. • Design Engineer, ACCO, San Leandro, CA, • Application Engineer, Syserco Inc., Fremont, CA
2012	M.S. (Mechanical Engineering), University of California, Los Angeles.
2011-2014	• Instructor (Summer 2011), UCLA BREES Program, MAE101, *Statics and Strength of Materials* • Teaching Assistant (Fall 2011), UCLA, ENGR113, *Product Strategy* • Instructor (Fall 2012, Winter 2013, Winter 2014), *UCLA Engineering MATLAB Workshop* • Teaching Assistant (Winter 2012, Spring 2012), UCLA, MAE107, *Modeling and Analysis of Dynamics Systems* • Teaching Associate (Winter 2013), UCLA, MAE171A, *Feedback and Control Systems*
2014–Present	Lecturer (Winter 2014, Fall 2014), Cal Poly Pomona, ME439/L, *Control of Mechanical Systems*

CHAPTER 1

Introduction

All mechanical systems are subject to external disturbances which can adversely affect the system's performance if left unaccounted for. Disturbances may come in many forms including: mechanical vibrations, turbulence, and sensor / actuator noise and they can be narrowband or broadband across the frequency spectrum as well as time-varying in nature. For the class of control problems known as *Disturbance Rejection*, the goal of the controller is to reduce the effects of these disturbances and maintain favorable system performance. The counterpart to this most general classification of control problems is known as *Reference Tracking*, however, for the experiments in this dissertation, the primary control objective is disturbance rejection and a receding-horizon adaptive controller is used to address a variety of complex disturbances.

One of the primary benefits of using an adaptive control scheme is to extend the disturbance rejection bandwidth beyond the relatively limited bandwidths achieved by classical linear time-invariant (LTI) feedback controllers. While both least-mean-squares (LMS) [1–3] and recursive-least-squares (RLS) [4–8] adaptive filters have been used to minimize output error variance, this research is based on an RLS adaptive lattice filter that predicts a specific linear combination of future disturbances, determined by the receding-horizon performance index discussed at length in [9] and [10], as well as in Chapter 2.

While minimum-variance adaptive control based on an RLS adaptive filter has been applied in practice [6], receding-horizon adaptive control (based on the same RLS filter) offers a wider range of tuning capabilities, including the ability to more easily handle coupled-channel systems. The receding-horizon controller also enjoys a more natural method of incorporating frequency weighting. A form of frequency weighting has been previously used to reduce sensitivity to high frequency noise [6], but receding-horizon controller design allows for this functionality in a more direct and flexible way. A similar receding-horizon adaptive controller has also been applied in an adaptive optics experiment [9] to correct aero-optical wavefront error. In that experiment, however, the challenge was more rooted in the complex optics and volume of control channels (30+) and less in the plant dynamics, which were simple and modeled by many single-input-single-output (SISO) transfer functions. Both experimental setups in this research have significantly more complex plant dynamics and the challenge of dealing with significant plant modeling error is addressed by the receding-horizon adaptive controller for the first time.

From a broader perspective, the receding-horizon (model predictive) control studied in this research is used in a fundamentally different way than is has been in the past. In practice, the most common application of receding-horizon control is in process control, where it is important to avoid violating constraints on the control command and/or certain terms in the state vector. By applying receding-horizon control in these applications, the controller can determine whether a constraint will be violated before it occurs and adjust accordingly to remain within an acceptable operating region. The application of receding-horizon adaptive control used throughout this research is based on formulating the algorithm as a disturbance rejection problem, where the adaptive portion is used to predict a linear combination of a fixed number of future disturbances.

This dissertation focuses on two primary avenues corresponding to two distinct experimental setups. In the beam steering experiment, the plant is comprised of two decoupled single-input-single-output (SISO) plants, but the adaptive controllers are confronted with a complex mix of narrow and broadband disturbances, resembling actual disturbances encountered in practice. These disturbances are injected experimentally by a dedicated disturbance mirror, but are not available to the controller, which only has access to the measured output. In contrast to the beam steering experiment, the magnetic bearing experiment possesses inherent coupling between two pairs of coupled channels and lends itself to a multi-input-multi-output (MIMO) approach. While this system is more complex, the disturbance in this case comes primarily from the shaft imbalance during spinning and manifests in the output as a single, dominant frequency. Additional experiments were conducted to test the performance when subjected to a time-varying disturbance as well as an additional broadband disturbance component. Both avenues apply the receding-horizon adaptive controller discussed in [9] and [10], but each experiment highlights different benefits and limitations of the adaptive controller. In both cases, plant modeling error is more prevalent at high frequencies, and can cause problems with performance and stability.

Chapter 2 describes the adaptive control algorithm in detail and highlights some of its key features, including the method of frequency weighting used. Chapter 3 presents the results from applying the receding horizon adaptive controller to a 2-axis laser beam steering experiment. Some important elements discussed in this section include: a detailed discussion on the plant nonlinearities, a study on the effect of the adaptive filter order, and a study on the effects of the weighting parameters, R_u and R_f. Chapter 4 presents results from applying a modified version of the receding-horizon controller to a magnetically-levitated bearing system. Because this open-loop unstable multi-channel plant is inherently challenging to control, a detailed discussion on the plant identification is provided, including an analysis on the actuator saturation which contributes to the system's nonlinearities. In this section, the adaptive controller is challenged to reject time-varying disturbances, as well as operate in a multi-input-multi-output (MIMO) framework. A study on the forgetting factor λ is also conducted to analyze the controller's efficacy in handling non-stationary disturbances. Three distinct control architectures are also studied to reveal the ability of the adaptive controller to handle severely coupled channels. Chapter 5 presents the results from a collaborative control effort in which the receding-horizon adaptive controller along with a peak resonator work in parallel to reject complex disturbances on the magnetic bearing system. Independence of the controllers is demonstrated and it is shown that this parallel combination control structure is capable of rejecting the disturbances to a greater degree than either controller operating alone.

In Chapter 6, a novel method of *online* system identification is analyzed. The algorithm constructs a plant model in real-time, in the presence of a disturbance. This is essential to this research because the adaptive controller requires a plant model of the system it controls. Prior to this online method of identification, a one-time offline identification was usually performed under ideal circumstances with no disturbances present. This is often impractical in application, as engineers do not have the luxury of obtaining a clean, noise-free identification prior to closing an adaptive control loop. The online identification method eliminates the need for an offline procedure and is capable of generating plant models in real-time and with disturbances acting on the system.

CHAPTER 2

Receding-Horizon Adaptive Control

2.1 Overview

Fundamentally, the receding-horizon control algorithm used throughout this research is based on the linear-quadratic regulator (LQR) optimal control structure which generates an optimal control law, $u(t)$, based on minimizing a given cost function. Figure 2.1 illustrates the receding-horizon control structure. The $G(z)$ block has the state-space representation:

$$x_p(t+1) = A_p x_p(t) + B_p u(t), \tag{2.1}$$

$$e(t) = C_p x_p(t) + w(t) \tag{2.2}$$

which represents the plant, where t is the sample number, $x_p(t)$ is the state vector, $e(t)$ is the measured output, $w(t)$ is the disturbance signal, and A_p, B_p, C_p, D_p are the state-space matrices with the direct feedthrough term equal to zero $D_p = 0$. It can either represent an open-loop or closed-loop plant, however, in both experimental applications studied so far, $G(z)$ represents a closed-loop system consisting of an open-loop system with a classical LTI feedback controller.

The dashed region shown in Fig. 2.1 represents the receding-horizon adaptive controller. It is apparent that the controller requires an estimate of the plant, $\hat{G}(z)$. Since only the error signal, $e(t)$, is measured in applications, estimates of the state, $\hat{x}(t)$, and

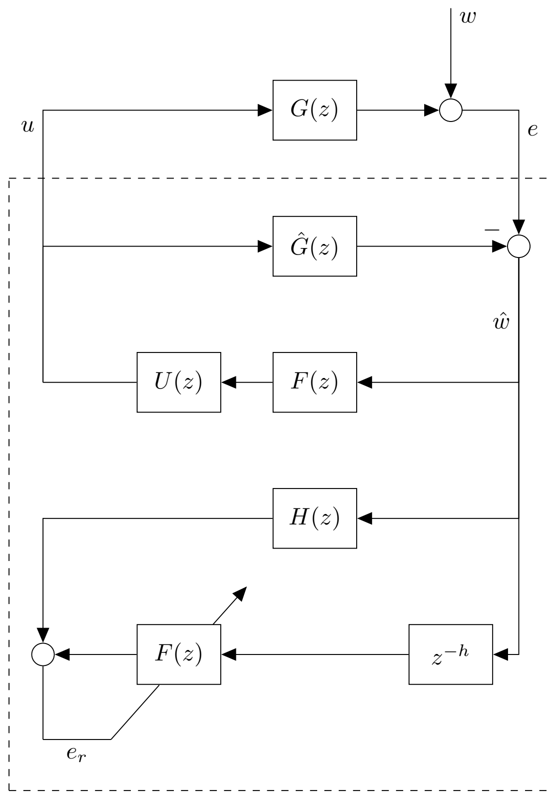

Figure 2.1: Block Diagram for Receding-Horizon Adaptive Control

disturbance, $\hat{w}(t)$, are used. Additionally, the plant itself is a linear model obtained by performing a one-time, offline system identification. Thus, the estimate of the plant available to the adaptive controller, $\hat{G}(z)$, has the $n-$dimensional state-space representation:

$$\hat{x}(t+1) = \hat{A}\hat{x}(t) + \hat{B}u(t), \tag{2.3}$$

$$e(t) = \hat{C}\hat{x}(t) + \hat{w}(t) \tag{2.4}$$

with $\hat{x}(t) \in \mathbb{R}^{n\times 1}$, $u(t) \in \mathbb{R}^{l\times 1}$, and $e(t), \hat{w}(t) \in \mathbb{R}^{m\times 1}$

As opposed to the more standard, infinite-horizon optimal controller, this research applies a model-predictive controller or receding-horizon controller, with fixed, finite horizon length, h. The quadratic performance index used in this research is:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \overbrace{\sum_{k=1}^{h} \Big[e^T(t+k)Q_1 e(t+k) + \hat{x}^T(t+k)Q_2\hat{x}(t+k) \Big]}^{①}$$
$$+ \underbrace{\sum_{k=0}^{h-1} \Big[u^T(t+k)R_u u(t+k) \Big]}_{②} + \underbrace{\hat{x}^T(t+h)Q_3\hat{x}(t+h)}_{③} \qquad (2.5)$$

with the Q_i matrices symmetric and positive semi-definite and the R matrix symmetric and positive definite and where the column vectors, $\hat{w}_h(t)$ and $u_h(t)$, are defined by:

$$\hat{w}_h(t) = [\hat{w}(t+1), \hat{w}(t+2), \ldots, \hat{w}(t+h)]^T \qquad \in \mathbb{R}^{mh\times 1} \qquad (2.6)$$
$$u_h(t) = [u(t), u(t+1), u(t+2), \ldots, u(t+h-1)]^T \qquad \in \mathbb{R}^{lh\times 1} \qquad (2.7)$$

Additionally, we define:

$$\hat{x}_h(t) = [\hat{x}(t), \hat{x}(t+1), \hat{x}(t+2), \ldots, \hat{x}(t+h)]^T \qquad \in \mathbb{R}^{n(h+1)\times 1} \qquad (2.8)$$
$$e_h(t) = [e(t+1), e(t+2), \ldots, e(t+h)]^T \qquad \in \mathbb{R}^{mh\times 1} \qquad (2.9)$$

Then, from equation (2.5), terms ① and ③ can be written as:

$$① + ③ = \hat{x}_h(t)^T\tilde{Q}_2\hat{x}_h(t) + e_h(t)^T\tilde{Q}_1 e_h(t) \qquad (2.10)$$
$$= \begin{bmatrix} \hat{x}_h(t) \\ e_h(t) \end{bmatrix}^T \begin{bmatrix} \tilde{Q} \end{bmatrix} \begin{bmatrix} \hat{x}_h(t) \\ e_h(t) \end{bmatrix} \qquad (2.11)$$

where, for $Q_1 \in \mathbb{R}^{m\times m}$ and $Q_2 \in \mathbb{R}^{n\times n}$:

$$\tilde{Q}_1 = \begin{bmatrix} Q_1 & & & \\ & Q_1 & & \\ & & \ddots & \\ & & & Q_1 \end{bmatrix} \in \mathbb{R}^{mh\times mh} \tag{2.12}$$

$$\tilde{Q}_2 = \begin{bmatrix} 0 & & & & \\ & Q_2 & & & \\ & & \ddots & & \\ & & & Q_2 & \\ & & & & (Q_2+Q_3) \end{bmatrix} \in \mathbb{R}^{n(h+1)\times n(h+1)} \tag{2.13}$$

and,

$$\tilde{Q} = \left[\begin{array}{c|c} \tilde{Q}_2 & 0 \\ \hline 0 & \tilde{Q}_1 \end{array}\right] \in \mathbb{R}^{[n(h+1)+mh]\times[n(h+1)+mh]} \tag{2.14}$$

From (2.3), (2.4), (2.6), and (2.7), $\hat{x}_h(t) \in \mathbb{R}^{n(h+1)\times 1}$ from (2.8) can be written as:

$$\hat{x}_h(t) = \tilde{A}_1 \begin{bmatrix} x(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B}_1 u_h(t) \tag{2.15}$$

where

$$\tilde{A}_1 = \left[\begin{array}{c|c} I & \\ \hat{A} & \\ \hat{A}^2 & \begin{bmatrix} 0 \end{bmatrix} \\ \vdots & \\ \hat{A}^h & \end{array}\right] \in \mathbb{R}^{n(h+1)\times(mh+n)} \tag{2.16}$$

and

$$\tilde{B}_1 = \left[\begin{array}{ccccc} 0 & \dots & & 0 & \\ \hline \hat{B} & & & & \\ \hat{A}\hat{B} & \hat{B} & & & \\ \hat{A}^2\hat{B} & \hat{A}\hat{B} & \hat{B} & & \\ \vdots & \vdots & & \ddots & \\ \hat{A}^{h-1}\hat{B} & \hat{A}^{h-2}\hat{B} & \dots & \hat{A}\hat{B} & \hat{B} \end{array}\right] \in \mathbb{R}^{n(h+1)\times lh} \tag{2.17}$$

Also, from (2.3), (2.4), (2.6), and (2.7), the term $e_h(t) \in \mathbb{R}^{mh\times 1}$ in (2.9) can be written as:

$$e_h(t) = \tilde{A}_2 \begin{bmatrix} x(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B}_2 u_h(t) \tag{2.18}$$

where

$$\tilde{A}_2 = \left[\begin{array}{c|c} \hat{C}\hat{A} & \\ \hat{C}\hat{A}^2 & \\ \vdots & \begin{bmatrix} I \end{bmatrix} \\ \hat{C}\hat{A}^h & \end{array}\right] \in \mathbb{R}^{mh\times(mh+n)} \tag{2.19}$$

and

$$\tilde{B}_2 = \begin{bmatrix} \hat{C}\hat{B} & & & & \\ \hat{C}\hat{A}\hat{B} & \hat{C}\hat{B} & & & \\ \hat{C}\hat{A}^2\hat{B} & \hat{C}\hat{A}\hat{B} & \hat{C}\hat{B} & & \\ \vdots & \vdots & \vdots & \ddots & \\ \hat{C}\hat{A}^{h-1}\hat{B} & \hat{C}\hat{A}^{h-2}\hat{B} & \dots & \hat{C}\hat{A}\hat{B} & \hat{C}\hat{B} \end{bmatrix} \in \mathbb{R}^{mh\times lh} \tag{2.20}$$

Combining (2.15) and (2.18), the term $\begin{bmatrix} \hat{x}_h(t) \\ e_h(t) \end{bmatrix}$ in (2.11) can be written as:

$$\begin{bmatrix} \hat{x}_h(t) \\ e_h(t) \end{bmatrix} = \tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B} u_h(t) \tag{2.21}$$

where

$$\tilde{A} = \begin{bmatrix} \tilde{A}_1 \\ \hline \tilde{A}_2 \end{bmatrix} \quad \text{and} \quad \tilde{B} = \begin{bmatrix} \tilde{B}_1 \\ \hline \tilde{B}_2 \end{bmatrix} \tag{2.22}$$

Also, for $R \in \mathbb{R}^{l \times l}$, term ② in equation (2.5) can be written as:

$$② = u_h(t)^T \tilde{R} u_h(t) \tag{2.23}$$

where

$$\tilde{R} = \begin{bmatrix} R_u & & & \\ & R_u & & \\ & & \ddots & \\ & & & R_u \end{bmatrix} \in \mathbb{R}^{lh \times lh} \tag{2.24}$$

Finally, from (2.11), (2.21), and (2.23), the performance index in (2.5) can be re-written as:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \left(\tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B} u_h(t) \right)^T \tilde{Q} \left(\tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} + \tilde{B} u_h(t) \right) + u_h(t)^T \tilde{R} u_h(t) \tag{2.25}$$

with $\tilde{A}$, $\tilde{B}$, $\tilde{Q}$, and $\tilde{R}$ defined above.

The minimizing $u_h(t)$ is a linear function of $\hat{x}(t)$ and $\hat{w}_h(t)$ and has length hl as defined in (2.7). The receding-horizon adaptive control command, $u(t)$, is chosen to be the first l terms in the minimizing $u_h(t)$ and is applied at each time step. The RH control command at the next time step is again the first l terms in the minimizing $u_h(t)$ computed again at that time step. In general, at each time step, h future control commands are produced, but only the first one is applied, then the whole process begins again at the next time step, applying only the first control command of that series of future control commands.

Thus, the receding-horizon control command can be written as:

$$u(t) = -K_x \hat{x}(t) - H_w \hat{w}_h(t)$$

where K_x and H_w can be computed in terms of the state-space matrices in (2.3) and (2.4) as well as $\tilde{A}$, $\tilde{B}$, $\tilde{Q}$, and $\tilde{R}$. To determine the K_x and H_w matrices, begin by expanding the quadratic performance index in (2.25):

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = u_h(t)^T \underbrace{[\tilde{B}^T \tilde{Q} \tilde{B} + \tilde{R}]}_{\bar{Q} \in \mathbb{R}^{lh \times lh}} u_h(t) + 2 \underbrace{\begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix}^T \tilde{A}^T \tilde{Q} \tilde{B}}_{b^T \in \mathbb{R}^{1 \times lh}} u_h(t)$$

$$+ \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix}^T \tilde{A}^T \tilde{Q} \tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix}$$

where,

$$b = \tilde{B}^T \tilde{Q}^T \tilde{A} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} \in \mathbb{R}^{lh \times 1} \tag{2.26}$$

This quadratic form (in $u_h(t)$) has the minimizing solution:

$$u_h(t) = -\bar{Q}^{-1} b = -\underbrace{[\tilde{B}^T \tilde{Q} \tilde{B} + \tilde{R}]^{-1} \tilde{B}^T \tilde{Q}^T \tilde{A}}_{\Gamma} \begin{bmatrix} \hat{x}(t) \\ \hat{w}_h(t) \end{bmatrix} \tag{2.27}$$

The matrix $\Gamma \in \mathbb{R}^{lh \times (mh+n)}$ can be partitioned such that:

$$u_h(t) = -\Gamma_n \hat{x}(t) - \Gamma_{mh} \hat{w}_h(t) \in \mathbb{R}^{lh \times 1} \tag{2.28}$$

where Γ_n is the first n columns of Γ and Γ_{mh} contains the remaining mh columns of Γ. Recall that $u_h(t)$ represents the h future control commands to be applied. However, the receding-horizon controller only implements the control command at time t, and repeats the entire process for each subsequent time step. To compute just the current control command, $u(t)$, the top l terms of $u_h(t)$ are used, so that:

$$u(t) = -K_x\hat{x}(t) - H_w\hat{w}_h(t) \in \mathbb{R}^{l\times 1} \tag{2.29}$$

where $K_x \in \mathbb{R}^{l\times n}$ is the top l rows of Γ_n, and $H_w \in \mathbb{R}^{l\times mh}$ is the top l rows of Γ_{mh}. It is important to note that the term $H_w\hat{w}_h(t)$ is a linear combination of h future values of the disturbance $\hat{w}(t)$.

The block diagram in figure 2.1 generates the adaptive control command as $u(t) = U(z)F(z)\hat{w}(t)$. To see how this is equivalent to the control command in equation (2.29), first note that the term $H_w\hat{w}_h(t)$ can be written as

$$\begin{aligned}
H_w\hat{w}_h(t) &= \begin{bmatrix} H_w(1) & H_w(2) & \dots & H_w(h) \end{bmatrix} \begin{bmatrix} \hat{w}(t+1) \\ \hat{w}(t+2) \\ \vdots \\ \hat{w}(t+h) \end{bmatrix} \qquad (2.30) \\
&= \sum_{k=1}^{h} H_w(k)\hat{w}(t+k) \\
&= \sum_{k=1}^{h} H_w(k)z^k\hat{w}(t) \\
z^{-h}H_w\hat{w}_h(t) &= \underbrace{\sum_{k=1}^{h} H_w(k)z^{(k-h)}}_{H(z)}\hat{w}(t)
\end{aligned}$$

$$H_w\hat{w}_h(t) = z^h H(z)\hat{w}(t) \tag{2.31}$$

Note that $H(z)$ is a causal FIR filter since $k - h \leq 0$.

Consider then the signal $e_r' = z^h H(z)\hat{w}(t) + F(z)\hat{w}(t)$ and suppose the bottom copy of the adaptive filter $F(z)$ minimizes the variance of e_r'. This implies that the output of the top copy of $F(z)$, i.e. $F(z)\hat{w}$, would be a minimum variance prediction of $-z^h H(z)\hat{w}$, or $-H_w\hat{w}_h$, as in (2.31). However, because the signal e_r' is not available at time t (non causal), it cannot be processed in real time. To work around this problem, the block diagram constructs a delayed version of e_r' as $e_r = H(z)\hat{w}(t) + z^{-h}F(z)\hat{w}(t)$. Noting that, in steady-state, minimizing the variance of e_r' is equivalent to minimizing the variance of e_r, the adaptive filter $F(z)$ can minimize the variance of e_r because the signal e_r is causal and available at time t. The result is that the top copy of $F(z)$ outputs a minimum variance prediction of $-H_w\hat{w}_h(t)$ in real time.

It is then clear that the input to the $U(z)$ block is $-H_w\hat{w}_h$. Additionally, the desired output of the $U(z)$ block is $u(t)$ in equation (2.29). This output is achieved by plugging in the $u(t)$ from equation (2.29) into the plant model in (2.3). Writing the state-space equations such that the input is $-H_w\hat{w}_h$, and the output is $u(t)$, we have

$$\begin{aligned}\hat{x}(t+1) &= (\hat{A} - \hat{B}K_x)\hat{x}(t) + \hat{B}(-H_w\hat{w}_h)\\ u(t) &= -K_x\hat{x}(t) + (-H_w\hat{w}_h)\end{aligned}$$

Thus, the $U(z)$ block has the state-space realization

$$U(z) = (\hat{A} - \hat{B}K_x, \hat{B}, -K_x, I). \tag{2.32}$$

Due to the direct feedthrough term in (2.32), the output of the $U(z)$ block is

$$u(t) = -K_x\hat{x}(t) - H_w\hat{w}_h(t)$$

which agrees with (2.29).

Additionally, it is clear that the weighting matrices in (2.5) should be chosen so that the resulting K_x matrix yields a stable $\hat{A} - \hat{B}K_x$ matrix. One way to guarantee that $\hat{A} - \hat{B}K_x$ is stable is to chose Q_3 in (2.5) such that

$$\hat{C}^T Q_1 \hat{C} + Q_2 + Q_3 = P \tag{2.33}$$

where the symmetric P is a stabilizing solution to the algebraic Riccati equation

$$P = \hat{C}^T Q_1 \hat{C} + Q_2 + \hat{A}^T [P - P\hat{B}(R + \hat{B}^t P \hat{B})^{-1} \hat{B}^T P] \hat{A} \tag{2.34}$$

For this case,

$$K_x = (R + \hat{B}^T P \hat{B})^{-1} \hat{B}^T P \hat{A} \tag{2.35}$$

However, for the applications in this research, K_x and H_w are directly computed as the solution to a finite-dimensional linear-quadratic optimization problem with the objective function in (2.25) as shown in (2.29).

To clarify, the adaptive part of the control algorithm minimizes the variance of e_r, defined above, and is ultimately used to predict the linear combination of h future disturbances as in (2.29). It does so using a multi-channel recursive least-squares (RLS) lattice filter [11] with finite-impulse response. In principle, a transversal FIR filter could be paired with any multi-channel RLS algorithm to update the filter gains, however, the lattice filter and RLS updating algorithm used in the is research has demonstrated superior numerical stability and efficiency due to orthogonalization of the data channels [11].

2.2 Frequency Weighting

One of the primary benefits of the receding-horizon performance index in (2.5) is the large design space parameterized by the weighting matrices Q_i and R_u. The receding-horizon control design also makes it straightforward to expand the state vector and plant model in

(2.3) and (2.4) to incorporate frequency weighting. This state-space weighting filter can be applied to the control signal, $u(t)$, or the output error, $e(t)$, or both. In past applications of this type of frequency weighting [9], a weighting filter with a cutoff frequency above that of the wavefront error bandwidth was used to penalize the control effort from amplifying high frequency sensor noise. In this research, the frequency weighting filter is still used to limit the control $u(t)$ at high frequencies, but the motivation is slightly different. In previous adaptive optics experiments [9], while the number of control channels is large, the plant dynamics were relatively simple, modeled by just a delay. In the research presented here, the plant dynamics for both experiments are significantly more complex. Because the adaptive controller requires a model of the plant in each case, an accurate system identification is required. However, due to inherent system nonlinearities and limitations on identification methods, plant modeling errors tend to be significant at high frequencies. Thus, the goal for the research here is to use the frequency weighting to limit the control from acting at higher frequencies where the plant model is poor, but still allowing the control to reject lower frequency disturbances as demonstrated in Chapters 3 and 4.

The method of incorporating frequency weighting depends on expanding the state vector, $\hat{x}$, in (2.3) and (2.4) as:

$$\hat{x} = \begin{bmatrix} \hat{x}_p \\ x_f \end{bmatrix} \tag{2.36}$$

where $\hat{x}_p$ is the estimate of the plant state vector, and x_f is the state vector of the frequency weighting filter. The frequency weighting filter has the state-space realization $(A_f, B_f, C_f, 0)$, so that the state-space matrices in (2.3) and (2.4), and the Q_2 matrix in (2.5) have the form:

$$\hat{A} = \begin{bmatrix} \hat{A}_p & 0 \\ 0 & A_f \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} \hat{B}_p \\ B_f \end{bmatrix}, \quad \hat{C} = \begin{bmatrix} \hat{C}_p & 0 \end{bmatrix}, \qquad Q_2 = \begin{bmatrix} 0 & 0 \\ 0 & C_f^T R_f C_f \end{bmatrix} \tag{2.37}$$

Next, note that the receding-horizon performance index in (2.5) contains the term:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \sum_{k=1}^{h} \left[\cdots + \hat{x}^T(t+k) Q_2 \hat{x}(t+k) \right] + \ldots$$

Using the expanded state vector and matrices above, the term $\hat{x}^T Q_2 \hat{x}$ can be written as

$$\begin{aligned} \hat{x}^T Q_2 \hat{x} &= \begin{bmatrix} \hat{x}_p^T & x_f^T \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 0 & C_f^T R_f C_f \end{bmatrix} \begin{bmatrix} \hat{x}_p \\ x_f \end{bmatrix} \\ &= x_f^T C_f^T R_f C_f x_f \\ &= R_f y_f^T y_f \\ &= R_f |y_f|^2 \end{aligned} \tag{2.38}$$

where y_f is the output of the weighting filter, with $u(t)$ as the input as described in (2.37).

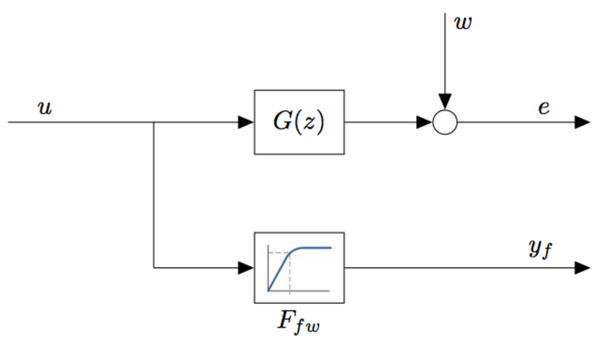

Figure 2.2: Frequency Weighting Filter Schematic

Thus, if the weighting filter is chosen to be a *high-pass* filter as in Fig. 2.2, the term $|y_f|^2$ will be large due to any high frequency content in $u(t)$. As a result, the $u(t)$ that minimizes the receding-horizon performance index will contain *less* high frequency content. By then applying this penalized $u(t)$ to the system, the high frequency plant modeling errors and/or high frequency sensor noise are not amplified, which is desirable because experimental results have demonstrated that sufficient amplification of high frequency modeling errors can cause system instability.

Note also, that the performance index in (2.5) contains the term:

$$J_h(\hat{x}(t), \hat{w}_h(t), u_h(t)) = \cdots + \sum_{k=0}^{h-1} \left[u^T(t+k) R_u u(t+k) \right] + \ldots$$

The R_u weighting parameter here is used in a similar manner to penalize the control command, $u(t)$, but does so without the use of a weighting filter. By scaling R_u to be larger, the value of the performance index also increases. The minimizing $u(t)$ then attempts to reduce (minimize) the value of the performance index by simply reducing the amplitude of the control command (i.e. the control command is penalized equally across all frequencies). The result is the same as with the frequency weighting in that the control command is penalized at higher frequencies, and consequently does not amplify high frequency plant modeling errors. It is different, however, in that it also penalizes the control command at lower frequencies, where legitimate disturbances reside. The effects of penalizing the control command using R_u and R_f are thoroughly discussed in Chapter 3.

Finally, it is important to note that in the derivation above, the measurement is regarded as an output *error* denoted $e(t)$. However, in subsequent chapters, the measurement is considered to be the system output and is appropriately denoted $y(t)$; it is interchangeable with the $e(t)$ signal used in the derivation.

CHAPTER 3

Laser Beam Steering Experiment

3.1 Background

The ability to propagate laser beams over long distances and focus high-intensity light onto precise targets is becoming increasingly important in applications such as laser communications and directed energy systems. Regardless of the application, two primary types of disturbances can severely degrade the performance of the system. Atmospheric turbulence can produce inconsistencies in air density which can refract the light in different directions. This type of disturbance is generally broadband in nature and affects ground-based systems that reside within the earth's atmosphere. The second type of disturbance comes from mechanical vibrations which can excite the natural frequencies of the platforms on which the laser source or target is mounted. These disturbances are generally narrowband and affect not only ground-based systems, but free-space systems, such as satellite-to-satellite communication, wherein atmospheric disturbances are of lesser concern. These sources of disturbance *(jitter)* cause erratic fluctuations *(beam error)* in the spot location at the target which corresponds to significantly less data transfer (laser communications), lower levels of focused energy (directed energy systems), or similar undesired effects in other applications.

The goal of the adaptive control scheme for the beam steering experiment is to maximize the light intensity at the target by rejecting the broad- and narrow-band jitter discussed above, which may have changing statistics. The two primary adaptive filtering methods that have recently been developed are the least-mean-squares (LMS) and the recursive least-squares (RLS) adaptive filters. While the LMS algorithms are stable and simple to implement, RLS algorithms have been shown to converge to their optimal gains more quickly, despite being numerically more complex. The research addressed in this dissertation deals with a lattice realization of the RLS algorithm, as in [5, 6, 12], at the heart of the adaptive controller.

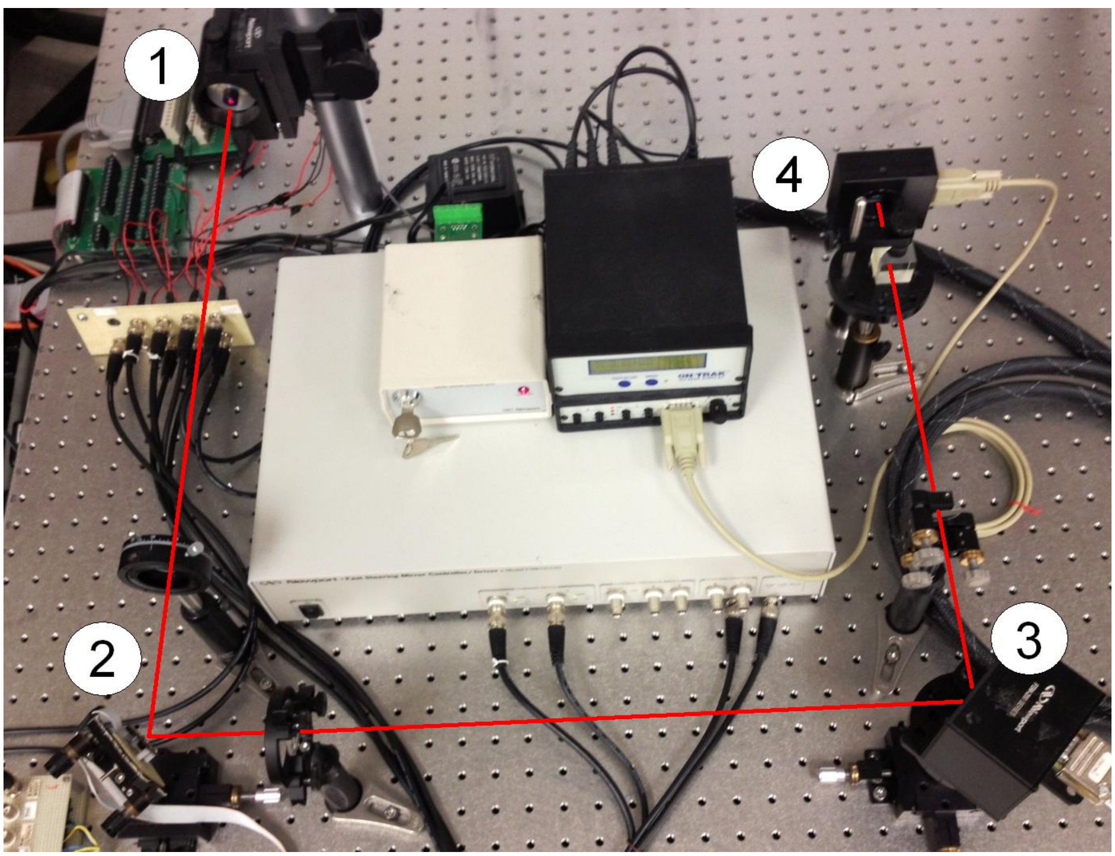

Figure 3.1: Laser Beam Steering Experiment

The main components of the beam control experiment shown in Figure 3.1 are: a laser source at position ①, a MEMS fast steering control mirror (FSM-C) at position ②, a fast steering disturbance mirror (FSM-D) at position ③, and an optical sensor at position 4. FSM-C is a Texas Instruments Micro Mirror and FSM-D is a 1" Newport fast steering mirror. Both mirrors are capable of motion about horizontal and vertical axes. The optical position sensor computes the position of the centroid of the incoming light and outputs two voltages corresponding to the horizontal and vertical position of the centroid. These signals are used in both the inner classical LTI control loop as well as in the outer adaptive control loop.

A dedicated PC running xPC Target executes the control algorithm, receiving position data and sending appropriate control signals to the actuators in real time. The adaptive controller only has access to the position measurements, despite the fact that the disturbance commands to FSM-D are also provided by this machine. In earlier experiments [10], the sample rate was 2kHz however, the most recent experimental results are based on a faster sample rate of 5kHz.

3.2 System Identification

3.2.1 Open-Loop Plant

The transfer function, $P(z)$, represents the open-loop mapping from a 2x1 input command at FSM-C to a 2x1 output measurement at the optical sensor, both in volts. The true open-loop plant, $P(z)$, is a 2x2 coupled-channel system, with control inputs each affecting both output measurements.

Careful calibration in the experimental setup yielded off-diagonal (coupling) terms that were extremely small relative to the diagonal terms. For this reason, the experiment is treated as two uncoupled single-input-single-output (SISO) systems throughout this research. Fig. 3.2 (top) shows the Bode plots for the two open-loop SISO transfer functions representing the vertical and horizontal axes, identified using Matlab's *n4sid* subspace identification function with 80,000 data points, or 16 seconds of data sampled at 5kHz.

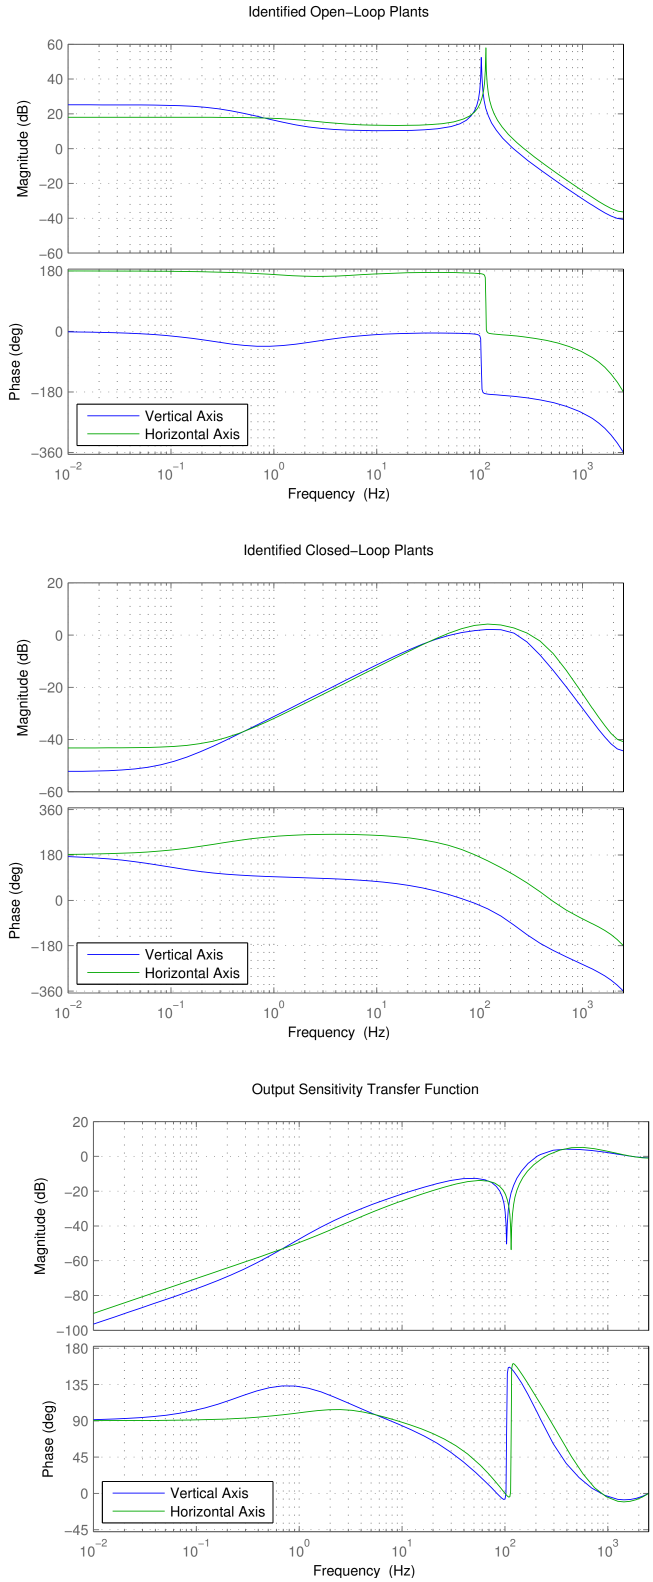

Figure 3.2: Bode Diagram of Identified Open-Loop Plants, $\hat{P}(z)$ (top), Bode Diagram of Identified Closed-Loop Plants, $\hat{G}(z)$ (middle), Bode Diagram of Output Sensitivity Transfer Function, $\hat{S}(z)$ (bottom)

3.2.2 Closed-Loop Plant

A classical LTI PID controller, denoted $C(z)$, is used in the feedback path to close the loop around $P(z)$ such that the closed-loop transfer function is: $G(z) = \frac{P(z)}{1+C(z)P(z)}$. Figure 3.2 (middle) shows the Bode plots for the closed-loop plants, obtained from performing a separate, one-time identification using input-output data. The output sensitivity transfer function (Fig 3.2 (bottom)), $\hat{S}(z) = \frac{1}{1+C(z)\hat{P}(z)}$, determines how the output error will respond to a disturbance signal, w, as in Fig. 2.1 when only the LTI loop is closed. The goal in designing $C(z)$ was to maximize the disturbance rejection bandwidth while minimizing the amount of disturbance amplification. Figure 3.2 (bottom) shows that the PID controller designed for this experiment yields a 6dB disturbance-rejection bandwidth of approximately 160Hz and amplifies output disturbance and noise by less than 5dB between 300Hz and 700Hz. Attempting to increase the error-rejection bandwidth any further resulted in unacceptably high levels of noise amplification at higher frequencies.

3.3 Laser Beam Steering System

3.3.1 System Nonlinearities

The configuration of the optical path used in this experimental setup is such that the angular range of the FSM-C is larger than in previous experiments [13]. Put another way, in order to achieve the same output displacement at the optical sensor in this research, the FSM-C must be driven to a larger angle than in previous experiments. Because of this, the piezo-electric actuators that drive the FSM-C mirror became saturated at times, and the plant exhibited fairly severe nonlinear behavior. This actuator saturation was was a function of both amplitude and frequency, as seen in Fig. 3.3.

Figure 3.3 shows the results of a test performed to identify and characterize the nonlinearities. In each of the three plots shown, two sinusoids (small and large amplitude)

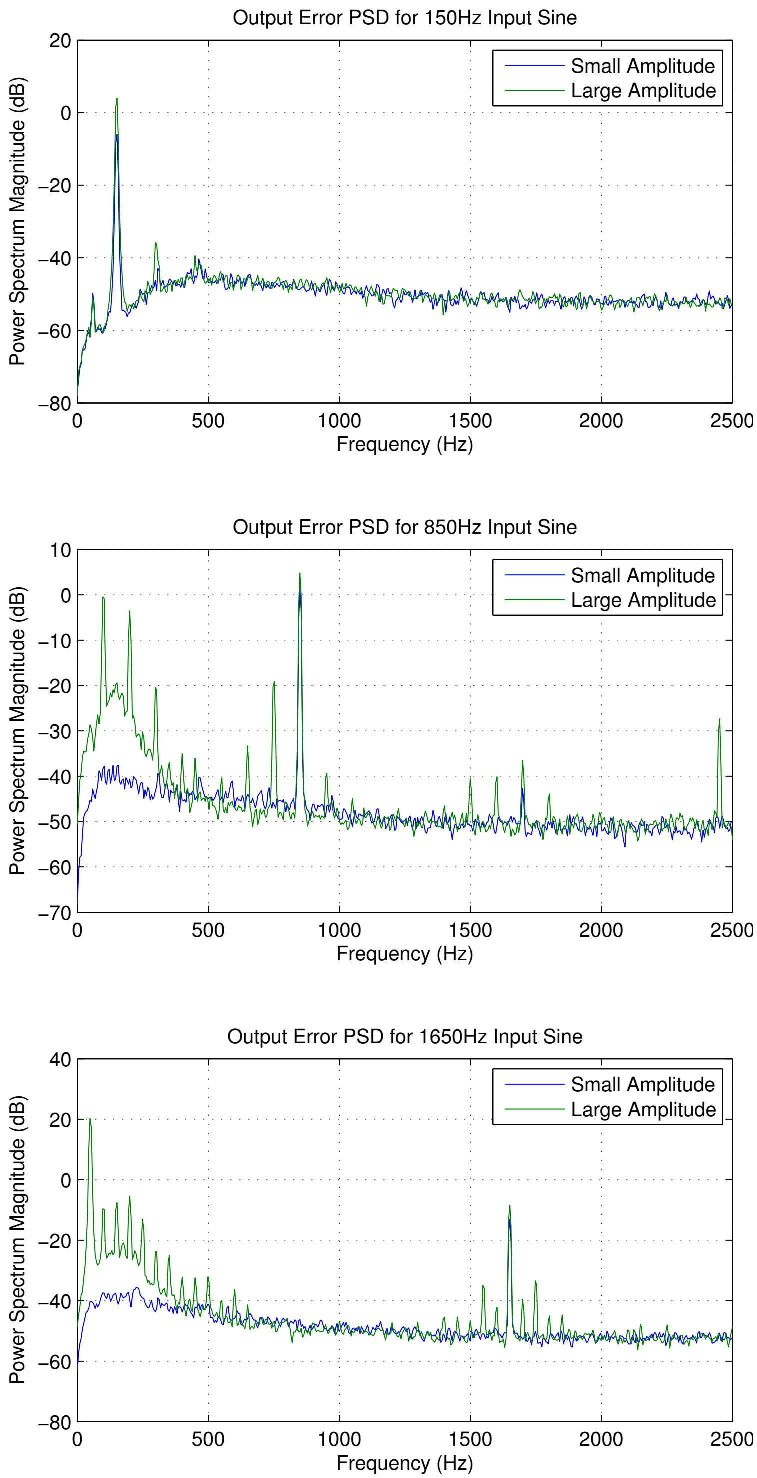

Figure 3.3: Output PSDs of sinusoidal signals passed though $G(z)$ at three different frequencies and at two different amplitudes each.

were passed through the closed-loop plant, $G(z)$, and the PSDs of the output signals were plotted. "Small" and "large" amplitudes are in reference to the actual operating point, i.e. they are smaller and larger compared to the adaptive control command used in the actual experiment. In a linear system, the output PSD should reveal a single, dominant spike at the frequency of the input sinusoid. This is the case at 150Hz for both a small and large amplitude (Fig. 3.3, bottom). However, for larger amplitude sine waves at 850Hz and 1650Hz, the output PSDs show numerous harmonics, indicating the presence of nonlinear behavior. The conclusion that the cause of the nonlinearities is actuator saturation was drawn from a simulation where a saturation block was placed in front of the plant. The output PSDs of the simulated signals exhibited the same harmonic spiking phenomenon as in the above test.

3.3.2 Experimental Results

The experimental results in this section compare the effects of penalizing the adaptive control command across all frequencies equally (R_u), and penalizing according to the weighting filter by using the R_f parameter. Additionally, the effect of the adaptive filter order is studied with regards to the minimum order required to effectively reduce the disturbance as well as the robustness of the control system when confronted with a severely over-parametrized adaptive controller (higher filter order than necessary).

Because the nonlinearities are most prevalent at frequencies above 700Hz, the frequency weighting filter used in experiment is a high-pass filter with cutoff frequency of 1000Hz. The rolloff of this fourth-order filter yielded a sufficient amount of control command penalty at 700Hz (where the nonlinearities began to appear), with a greater effect at higher frequencies (where the nonlinearities were more pronounced). Figure 3.4 shows the bode plot of the weighting filter used throughout the experiment.

Figure 3.5 shows the gains in the FIR filter $H(z)$ for two different combinations of R_u

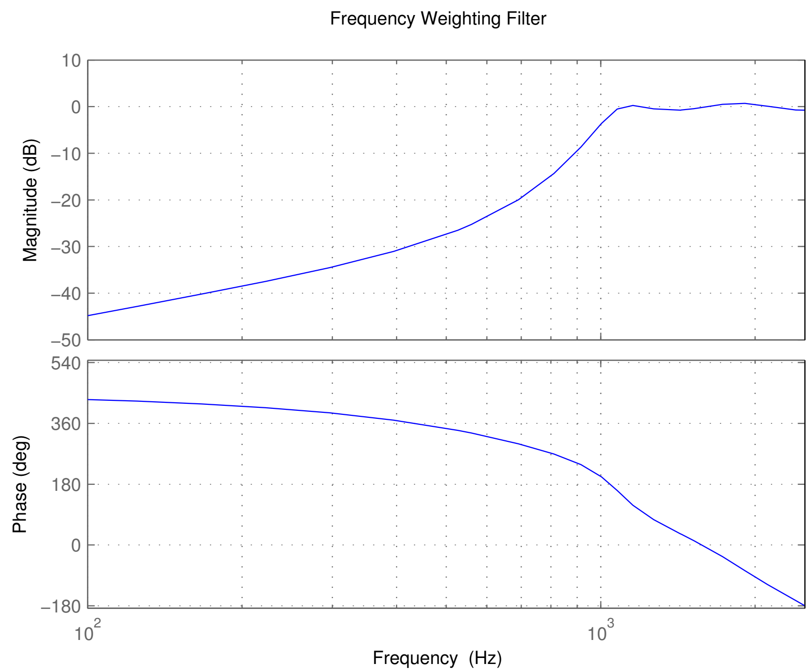

Figure 3.4: Bode Plot of High-Pass Frequency Weighting Filter.

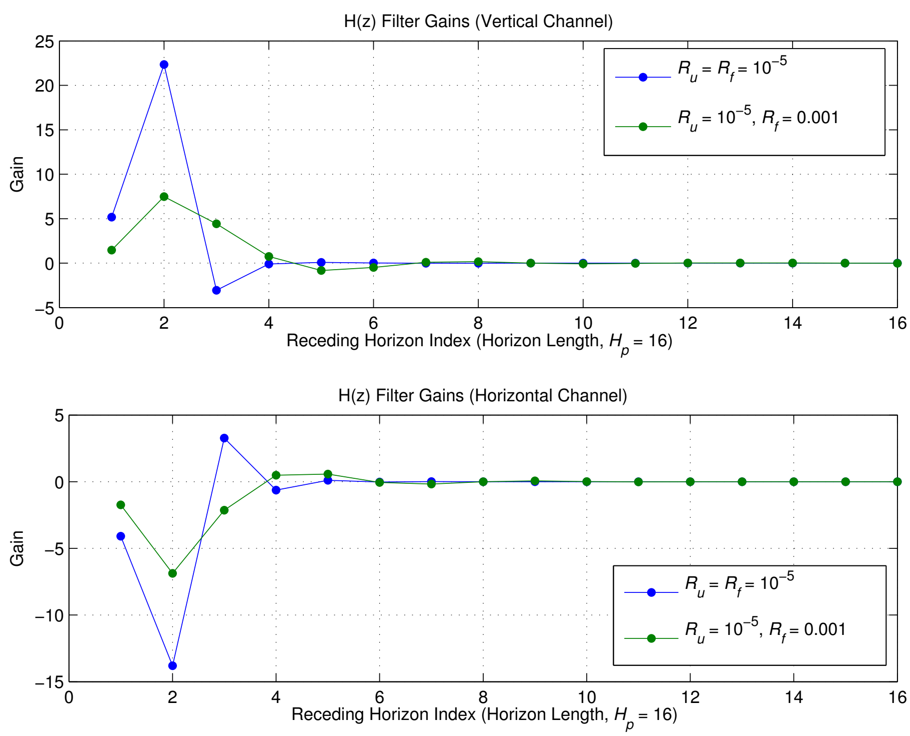

Figure 3.5: Gains in FIR filter $H(z)$. Number of taps = horizon length $h_p = 16$. Top: vertical channel. Bottom: horizontal channel.

and R_f as a function of the horizon length, h_p. In all cases, the gains become insignificant above $h \approx 8$. Because experimental tests confirmed that the results are nearly identical for $h_p = 8$,16, and 32, a horizon length $h_p = 16$ is used throughout this experiment.

The disturbance applied to FSM-D was generated by passing broadband noise through various filters in parallel, and is used throughout the experiment. The frequency components of the disturbance signal are summarized in Table 3.1.

Jitter Bandwidths
30 Hz - 80 Hz
450 Hz - 500 Hz
Peak at 875 Hz
Peak at 1100 Hz
0 Hz - 500 Hz

Table 3.1: Jitter Bandwidths used in FSM-D to Generate Disturbance

Figure 3.6 shows the results of three independent experimental runs that compare the performance of the LTI controller alone to the LTI loop augmented by two versions of the adaptive controller. The first 6 seconds of each experiment are identical; the system runs open-loop for $0 \le t < 3s$, then the LTI loop is closed for $3s \le t < 6s$. For $6s \le t \le 10s$, the legend indicates the control configuration. In one experiment, the adaptive loop is never closed. This serves as a baseline case against which the adaptive loop performance can be compared. In one of the experiments with the adaptive loop closed, the control penalties are relatively small ($R_u = R_f = 10^{-5}$) and the order of the adaptive filter, $F(z)$, is 2. In the other experiment with adaptive control, there is a higher penalty on the adaptive control command ($R_u = 10^{-5}$, $R_f = 10^{-3}$), as well as a higher adaptive filter order of 32.

With $R_u = R_f = 10^{-5}$, the nonlinearities in the plant caused frequent instability at adaptive filter orders greater than 2. However, when R_f is increased to 10^{-3}, the frequency weighting filter limited the adaptive control command at higher frequencies allowing the adaptive filter order to be increased up to 90 without going unstable. Despite a higher

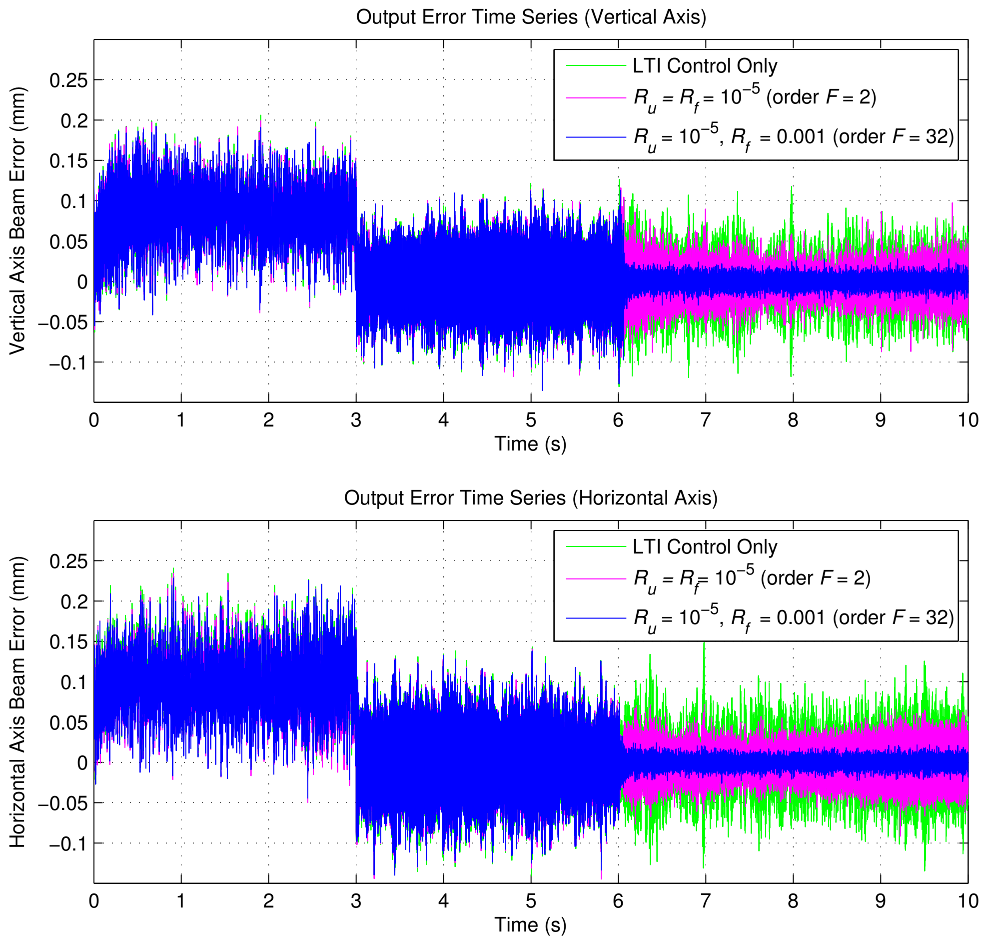

Figure 3.6: Output error time series. Open-loop ($0 \leq t < 3s$), classical LTI control only ($3s \leq t < 6s$), classical LTI control + adaptive control ($6s \leq t \leq 10s$). Top: vertical position error. Bottom: horizontal position error.

penalty, the output error for the $R_u = 10^{-5}$, $R_f = 10^{-3}$ case with adaptive filter order 32 is much smaller than the case with adaptive filter order 2.

Figure 3.6 also illustrates the performance of the LTI controller. From $3s \leq t < 6s$, the LTI controller is closed. It is apparent that the integrator is working to eliminate the bias, however, during this time range the variance is actually higher than that of the open-loop only interval ($0s \leq t < 3s$). For the the vertical channel, the LTI controller increases the output error variance by 27%, and for the horizontal channel, the output error variance is increased by 36% when only the LTI controller is applied.

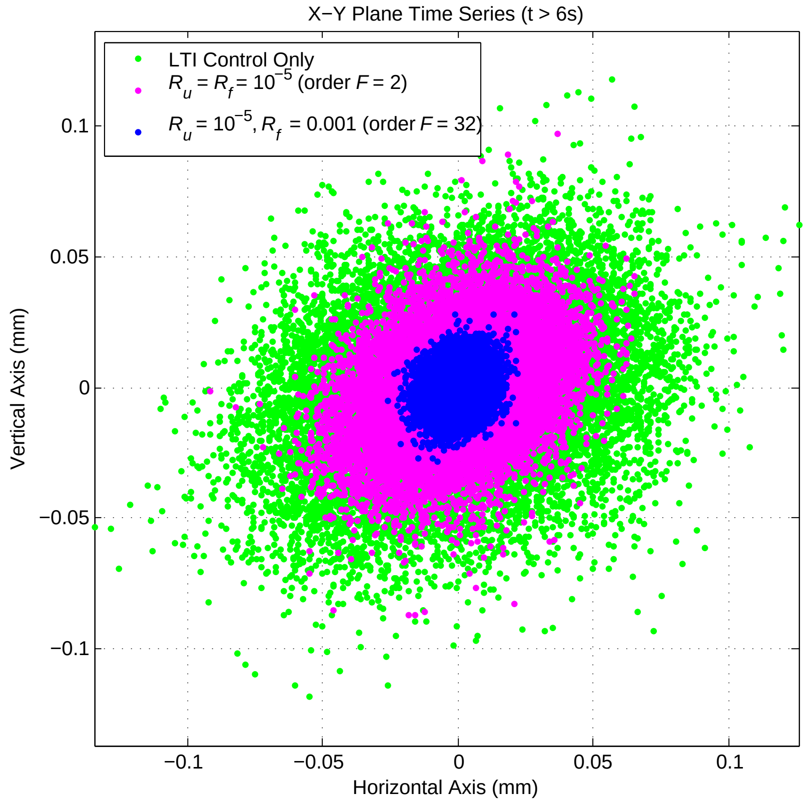

Figure 3.7: Output error scatter plot. Represents spatial time series for $6s \leq t \leq 10s$ for all three experiments in Fig. 3.6 .

Figure 3.7 depicts a spatial time series for the three cases (one LTI-only, and two adaptive), for $6s \leq t \leq 10s$. It shows the effect of the adaptive controller in reducing the variance of the output and highlights the ability of the adaptive controller to hold the laser spot tightly on a target, in the presence of a complex disturbance.

The power spectral densities (PSDs) in Fig. 3.8 reflect the output spectra for three cases: (i) Open-Loop, (ii) LTI control only, and (iii) LTI control augmented by adaptive control with order of $F = 32$. Each PSD is based on 3 seconds of data (15,000 samples). While cases (ii) and (iii) are taken directly from the last three seconds of the time series in Fig. 3.6, the open-loop PSD was generated by running a separate experiment for 10 seconds with no control and using the last 3 seconds of that time series.

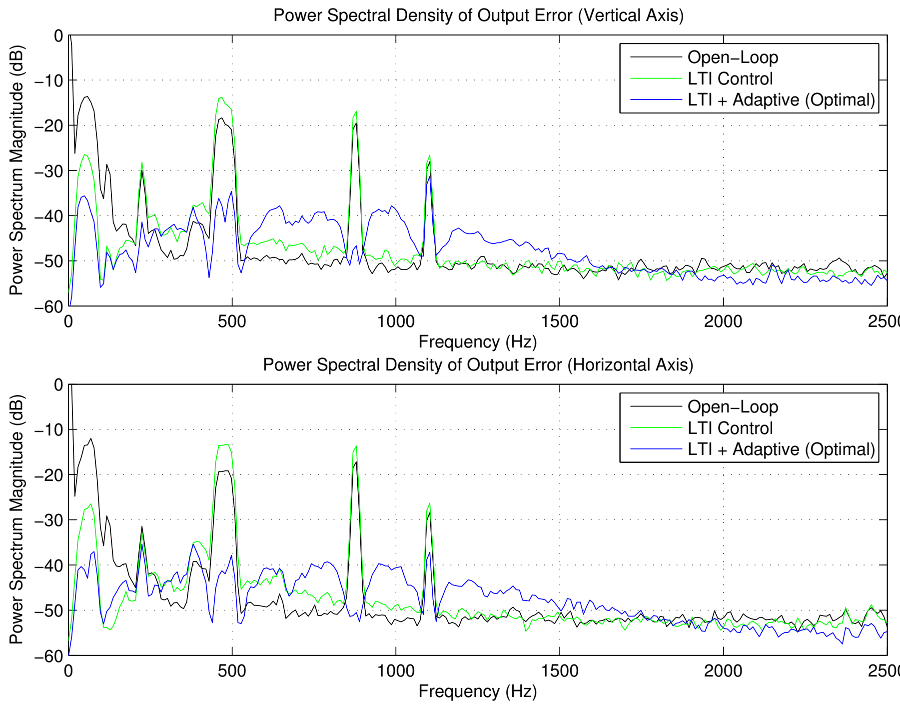

Figure 3.8: Power spectral densities of steady-state output errors for three cases: open-loop, LTI control only, and LTI control + adaptive control with $R_u = 10^{-5}, Rf = 10^{-3}$, order of F = 32. Top: vertical error. Bottom: horizontal error.

The PSDs in Fig. 3.8 agree quite well with the output sensitivity transfer functions in Fig. 3.2 (bottom). Comparing the open-loop case to the LTI-only case, it is apparent that the LTI controller does a decent job attenuating the disturbances below about 200Hz, namely the 30Hz-80Hz band in the disturbance but also the lower portion of the disturbance band from 0Hz-500Hz. Additionally, as the sensitivity transfer function predicts, there is some level of amplification of the disturbance from 300Hz-1100Hz. This is clear when comparing the LTI-only PSD to the open-loop PSD at the 450Hz-500Hz band and the peaks at 875Hz and 1100Hz and noting that the amplitude is larger than the open-loop case.

The PSD in Fig. 3.8 associated with when the adaptive loop is closed reveals just how effective the adaptive controller is in reducing the output error by targeting the disturbance components. The frequency components in the output due to the disturbances at 30Hz-80Hz, 450Hz-500Hz, and 875Hz are significantly reduced with the adaptive controller closed. The peak due to the disturbance at 1100Hz is only mildly affected by the adaptive controller, however, this is expected because this is a case where the adaptive control command is penalized primarily according to the frequency weighting filter, which has a cutoff frequency of 1000Hz. Put another way, by penalizing the control command using a higher value of R_f, the adaptive control has less authority at frequencies above 1000Hz, and consequently does not target the 1000Hz disturbance peak as aggressively as the lower frequencies. This is a visual confirmation that the frequency weighting filter is effective in penalizing the adaptive control at higher frequencies.

3.3.3 Effect of Adaptive Filter Order

In general, as the complexity of a disturbance signal increases, so does the adaptive filter order required to effectively predict that disturbance. Figure 3.9 shows the RMS values of the output error signal for a range of adaptive filter orders ($order_F$). For this study, the control penalties are kept constant ($R_u = 10^{-5}$ and $R_f = 10^{-3}$) to isolate the effect of an increasing value of $order_F$. There is a data point (green) that represents the RMS value of the output error when only the LTI loop is closed. When closing the adaptive loop and using an adaptive filter order of 2, there is a significant amount of reduction in the output error. However, as expected, the output error can be further reduced by increasing $order_F$. It is apparent, for this particular disturbance profile, that the best performance (smallest output error RMS) is achieved when $order_F$ is about 30. Increasing the filter order beyond this value yields no significant improvement in performance.

Figure 3.9 also shows that, despite a negligible improvement in performance, the adaptive filter order can be increased beyond what is necessary to predict the disturbance (over-parametrization) without causing detrimental effects on the performance. In early experiments, a filter order too large would often cause an increase in output error, or

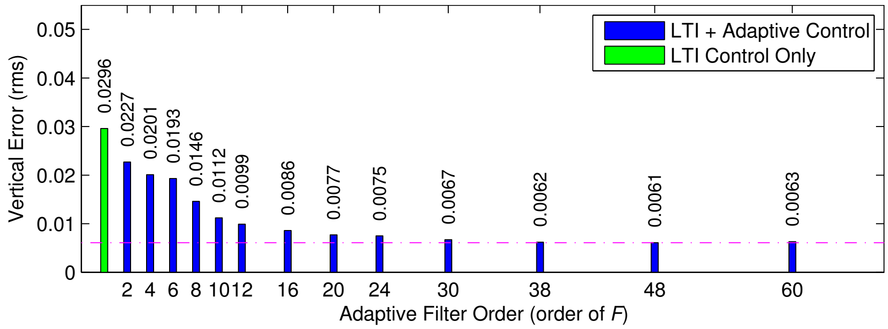

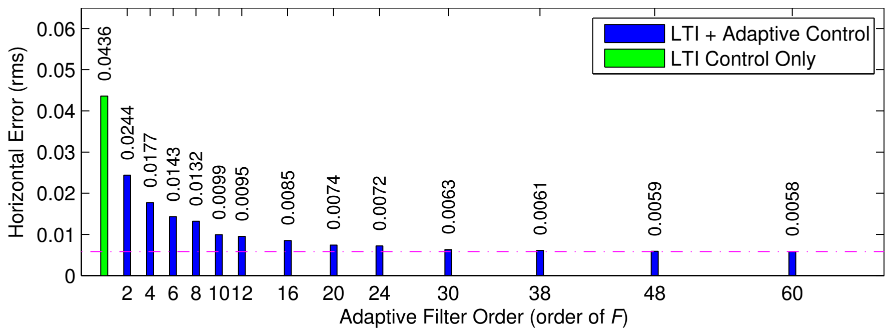

Figure 3.9: RMS values of steady-state output errors vs. adaptive filter order. Top: vertical error. Bottom: horizontal error. $R_u = 10^{-5}, R_f = 10^{-3}$.

in some cases, instability. Here, R_f is sufficiently large and allows for a reliably-stable controller, even when the adaptive filter order is much larger than necessary. Table 3.2 provides a detailed study on the effect of the control penalties, R_u and R_f, on the maximum allowable adaptive filter order for constantly stable operation. It shows that a sufficiently large control penalty must be used in order to be able to use a filter order high enough to effectively predict the disturbance. In practice, it is valuable to be able to use an over-parametrized adaptive controller ($order_F$ higher than necessary), because it is typically difficult to know the optimal $order_f$ beforehand.

R_u	R_f	Maximum order of F
1×10^{-5}	1×10^{-5}	2
2×10^{-4}	1×10^{-5}	2
4×10^{-4}	1×10^{-5}	26
6×10^{-4}	1×10^{-5}	58
8×10^{-4}	1×10^{-5}	80
1×10^{-3}	1×10^{-5}	100+
1×10^{-2}	1×10^{-5}	100+
1×10^{-5}	2×10^{-4}	2
1×10^{-5}	4×10^{-4}	2
1×10^{-5}	6×10^{-4}	8
1×10^{-5}	8×10^{-4}	64
1×10^{-5}	1×10^{-3}	90
1×10^{-5}	1×10^{-2}	100+

Table 3.2: Effects of R_u and R_f on the maximum adaptive filter order for consistent stability.

3.3.4 Effect of R_u vs. R_f

As previously discussed, R_u penalizes the adaptive controller equally across all frequencies whereas R_f penalizes the adaptive control command according to the frequency weighting filter, allowing for less control effort to be applied at higher frequencies where plant modeling errors are more abundant.

Figure 3.10 illustrates the effect of each control penalty (R_u, R_f), on the RMS of the output error in each channel. For all cases $order_F = 32$, but to isolate the effect of each penalty term, either R_u or R_f was held constant and the other was allowed to vary. In both plots, the RMS value of the output error is smaller when R_u is kept small and a larger value of R_f is used. This behavior is expected because when penalizing the adaptive control command with R_f, it only suppresses the control command at higher frequencies, but allows the control to have higher authority at the lower frequencies where the plant model is better and where most of the disturbance resides. By contrast, when keeping R_f small and using the R_u parameter to penalize the control signal, the output

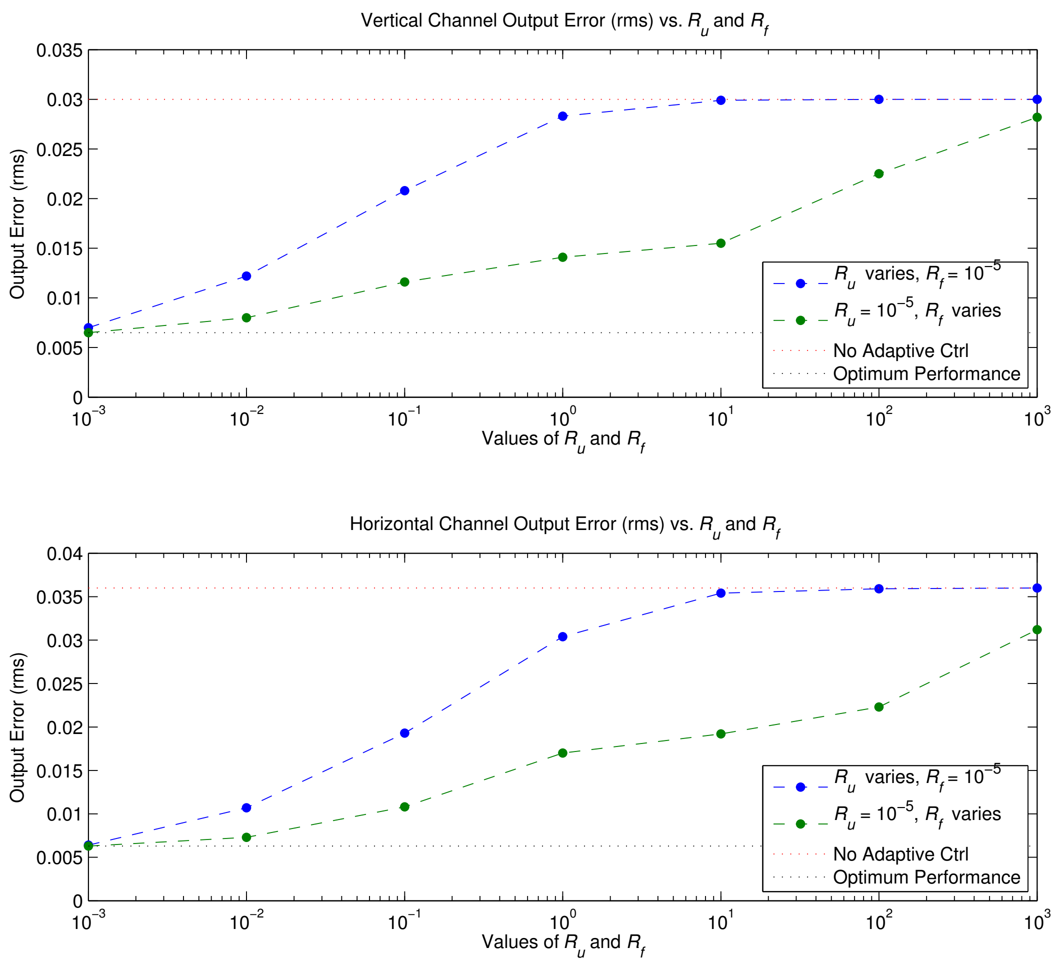

Figure 3.10: RMS output error as a function of R_u and R_f . Order of F = 32 for all cases. Top: vertical error. Bottom: horizontal error.

error RMS tends to be larger. This is because, in addition to penalizing the higher frequency components of the adaptive control command, the lower frequency range is equally penalized, effectively restricting the control command from fully reducing the amount of disturbances in the output.

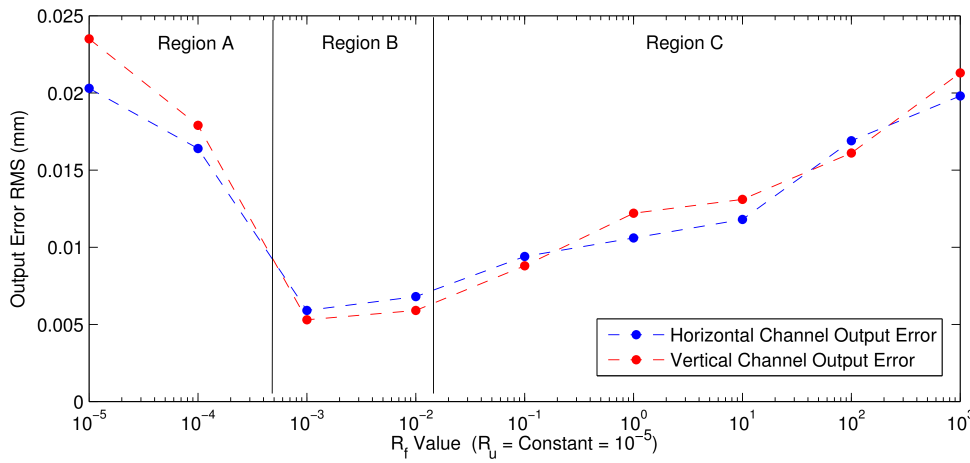

Figure 3.11: Visualization of Performance vs. Adaptive Control Command Penalty

Figure 3.11 depicts a visualization of the performance as a function of the severity of the adaptive control command penalty, R_f. In region A, the penalty is extremely small which, in theory, should yield nearly optimal performance. Experimental results, however, yield a controller that is not robust to modeling errors and thus, cannot operate near the optimal adaptive filter order (≈30) without consistent instability. The highest adaptive filter order used in Region A is order 3. In region C, the controller is robust to modeling errors because a sufficient penalty is applied to the adaptive control command. As the amount of control penalty increases, the controller becomes more robust, allowing for adaptive filter orders exceeding order 90. However, as a larger penalty is used, the adaptive control command becomes too severely penalized and its ability to influence the control mirror and significantly reject jitter becomes diminished. In region B, a balance is achieved where there is enough of a control penalty to yield a robust controller capable of operating at sufficiently large filter orders (> 30), but not so much as to over-penalize the control command and significantly restrict the controller's influence.

While the plot in Fig. 3.11 depicts a varying R_f value while R_u is held constant, the trend is similar for when R_u varies and R_f remains constant. Additionally, in Fig. 3.11, the adaptive filter order varies and was chosen to be near its maximum value for reliably stable operation for the given R_u and R_f values. This is in contrast to Fig. 3.10, which is intended to highlight the difference between using R_u or R_f by holding the adaptive filter order at a constant value.

The experimental results show that, when compared to a linear-time-invariant feedback loop, the adaptive controller significantly reduces the jitter in the position of the laser spot on the target sensor. The receding-horizon performance index provides a large design space for improving the robustness of the adaptive controller respect to plant modeling error due to plant nonlinearities and unmodeled sensor noise. The experimental results in this research demonstrate how control weighting in the performance index, especially frequency weighting on the adaptive control signal, significantly improves the robustness of the adaptive controller to unmodeled plant nonlinearities that are most prominent at high frequencies.

While this experimental research has demonstrated a successful implementation of the receding-horizon adaptive controller on a SISO plant, the controller is capable of implementation in a MIMO setting as well. Preliminary tests on the beam steering experiment have shown that a MIMO adaptive control architecture works in implementation. In one test, the two output channels were rotated in software, effectively generating a set of coupled channels. A separate 2x2 system identification was performed and the resulting model was using by the adaptive controller in a MIMO experiment with promising results. Additionally, the adaptive controller is theoretically capable of adapting to non-stationary disturbances, i.e. time-varying disturbances. Thus, the magnetic bearing discussed in the next chapter was the natural choice for a new plant on which to implement the controller because it is MIMO by nature and it incorporates an inherent time-varying disturbance as it rotates from rest to a final angular velocity.

CHAPTER 4

Magnetic Bearing Experiment

4.1 Background

Active magnetic bearings (AMBs) have recently become a popular alternative to traditional ball, roller, and journal bearings in applications where mechanical wear and bearing maintenance are of concern. Perhaps the greatest benefit of an AMB over a traditional bearing is the lack of physical contact between moving surfaces. The requirements of expelling heat and lubricating contact points in traditional bearings are eliminated with an AMB. Additionally, the mechanical friction associated with the traditional bearing is entirely eliminated in an AMB, improving the efficiency of many systems. Some applications that currently benefit from AMBs include: flywheel energy storage systems, wind turbines, and jet engines.

While the benefits of the magnetic bearing are apparent, there are several challenges and technological requirements that accompany the use of an AMB. Noting that any magnetic levitation system is open-loop unstable, a stabilizing controller must first be implemented simply to hold the bearing stationary within the bearing housing. Many forms of this stabilizing controller have been used, including the built-in PD-controller in the MBC500 described below as well as a similar lead compensator [14], and an LQG controller [15].

Even after the bearing is stabilized, mass imbalances in the bearing shaft can cause severe disturbances in the output when the shaft is spinning at high frequencies. The control problem can then be formulated as a disturbance rejection problem, in which an additional outer-loop controller can be implemented for regulation. While the disturbance is relatively simple (single-frequency sinusoid), it's frequency is time-varying, and it modulates with the bearing's angular velocity. Consequently, it is generally a non-stationary disturbance. In order to reject this type of disturbance, several adaptive control schemes have been implemented. One such control scheme is that in [16], in which a controller based on the internal model principle was used to reject sinusoidal disturbances at predetermined frequencies. Variations using repetitive control have also been used due to the periodic nature of the disturbance [17]. In this dissertation, the same inner-loop/outer-loop control scheme is used, with the outer-loop controller being the receding-horizon adaptive controller discussed in Chapter 2.

The Magnetic Moments MBC500 magnetic bearing system pictured in Fig. 4.1 is a four degree-of-freedom (DOF) magnetic levitation experiment in which electromagnets at each end of a 303-stainless steel rotor provide force in both horizontal and vertical directions. The Turbo edition, used in this research, features a shortened shaft, designed to push the resonant modes of the system past the speed range of the rotor. Hall-effect sensors measure the horizontal and vertical displacements at both ends of the bearing shaft and are recorded as voltages. The 4x1 endpoint displacement measurement vector, $\begin{bmatrix} x_1 & y_1 & x_2 & y_2 \end{bmatrix}'$, is transformed into a vector containing rotation and translation elements, $\begin{bmatrix} x_r & x_t & y_r & y_t \end{bmatrix}'$.

To describe the general control scheme, the open-loop unstable plant is stabilized using an LQG controller, producing a stable closed-loop system referred to as $G(z)$. Then, the adaptive loop is wrapped around this $G(z)$. The disturbances come from a shake-table support system that introduces broadband disturbances into the output measurement, as well as from mass imbalances in the shaft as it is spun to speeds up to 240Hz by an attached turbine and compressed air controlled by a separate PID loop actuated by a

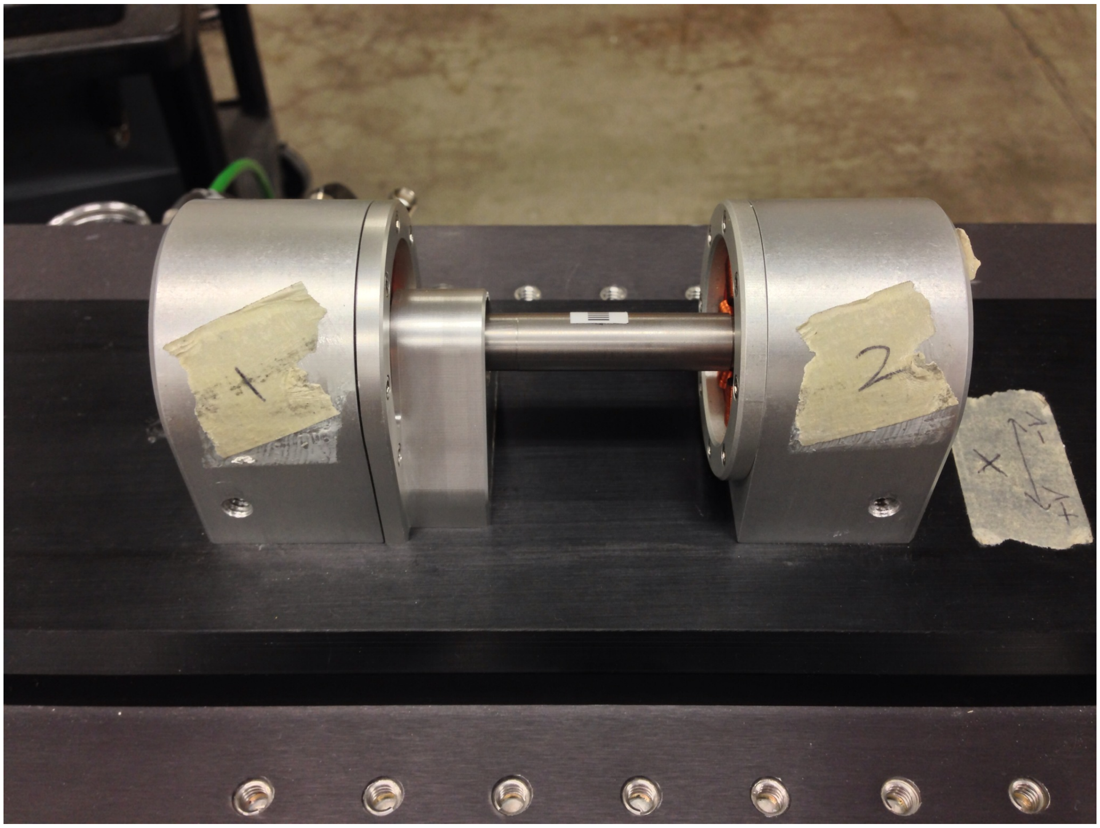

Figure 4.1: Magnetic Moments Magnetic Bearing

mechanical servo valve.

To spin the rotor, the rotor shaft is also fitted with turbine blades and a 2-way servo valve can be controlled to direct pressurized air. Regulating the speed of the rotor is an important aspect of the experiment to the extent that repeatability of the process is required. The pulses of the onboard encoder are used to trigger a high clock rate (20MHz) counter card to determine the speed of the rotor. The measured speed is used to close the loop on a Proportional-Integral-Derivative (PID) controller for servo valve command.

Stiction was observed to cause hunting of the speed set point, preventing the rotor from reaching any useful constant speed. To counteract stiction in the servo valve, a high frequency zero-mean dither signal was added to the valve command.

Data acquisition and control were implemented using the Mathworks xPC Target platform at a sampling frequency of 10kHZ. National Instruments PCI-6052e A/D and

PCI-6601 Counter cards were used for data acquisition and command. A shake table was placed under a corner of the system to introduce a broadband disturbance also controlled through the xPC environment. Figure 4.2 illustrates the primary hardware components used in this research.

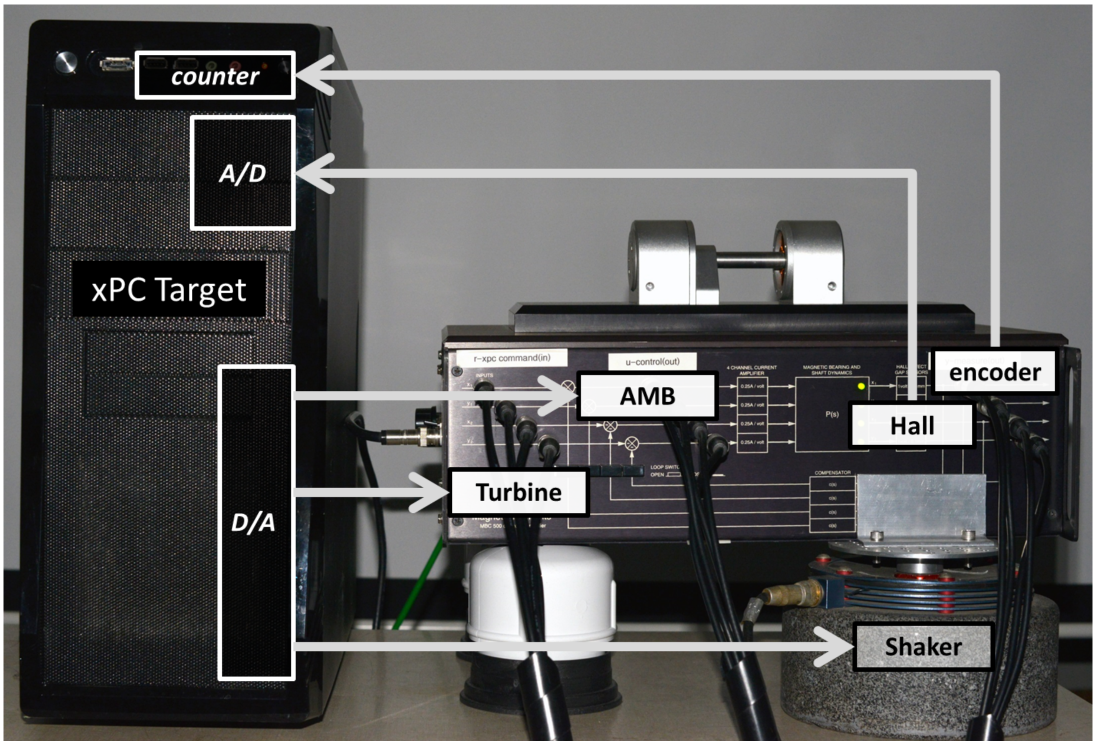

Figure 4.2: MBC500 Turbo magnetic bearing system hardware

4.2 System Identification

The adaptive controller requires a plant model, $\hat{G}(z)$, which was obtained from a system identification based on recorded input/output data. The identified plant $\hat{G}(z)$ represents a closed-loop stable system in which a 4-channel open-loop unstable plant $P(z)$ is stabilized by a robust LQG controller with an integrator $C(z)$, designed to stabilize the plant across a broad range of disturbance frequencies.

While the identification of $\hat{G}(z)$ was performed exclusively by the subspace methods utilized by Matlab's *n4sid* function, the choice of the plant model structure was driven by a test to determine the amount of coupling present in the system. Impulsive inputs were applied to each channel of the closed-loop system, $G(z)$, and the outputs were recorded. Figure 4.3 shows that when an impulsive input is applied to the x-translation channel, the output is seen predominantly in the same channel, but also a small amount is recorded in the x-rotation channel. This indicates that there is essentially no cross-plane coupling (into the y-channels), but there is some non-negligible in-plane coupling. This behavior is typical of applying impulses to the other channels. Additionally, while the tests depicted by Fig. 4.3 were conducted when the shaft was at rest, the results were similar when the shaft was spinning, indicating the absence of any significant gyroscopic coupling.

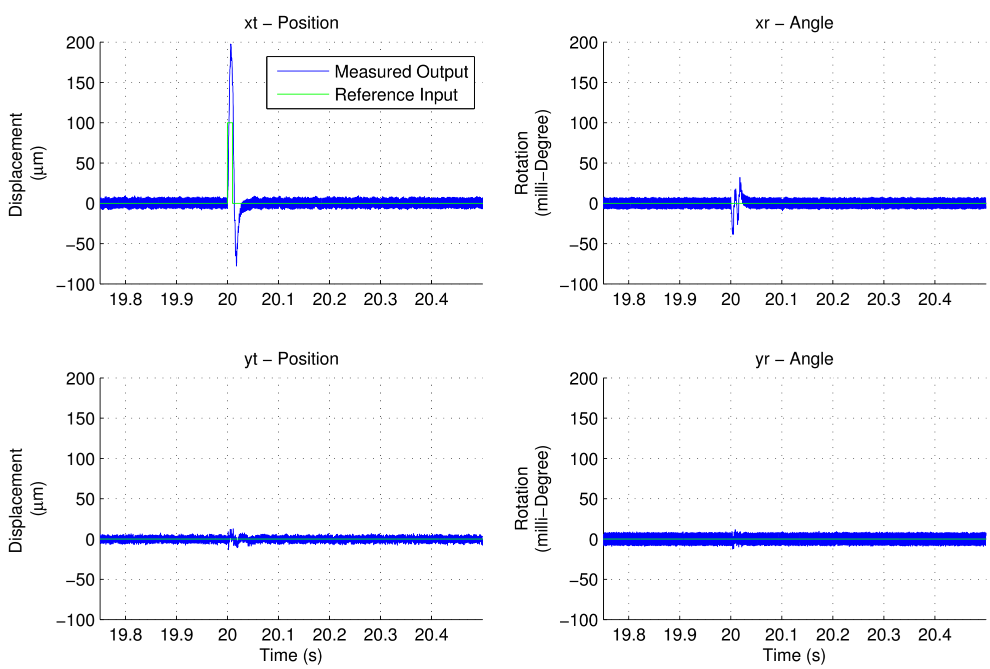

Figure 4.3: Impulsive input applied to x-translation channel

The likely reason for this in-plane coupling is that the transformation of coordinates embedded within $P(z)$ assumes that the center of mass of the shaft lies at half the shaft's length:

$$x_{trans} = \frac{yx_1 + yx_2}{2}, \quad x_{rot} = \frac{yx_1 - yx_2}{2} \tag{4.1}$$

$$y_{trans} = \frac{yx_1 - yx_2}{2}, \quad y_{rot} = \frac{yx_1 + yx_2}{2} \tag{4.2}$$

However, this is not truly the case because the turbine that is attached to the shaft adds a significant amount of mass to one of the shaft's ends, shifting the center of mass slightly to one end. While it may have been possible to define a different coordinate transformation to account for the true location of the center of mass, the simple transformation above is used throughout this work to test the adaptive controller's ability to handle coupled channels.

Because of the presence of some in-plane coupling and a lack of any significant cross-plane coupling, the control structure is treated as two independent adaptive control loops, one being for both horizontal channels, and the other being for the two vertical channels. Thus, the required plant models are two 2x2 MIMO models; one for the horizontal channels, $\hat{G}_{hor}(z)$ and another for the vertical channels, $\hat{G}_{vert}(z)$. Figure 4.4 shows the identified 2x2 plant models used.

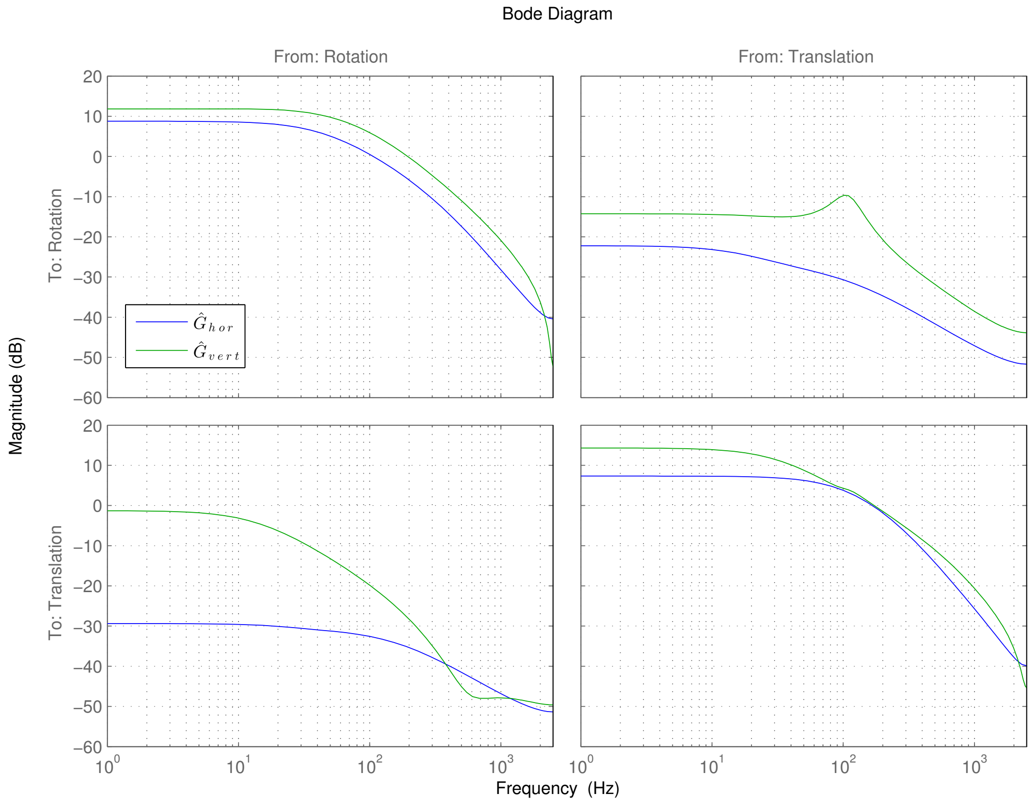

Figure 4.4: Bode Plots for Identified Closed-Loop Plant Models

4.3 Magnetic Bearing System

4.3.1 System Nonlinearities

Initial experimental tests showed that the adaptive controller was effective in reducing the disturbances caused by shaft imbalances while spinning. However, as shaft rotation rates were increased to above approximately 250Hz, a severe reduction in performance occurred in the x-plane channels. In the y-plane, a similar reduction in performance also occurred, but at a higher threshold frequency (approximately 500Hz) than for the x-plane. A series of tests ruled out the potential for a bad plant model as well as any malfunctioning hardware. Additionally, a 3-dimensional mapping of control effort required to hold the shaft at positions on a grid of x_r and x_t values revealed no problematic or nonlinear regions within the electromagnetic field.

The problem was ultimately identified to be a hardware limitation in the form of actuator saturation. Figure 4.5 shows results for a test performed at a low frequency (100Hz) and Fig. 4.6 shows results for the same test performed a high frequency (400Hz). A sinusoidal signal ($u = A_{\text{in}} \sin(\omega t)$) was fed in parallel through the identified closed-loop plant, $\hat{G}(z)$, and the true closed-loop plant, $G(z)$. The amplitudes of the output sinusoids for the simulated (blue) and experimental (magenta) cases were then plotted for a range of A_{in} values. Immediately apparent was the flattening-out of the experimental (magenta) output amplitudes for increasing values of A_{in} in the $\omega = 400$Hz case. This indicates actuator saturation, a phenomenon that occurs when the actuator simply cannot provide the necessary force to drive the output to the proper amplitude at a given frequency. The effect of this saturation is a nonlinear plant response, and consequently, poor adaptive controller performance at frequencies above 250Hz.

The blue and magenta plots in Fig. 4.5 and Fig. 4.6 reveal only that there exists actuator saturation at sufficiently large input amplitudes. In order to determine if the magnetic bearing experiment was actually operating in the nonlinear region, a second test was performed. First, in simulation, a sinusoidal disturbance was added without closing the adaptive loop. The amplitude of this disturbance sine wave (A_d) was adjusted until the output amplitude matched that of the output when the shaft was spinning at that frequency. This served to emulate the disturbance due to the shaft imbalance at a given frequency and is plotted as the green line in Fig. 4.5 and Fig. 4.6. Then, again in simulation, the adaptive loop was closed and the adaptive control command was recorded. The amplitude of the this signal was denoted $A_{\text{in, req}}$, the input amplitude of the control command required to sufficiently reduced the output error. This is plotted as the orange line on Fig. 4.5 and Fig. 4.6. In theory, the green and orange lines should intersect at the simulated (blue) plot, but because the control penalties cannot be exactly zero even in simulation, the vertical $A_{\text{in, req}}$ line is always slightly to the left of this intersection point, as expected. Once this test was performed, it became clear that at 400Hz (and really, at frequencies above $\approx$ 250Hz), in the x-Rotation channel, the actuators were saturated and the experiment was operating in a severely nonlinear region. The performance degradation was also seen in the x-translation channel, and is likely due to the coupling between the rotation and translation channels in the x-plane. It may also result from the fact that the operating point for the x-translation channel seems to be on the verge of the nonlinear region.

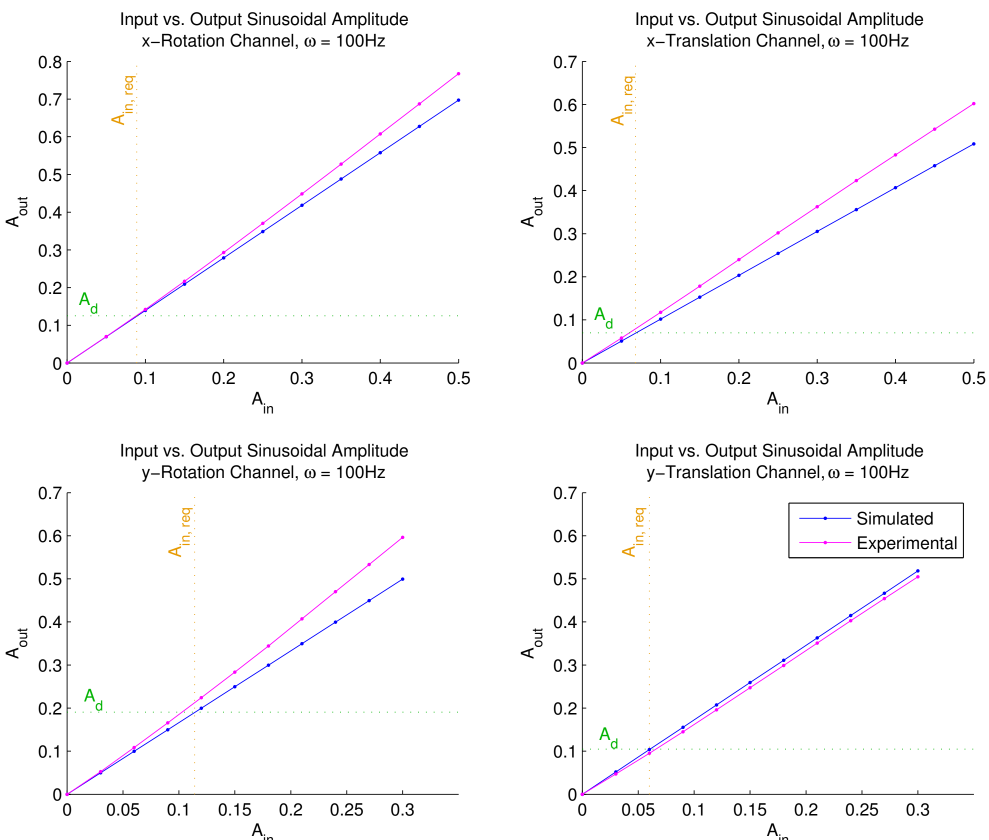

Figure 4.5: Plant Nonlinearity Test (100Hz)

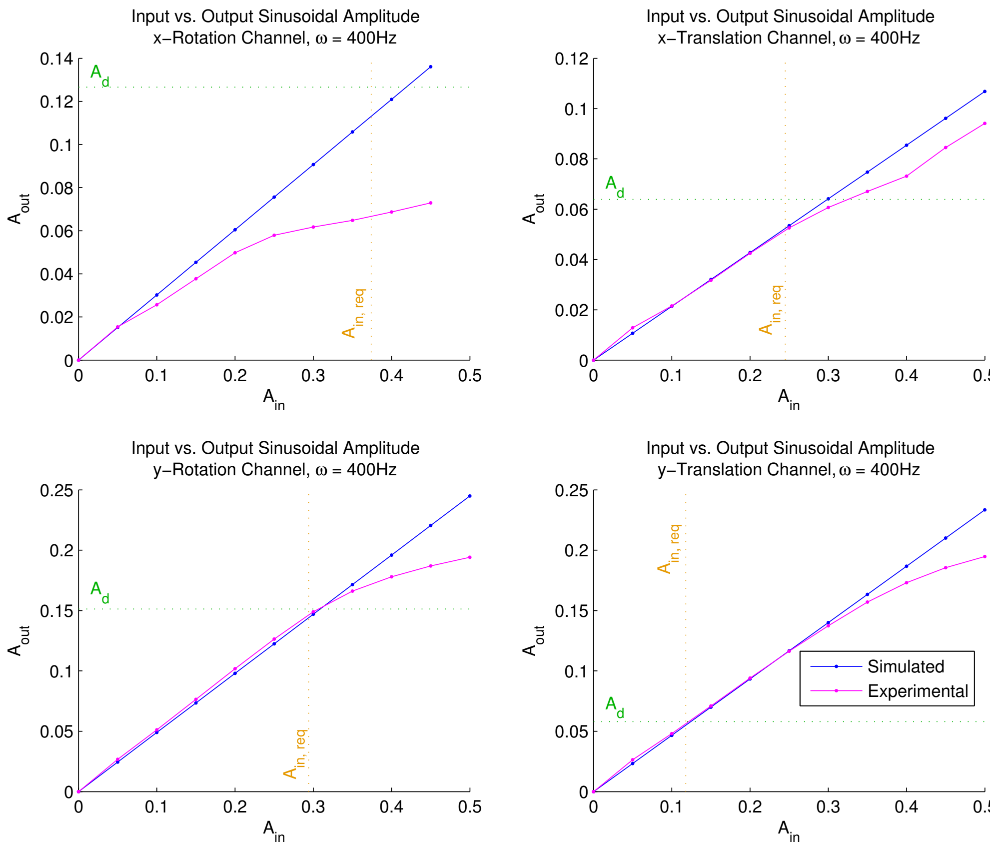

Figure 4.6: Plant Nonlinearity Test (400Hz)

4.3.2 Experimental Results

Because of the hardware limitation, the bearing's angular velocity was kept below a maximum spin rate of 250Hz. However, to make the control problem more challenging, two additional disturbance elements were added. The first element was to add a piezo-electric shake table which is capable of broader-band disturbances. By designing the shake table support system asymmetrically with respect to the shaft rotation axis, significant disturbances in the translational channel in each plane was achieved, as evident in Fig. 4.7 by noting the increased amplitude in the blue time series plots between $15s \leq t < 20s$.

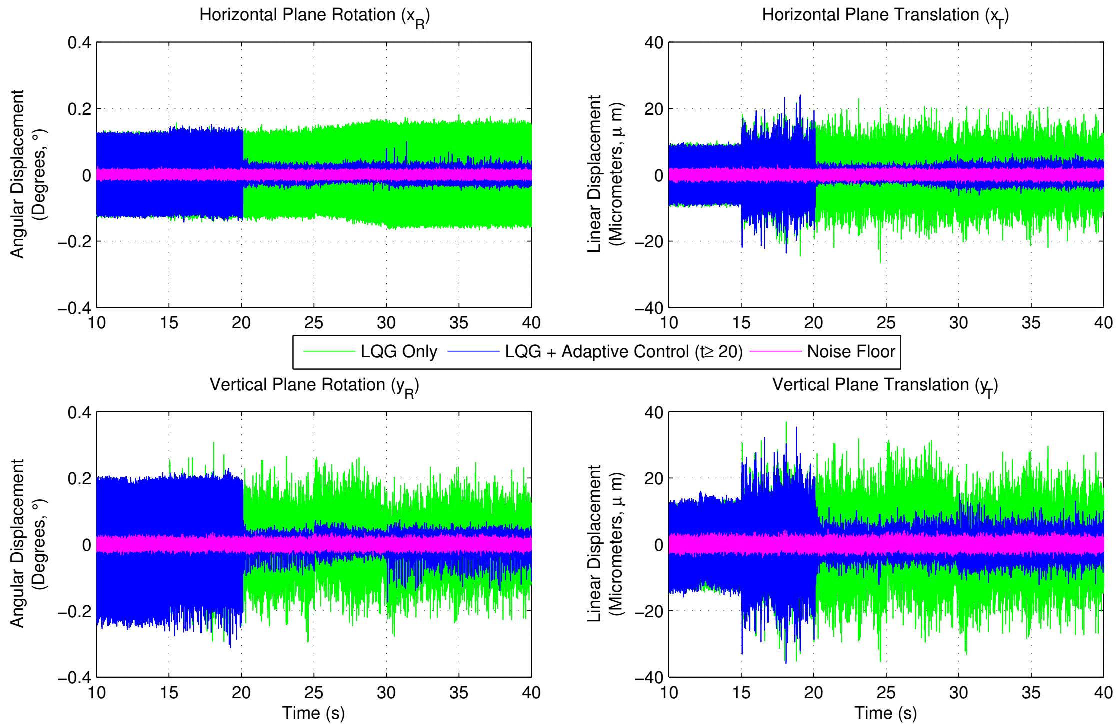

Figure 4.7: Time Series Plots for Broadband Disturbance + Time-Varying Disturbance

For the time series shown in Fig. 4.7, the blue plot shows the performance of the adaptive controller after $t = 20s$. The green plot in the background shows the system output in the absence of the adaptive control, and the magenta plot is the noise floor. In Fig. 4.7, the following sequence of events took place:

Time	Action
$0 \leq t < 5s$	Shaft reaches stable equilibrium
$t = 5s$	Shaft begins to spin up from rest
$10s \leq t < 15s$	Shaft spins at steady-state spin rate (120Hz)
$t = 15s$	Shaker disturbance is added (Remains on until end of experiment)
$t = 20s$	Adaptive control loop is closed (Remains closed until end of experiment)
$25s \leq t < 30s$	Shaft spin rate increased linearly from 120Hz to 240Hz
$30s \leq t < 40s$	Experiment reaches steady-state

Table 4.1: Sequence of events during experimental run

Typically, a plot illustrating the power spectral density (PSD) of the output with adaptive control vs. the PSD of the output with only the stabilizing loop would be included. However, because the output is non-stationary from $25s \leq t < 30s$, this comparison is omitted. Nonetheless, for $t \geq 20s$ it is clear in the blue time series that the adaptive controller is effective in reducing the amount of disturbance present in output signal on all four control channels.

4.3.3 Forgetting Factor, λ

The forgetting factor, $0 \leq \lambda < 1$, is a parameter built in to the RLS algorithm that allows data far in the past to be weighted exponentially less than more recent data when computing adaptive gains. The adaptive algorithm applies an attenuation multiplier defined as λ^n, to a data sample n steps back. As an example, because the sample rate for this experiment was 5000Hz, for $\lambda = 0.999$ a data sample taken 1 second in the past is weighted by a factor of $0.999^{5000} = 0.0067$. By contrast, for a small change in forgetting factor, if $\lambda = 0.9999$, a 1-second old data sample is weighted by 0.6065, which illustrates the sensitivity of the forgetting factor parameter.

Consequently, there exists an inherent trade-off when choosing a forgetting factor for a complex disturbance of this type. On one hand, a larger λ is desired when treating broadband disturbances (50-100Hz band produced by the shake table) because more data samples are required to capture all of the embedded frequencies. On the other hand, when a time-varying disturbance is present (i.e. shaft spin rate increasing), it is more important for the algorithm to weigh very recent data more heavily by using a smaller λ.

Figure 4.8 illustrates the problem with using a forgetting factor that is too large when dealing with only a time-varying disturbance:

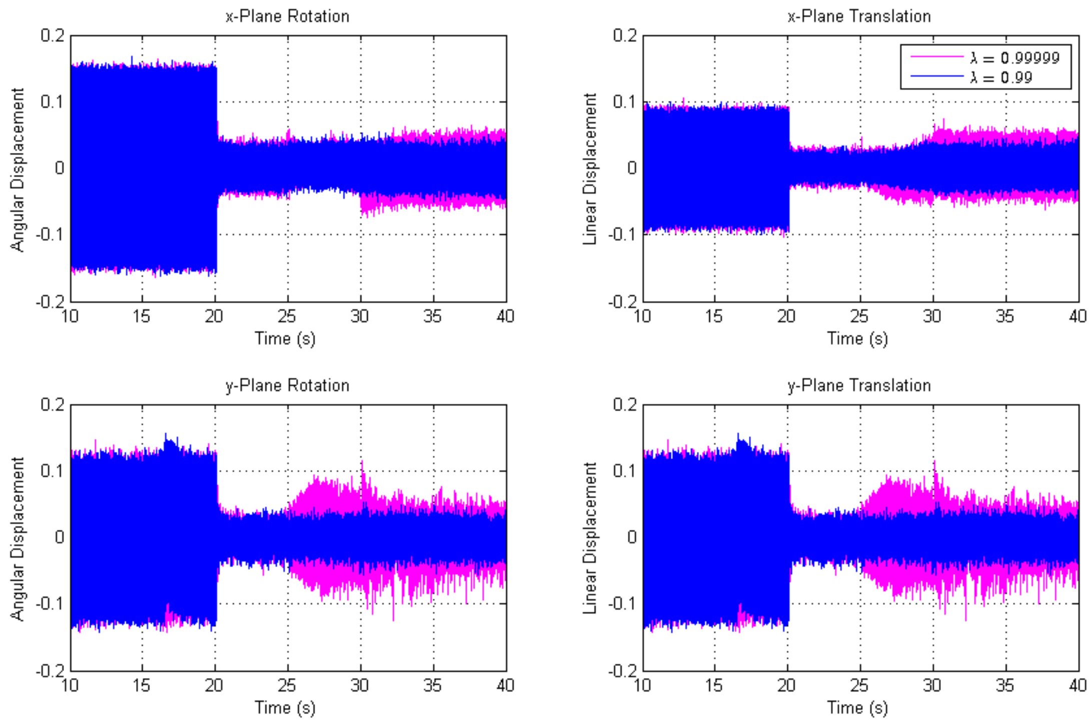

Figure 4.8: Forgetting Factor Comparison for Ramp Disturbance. Smaller values for λ yield improved performance by reducing output error.

For $\lambda = 0.99999$, the adaptive algorithm is heavily weighting data from several seconds back; even at 1 second in the past, the data is being weighted by a factor of $0.99999^{5000} = 0.9512$. However, during the spin from 120Hz to 240Hz between $25 < t < 30$, the adaptive controller should give very little weight to data far in the past and only consider very recent data. Thus, for $\lambda = 0.99$, which multiplies data even 0.1s in the past by a factor of only 0.0066, the performance is significantly improved.

The trade-off, of course, is that for broadband disturbances, a larger forgetting factor should be used to capture all of the frequencies contained. A similar plot is shown in Fig. 4.9 to show the opposite trend in performance as a function of λ when subject to only the shake table's disturbance. Here, the performance increases as λ gets larger because more past data is considered by the adaptive algorithm, yielding a more accurate prediction of the disturbance.

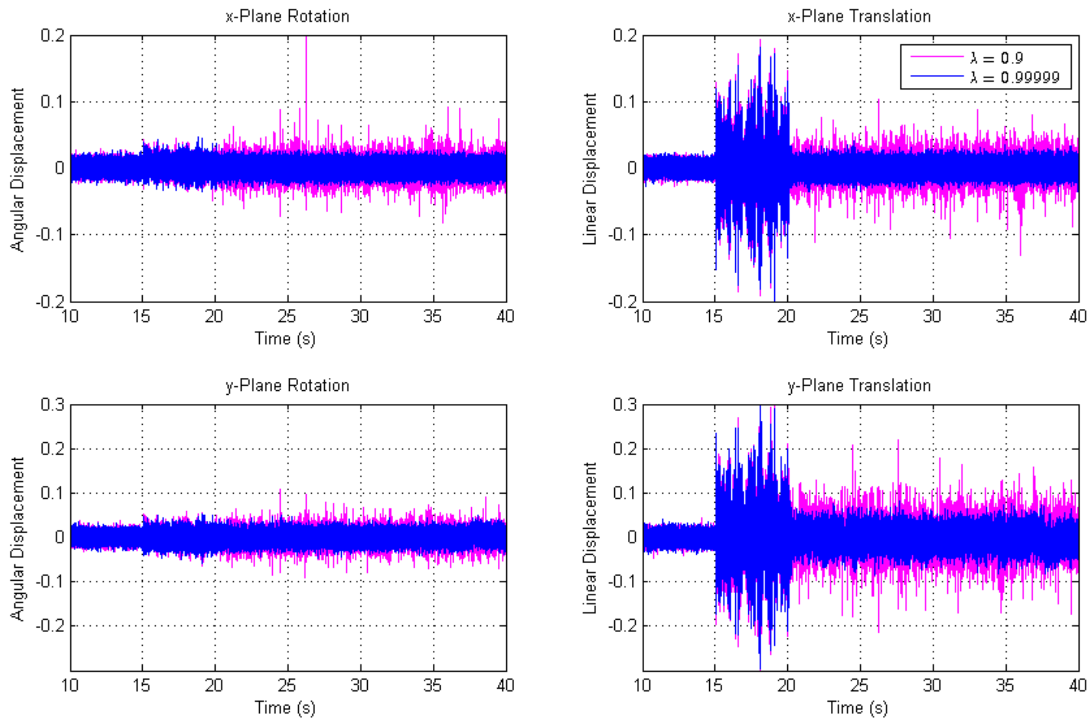

Figure 4.9: Forgetting Factor Comparison for Shaker Disturbance. Larger values for λ yield improved performance by reducing output error.

4.3.4 Coupled Channels in $\hat{G}(z)$

One of the primary goals of this paper is to highlight the controller's effectiveness in a MIMO configuration with plants containing severely coupled channels. While the Bode plots in Fig. 4.4 indicate that there is some amount of coupling between the channels in the 2×2 plants, the relative amount of coupling has not yet been discussed. One convenient way to visualize the amount of coupling between channels is to examine the $2 \times 2 \times h$ filter gains in $H(z)$. Figure 4.10 shows the $2 \times 2 \times h$ gains for three different controller configurations used for both the horizontal and vertical axes. The blue plots in Fig. 4.10 depict a small amount of coupling as evidenced by the small gain values on the off-diagonal elements. This indicates that the 2×2 plants in Fig. 4.4 are actually relatively decoupled.

In order to test the controller's performance in a configuration with severely coupled channels, two alternate versions of the experiment were conducted. The red plots in Fig. 4.10 depict the gains for the configuration in which the controller operates entirely in the raw (endpoint measurement) coordinates. In this configuration the input and output of the adaptive controller are in raw coordinates and the appropriate plant model for the controller is $\hat{\hat{G}} = T\hat{G}T^{-1}$, where $T = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & -1/2 \end{bmatrix}$. There is significantly more coupling in this case, as the off-diagonal gains are substantially larger than in the blue plots for which the input and output of the adaptive controller are in the transformed (rotation and translation) channels.

In another configuration, the input to the adaptive controller is in the raw coordinates, but the output of the controller is in the transformed coordinates. This extreme case for which $\hat{\hat{G}} = \hat{G}T^{-1}$ is depicted by the green plots in Fig. 4.10 and indicates the largest amount of coupling among the three different configurations.

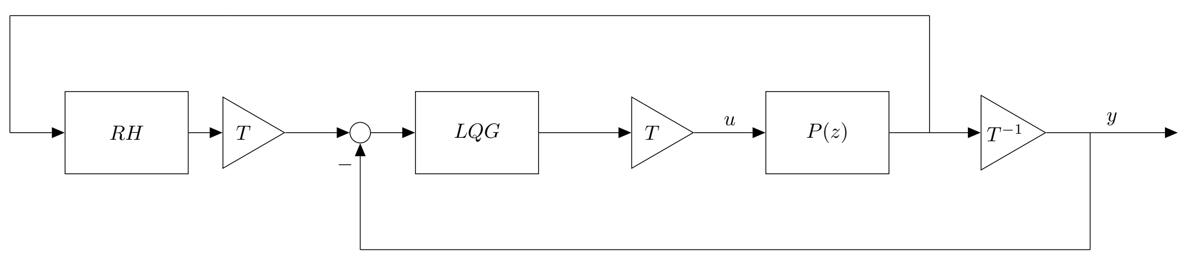

Figure 4.11: Block Diagram Control Structure for $\hat{\hat{G}} = T\hat{G}T^{-1}$

Figure 4.11 shows the block diagram structure for the raw coordinate configuration. Where typically the measurement, y is fed back into the adaptive controller as transformed signals representing translation and rotation, in this case the raw measurement is fed back into the adaptive controller. Additionally, the output bock of the adaptive controller, which is in raw coordinates, is transformed back into translation and rotation coordinates.

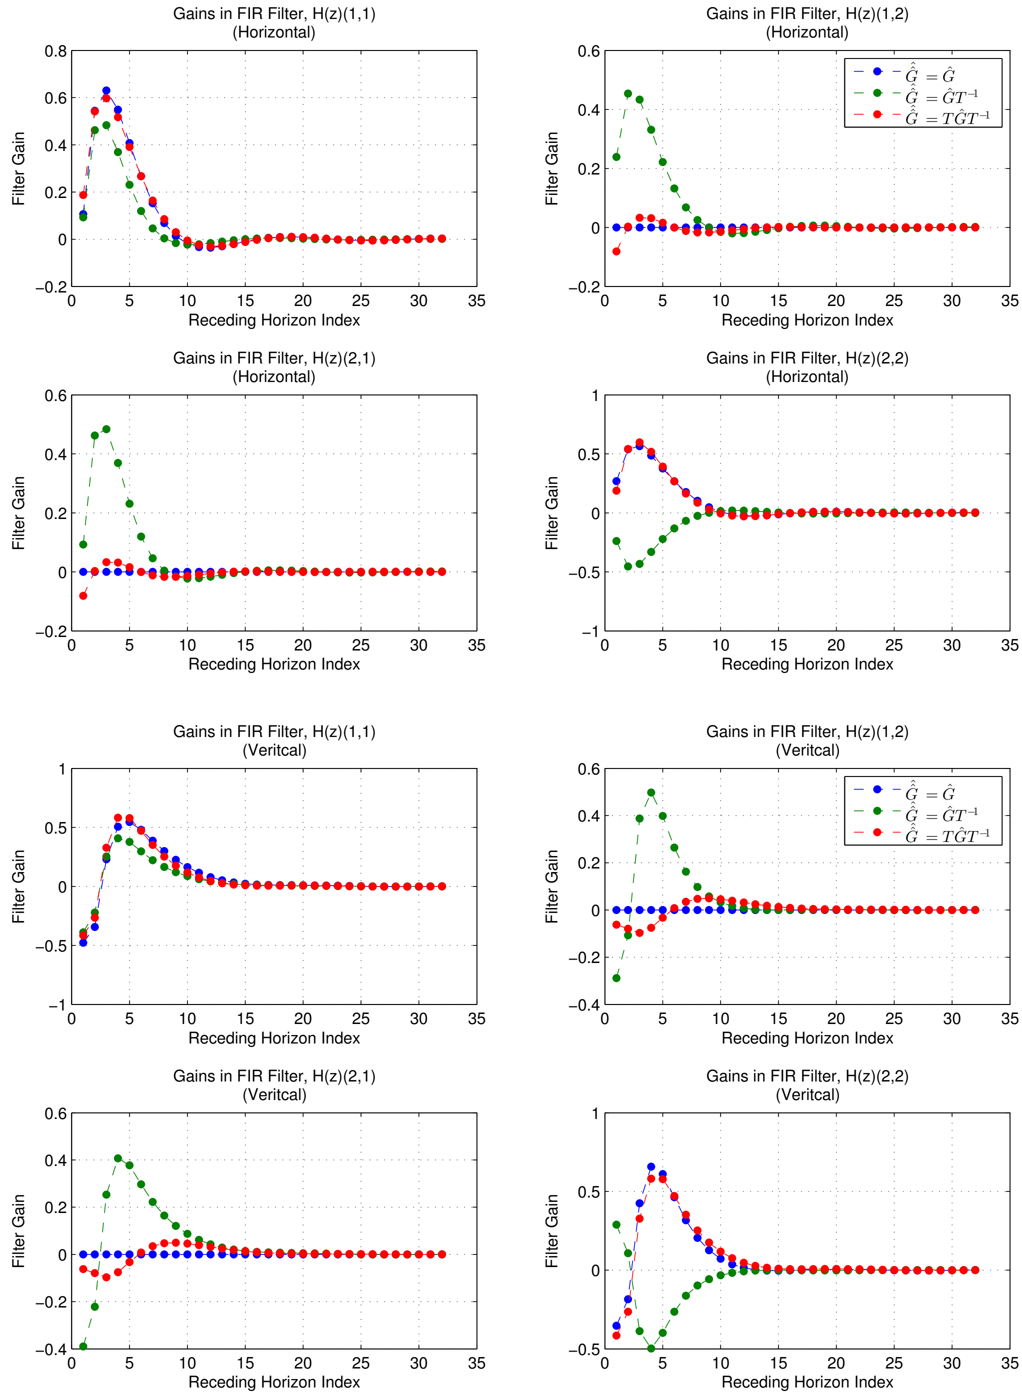

Figure 4.10: $2 \times 2 \times h$ filter gains of $H(z)$ for horizontal (top) and vertical (bottom) channels. In each 2×2 plot, larger gains in the off-diagonal elements indicate a stronger coupling between channels.

The gains for this configuration are shown in Fig. 4.10 in the red plots.

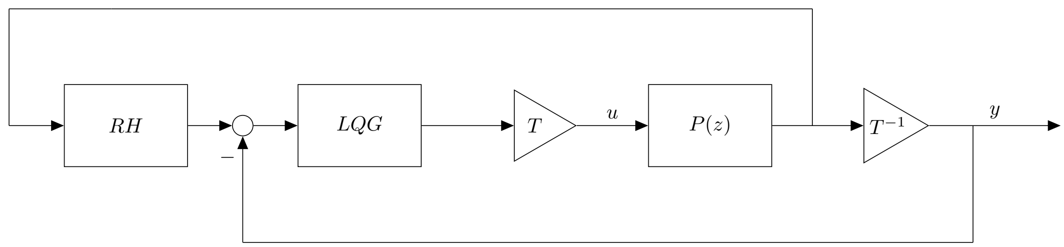

Figure 4.12: Block Diagram Control Structure for $\hat{\hat{G}} = \hat{G}T^{-1}$

Figure 4.12 shows the block diagram for the mixed coordinate configuration. Note that the input to the adaptive controller is in the raw coordinate frame, while the output represents an adaptive control command in the translation / rotation coordinates. The proper plant model for this scenario is $\hat{\hat{G}} = \hat{G}T^{-1}$ and the resulting adaptive filter gains for this case are depicted by the green plots in Fig. 4.10.

Bode plots of the proper plant models are shown in Fig. 4.13 for the raw coordinate case, referred to as *Configuration 1* and for the mixed coordinate case, referred to as *Configuration 2*. Indeed, Fig. 4.13 (Top) shows that when working entirely in raw coordinates, there is substantially more coupling between channels than in the purely transformed coordinates shown in Fig. 4.4. It is interesting to note that in the raw coordinates (Fig. 4.13, Top), there seems to be more coupling in one direction than the other within the vertical plane. In other words, an input at Endpoint 2 produces a larger output at Endpoint 1 than the output at Endpoint 2 due to a similar input at Endpoint 1. This is likely due to the fact that one end of the shaft is heavier than the other because of the attached turbine. This may explain why the phenomenon is more pronounced in the vertical plane, since the effect of gravity is present. In any event, Fig. 4.13 (Bottom) confirms that when working in mixed coordinates, the degree of coupling is the highest among all three configurations, which agrees with the green plots in Fig. 4.10. This is also

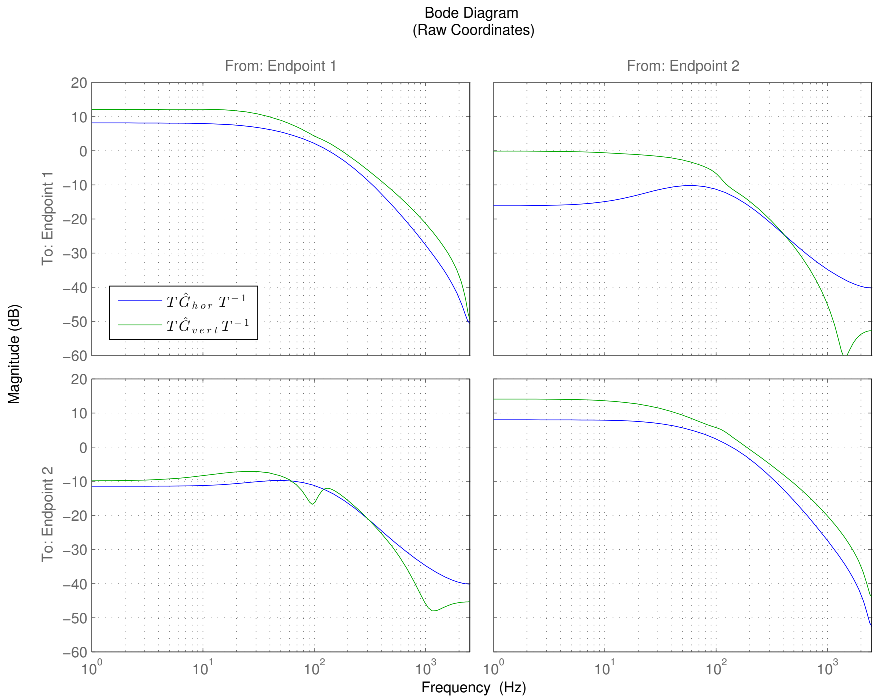

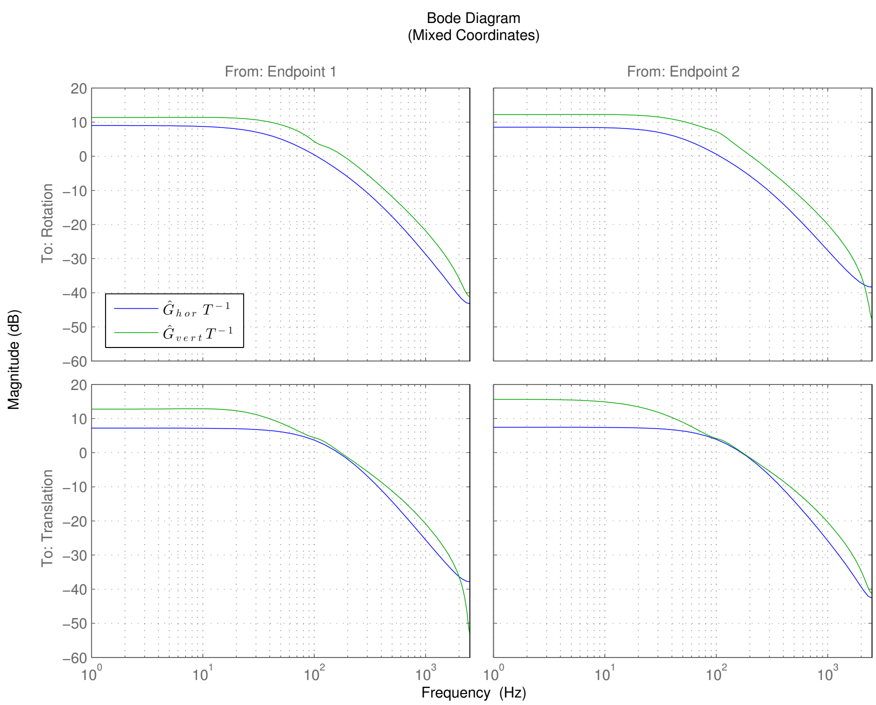

Figure 4.13: Bode magnitude plots for **Configuration 1:** $\hat{\hat{G}} = T\hat{G}T^{-1}$(Top), and **Configuration 2:** $\hat{\hat{G}} = \hat{G}T^{-1}$(Bottom).

logical because an endpoint displacement at either end ought to produce both rotation and translation of the shaft.

Figures 4.14 and 4.15 show time series results for the two configurations with the most coupling between channels (raw and mixed coordinates). The results compare the performance of the LQG control loop alone with that of the LQG control loop augmented by the adaptive control loop. In these experiments, the disturbance consisted of narrow-band unbalance of the spinning shaft combined with broadband vibration produced by a piezo-electric shaker under one corner of the platform on which the bearing was mounted, exactly as in Fig. 4.7. The shaker generated significant vibration both the horizontal and vertical translational channels. Table 4.1 shows the sequence of events for each experiment in which the adaptive loop is closed. It is important to note that the spin rate of the shaft increases linearly between 25 and 30 seconds, resulting in a time-varying disturbance over this range.

For the time series shown in Figs. 4.14 and 4.15, the blue plots show the performance of the adaptive controller by displaying the measured output for the 4 channels after $t = 20s$. The green plots in the background show the system output in the absence of the adaptive control and are included to show the effect of the adaptive controller. The magenta plots represent the noise floor generated by taking data with the stabilizing LQG controller closed, but with the shaft being stationary and free of disturbances. These nominal plots are included to depict the best possible performance achievable by the controller. The blue time series plots in Figs. 4.14 and 4.15 show that, when the adaptive control loop closes at 20 sec, the amplitudes of the output error signals are reduced quite substantially in all four control channels.

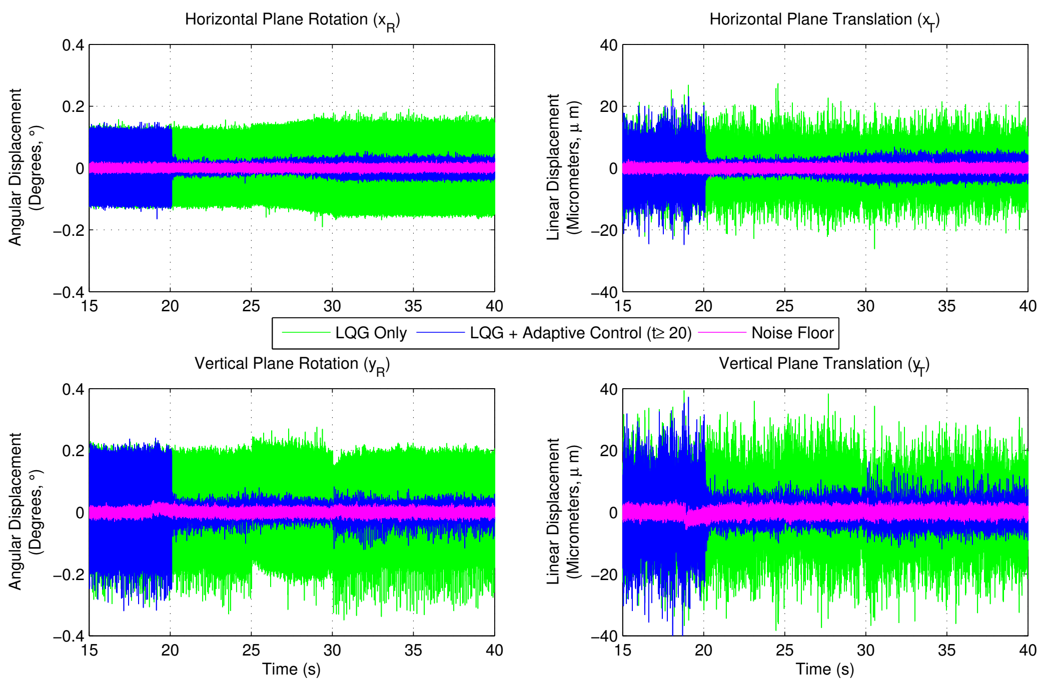

Figure 4.14: **Configuration 1: $\hat{\hat{G}} = T\hat{G}T^{-1}$ (raw input, raw output).** Time series plots for two experiments: one with LQG control only and one with LQG plus adaptive control. The LQG control loop is closed throughout both experiments. Adaptive control begins at 20 sec in one experiment. The disturbance consists of broadband vibration produced by the shaker + narrow-band unbalance of the spinning shaft. The spin rate increases between 25 sec and 30 sec. Table 4.1 gives the complete sequence of events.

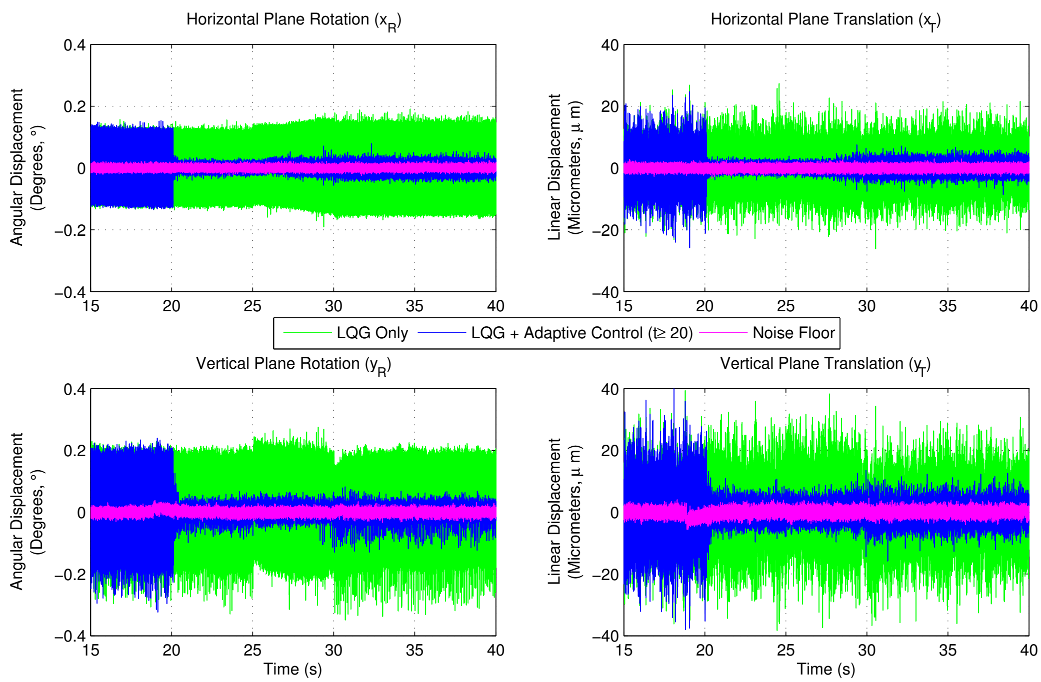

Figure 4.15: **Configuration 2: $\hat{\hat{G}} = \hat{G}T^{-1}$ (raw input, transformed output $\implies$ most coupling).** Time series plots for two experiments: one with LQG control only and one with LQG plus adaptive control. The LQG control loop is closed throughout both experiments. Adaptive control begins at 20 sec in one experiment. The disturbance consists of broadband vibration produced by the shaker + narrow-band unbalance of the spinning shaft. The spin rate increases between 25 sec and 30 sec. Table 4.1 gives the complete sequence of events.

Table 4.2 presents RMS values of the four output error measurements x_R, x_T, y_R, y_T for various controller configurations and for three distinct time ranges. For each time range, the *LQG Only* RMS value represents the magnitude of the output under the influence of the disturbance with no adaptive control, and the *Nominal* RMS value represents the output magnitude with no disturbances. As in Figs. 4.14 and 4.15, this depicts the noise floor, and hence the best possible performance achievable. The three highlighted rows in each time range indicate the performance for three controller configurations. The *Transformed* $\hat{G}$ row indicates the performance of the adaptive controller when operating entirely in the transformed coordinates, as in Fig. 4.7. The *Raw* and *Mixed* rows correspond to *Configuration 1* and *Configuration 2*, and it is apparent that the output RMS values for all three time ranges are roughly the same, indicating that the adaptive controller is capable of producing nearly the same results whether the plant channels are essentially decoupled or severely coupled.

Because the output is non-stationary from $25s \leq t < 30s$ in both configurations ($\hat{\hat{G}} = T\hat{G}T^{-1}$ and $\hat{\hat{G}} = \hat{G}T^{-1}$), the power spectral densities (PSD's) of the output with and without adaptive control are not compared over this range. However, Fig. 4.16 and 4.17 show the output error PSD's for $21s \leq t < 25s$ and $31s \leq t < 35s$, respectively for configuration 2, in which $\hat{\hat{G}} = \hat{G}T^{-1}$. Because the results are similar for both configurations 1 and 2, PSD's are only shown for configuration 2, in which the coupling is most severe. The disturbance in Fig. 4.16 consists of a frequency band from 50-100Hz, and a narrowband peak at the shaft spin rate of 120Hz. Also visible are 2 to 3 additional harmonics of this frequency, which are typical for this application. The most prominent reduction in the output is of the broadband disturbance as well as the primary peak due to the shaft unbalance. There is some amplification of the disturbances from approximately 250-400Hz, which is common for this type of adaptive control scheme which tends to whiten the output error. Overall, the reduction of the broadband and primary narrow-

RMS Values of Output Error				
$\boldsymbol{20 \leq t < 25s}$				
	Channel			
	x_R	x_T	y_R	y_T
LQG Only	0.0822	6.7486	0.0819	9.8267
Transformed ($\hat{G}$)	0.0176	1.5016	0.0270	2.9369
Raw ($T\hat{G}T^{-1}$)	0.0148	1.4531	0.0274	3.2796
Mixed ($\hat{G}T^{-1}$)	0.0148	1.5300	0.0275	3.1364
Nominal	0.0054	0.6219	0.0082	0.9808
$\boldsymbol{25 \leq t < 30s}$				
	Channel			
	x_R	x_T	y_R	y_T
LQG Only	0.0924	6.5029	0.0818	9.8201
Transformed ($\hat{G}$)	0.0125	1.2355	0.0181	2.6126
Raw ($T\hat{G}T^{-1}$)	0.0123	1.2430	0.0183	2.4524
Mixed ($\hat{G}T^{-1}$)	0.0118	1.4789	0.0176	2.5439
Nominal	0.0053	0.6173	0.0081	0.9672
$\boldsymbol{30 \leq t < 35s}$				
	Channel			
	x_R	x_T	y_R	y_T
LQG Only	0.1015	6.1483	0.0726	8.7081
Transformed ($\hat{G}$)	0.0147	1.7186	0.0238	2.9122
Raw ($T\hat{G}T^{-1}$)	0.0162	1.7768	0.0224	2.8661
Mixed ($\hat{G}T^{-1}$)	0.0183	2.1499	0.0223	2.9096
Nominal	0.0053	0.6211	0.0080	0.9585

Table 4.2: RMS values of each output error channel x_R, x_T, y_R, y_T for various controller configurations and for three different time ranges.

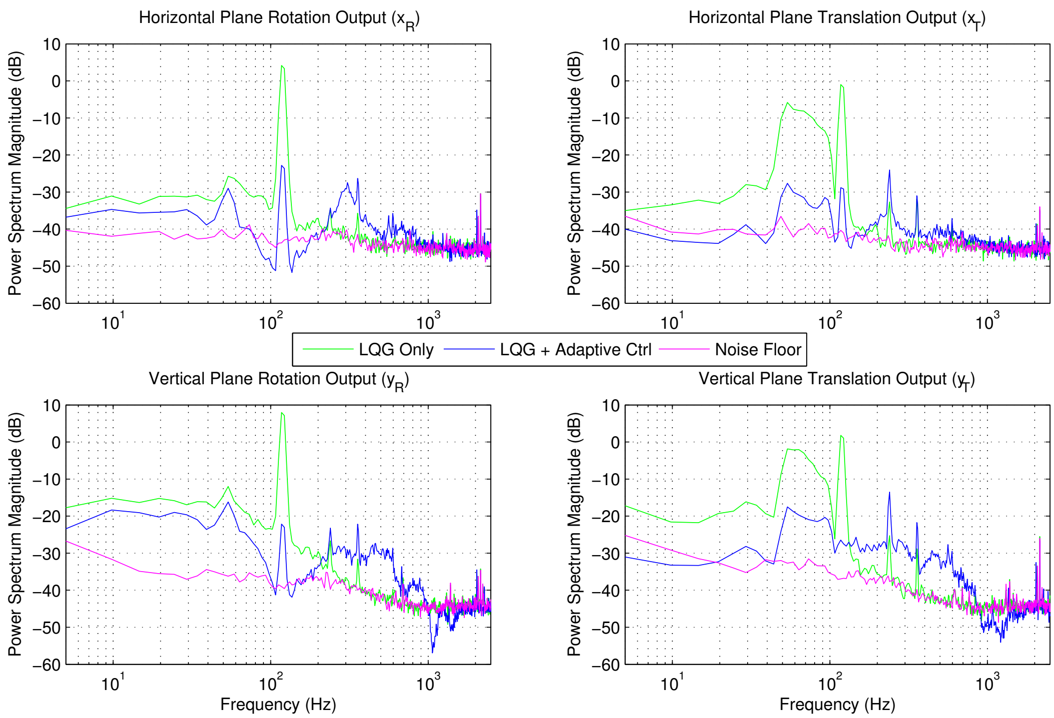

Figure 4.16: Power spectral densities of output error for **Configuration 2:** $\hat{\hat{G}} = \hat{G}T^{-1}$, **(**$21s \leq t < 25s$**)**. The broadband disturbance (50-100Hz) is due to the shaker. The narrowband peak and the associated harmonics are due to shaft unbalance while spinning at 120Hz.

band peak yields a significant reduction in the output error, which agrees with the time series plots in Fig. 4.14.

In Fig. 4.17, the broadband disturbance remains unchanged, but the shaft spin rate has increased and settled at 240Hz. Because this disturbance is very close to exceeding 250Hz, where the plant begins to exhibit nonlinear behavior, the adaptive controller had a more difficult time reducing the output due to the narrowband peak, but was still able to significantly reduce the output error due to the broadband disturbances.

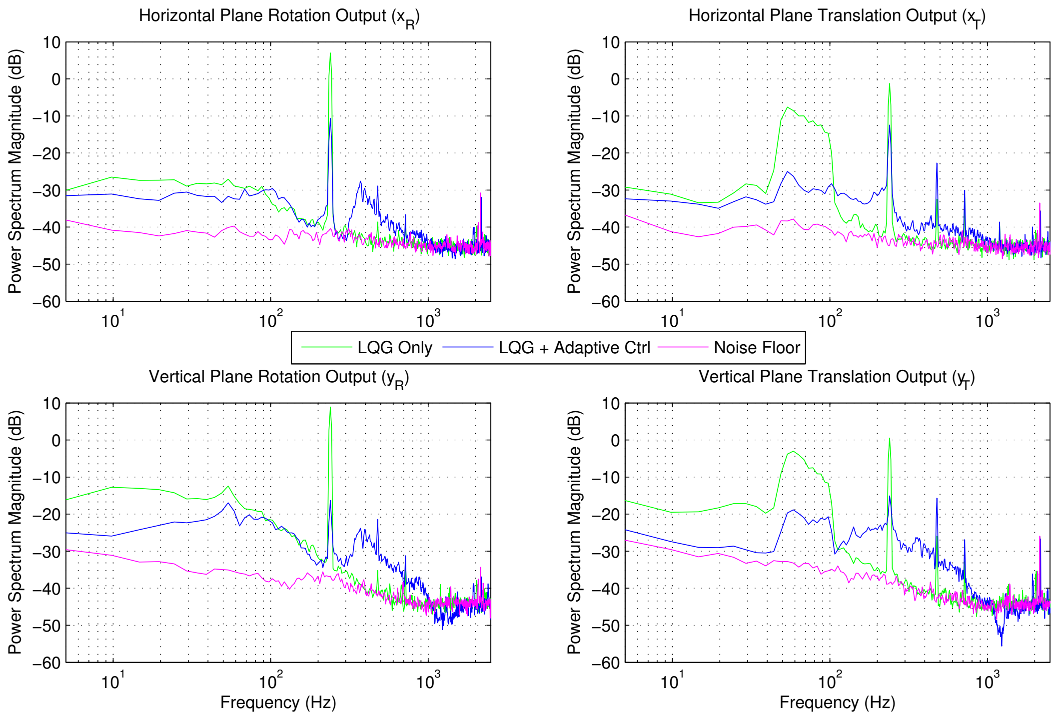

Figure 4.17: Power spectral densities of output error for **Configuration 2:** $\hat{\hat{G}} = \hat{G}T^{-1}$**,** $\mathbf{(31}s \leq t < \mathbf{35}s\mathbf{)}$. The broadband disturbance (50-100Hz) is due to the shaker. The narrowband peak and the associated harmonics are due to shaft unbalance while spinning at 240Hz.

It is apparent from Figs. 4.14-4.17 that the performance in both horizontal plane channels is notably better than in the vertical plane channels. This is due primarily to the fact that the vertical channel plants both contain an unstable zero, making system identification of the plant, and thus control, more challenging. Despite this, the adaptive controller provides significant reduction in the output error on all 4 channels in the presence of a complex disturbance which includes broadband and time-varying components, and regardless of the relative amount of coupling between channels.

CHAPTER 5

Parallel Control for Targeted Disturbance Rejection

5.1 Background

While the receding-horizon adaptive controller is shown to be capable of complex broadband disturbance rejection, in certain circumstances, it may prove advantageous to apply a combination of controllers that serve different purposes. In the case of the magnetic bearing experiment in Chapter 4, a typical disturbance profile consists of a broadband component generated by the shake table support system, as well as a sinusoidal component and its associated harmonics. In this study, a combination of a peak resonator controller and the receding-horizon adaptive controller discussed in Chapter 2 are applied in parallel to reduce a complex disturbance more effectively than either controller alone.

The resonator and receding-horizon controllers are implemented in the structure in Figure 5.1. The controller C (e.g. LQG with integral action) is used to stabilize the 4-channel open-loop unstable bearing system across the vast range of shaft rotation rates ($\geq$ 1000Hz). Because the primary objective of the stabilizing controller is to yield a robustly stable closed-loop bearing system G, the disturbance rejection bandwidth of C is nil, and essentially all of the disturbance rejection performance is the result of the outer-loop controllers C_{Res} and C_{RH}.

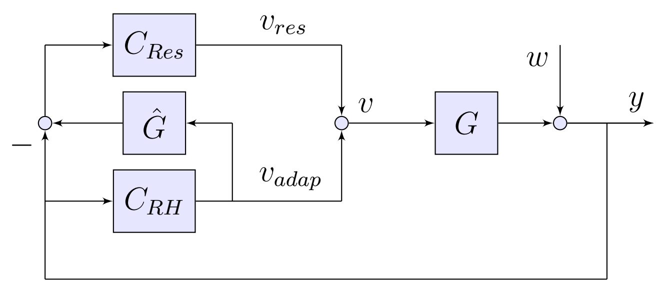

Figure 5.1: Block Diagram of Resonator and Adaptive Controllers

The configuration of the controllers shown in Figure 5.1 is such that the outer-loop controllers C_{Res} and C_{RH} are independent of one another, and each controller can be engaged or disengaged at will. This parallel structure is quite attractive in rejecting complex disturbances as there is no *fighting* between the two controllers. In this study, the resonator is designed to eliminate the sinusoidal components due the shaft rotation and its associated harmonics. Simultaneously and independently, the adaptive controller reduces the effects the broadband disturbances to improve the performance over any one controller alone. Prior to discussing the results of this experiment, an overview of the peak resonator is discussed.

5.2 Peak Resonator - Internal Model Control Design

To reject this harmonic disturbance, an internal model-based controller is examined [18, 19]. Its design flexibilities make it attractive in addressing the subtleties of rotor operations. Furthermore, stability in AMB-rotor application is paramount. To that end, a pre-existing feedback loop provides stability or satisfy other design criteria, while the harmonic resonator C_R employed as a plug-in unit supplies targeted control effort. If the adaptive loop is disconnected, the block diagram in Figure 5.1 can be reduced to that in Figure 5.2. The controller C_R is composed of two filters

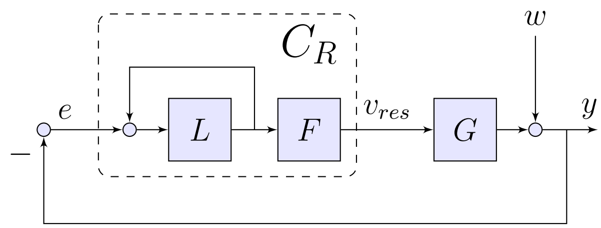

Figure 5.2: Block Diagram of Peak Resonator Controller

$$C_R = \frac{L}{1 - L} \cdot F \tag{5.1}$$

the internal model L and a stable inversion F.

5.2.1 Internal Model Design

The Internal Model Principle requires a model of the disturbance to be placed in the feedback path [20]. The internal model is designed by first designing a p-number of notch filter and the inverting them [18, 19].

Cascaded 2^{nd}-order notch filters can be defined by

$$H = \prod_{k=1}^{p} \frac{1 - 2\beta_k \cos\omega_k z^{-1} + \beta_k^2 z^{-2}}{1 - 2\rho_k \cos\omega_k z^{-1} + \rho_k^2 z^{-2}} \tag{5.2}$$

This formulation makes specifying key parameters of the internal model straightforward. The parameter ω_k in (5.2) represents the location of the notch in the Nyquist band and it is quite apparent that any frequency can be specified, including non-integer values. This characteristic is an attractive trait in rotor operations, as the rotor speed and thus the disturbance frequency can be general.

Furthermore, by using this 2^{nd}-order notch filter as the basis, we can take advantage of the direct control over the characteristic of each notch. Parameters of the filter, ρ and β, shape the depth and bandwidth. They can be tuned independently for each k^{th} frequency. The two factors are chosen to satisfy the relationship $0 << \rho < \beta \leq 1$. In the notch filter, $\beta = 1$ places the zeros on the unit circle producing zero filter output. Once it is inverted, the peak is at unity and will produce infinite gain in the feedback loop. The parameter ρ controls radius of the poles of H. As $\rho \to 1$, the filter becomes more ideal in that frequencies away from the notch are less affected while $\rho < 1$ is maintained to keep the poles inside the unit circle to produce a stable filter.

The relationship between the two factors can be used to approximate the -3dB bandwidth through the expression

$$BW_{Hz} \approx \frac{\pi(\beta_k - \rho_k)}{2\pi T_s} \tag{5.3}$$

where T_s is the sample time of the system. A wider notch provides faster convergence as well as robustness to mismatches between the disturbance and controller frequencies. This must be balanced by the so-called "waterbed effect", which describes the resulting amplification in another part of the spectrum resulting in a loss of performance and stability.

The peak filter is made by inverting the notch filter

$$L = 1 - H \tag{5.4}$$

In Figure 5.2, the internal model is generated by wrapping the peak filter in a positive feedback loop.

$$D = \frac{L}{1 - L} \tag{5.5}$$

5.2.2 Stability Condition and Inversion Filter F

Filter F is included to satisfy a stability condition derived from the Nyquist Stability Criterion. A detailed discussion can be found in [18]. The stability condition can be derived from Figure 5.2:

$$||L(1 - FG)||_\infty < 1 \tag{5.6}$$

This is easily satisfied in the majority of the spectrum by the design of L, where $L \approx 0$ for much of the band. Where $L(e^{j\omega}) = 1$, the condition is met by choosing the filter which satisfies $F(e^{j\omega}) \approx G^{-1}(e^{j\omega})$.

One approach is to invert the entire system. For stable minimum-phase systems, a complete model inversion is easily done. In the MBC500 Turbo, the nonminimum-phase systems of the Y-plane requires an approximate inversion method. The Zero Phase Error Controller (ZPEC) is a useful algorithm which inverts the nonminimum-phase zeros such that the phase is compensated [21]. For a factorization of the stable plant

$$G = \frac{z^{-d}B^{+}B^{-}}{A} \tag{5.7}$$

d represents the relative system order, and A, B^{+}, and B^{-} are the poles, stable zeros, and unstable zeros, respectively. The ZPEC inversion is defined as

$$F_{zpec} = \frac{A[B^{-}]^{*}}{\gamma z^{-d}B^{+}} \tag{5.8}$$

where $[B^{-}]^{*}$ is the complex conjugate of the unstable zeros. This produces the cascaded system $F_{zpec}G = \frac{B^{-}[B^{-}]^{*}}{\gamma}$ which has zero phase since the zeros are complex conjugates of one another.

Depending on the locations of the unstable zeros, the corresponding magnitude response may be unfavorable. Nonminimum phase zeros near the unit circle results in a filter which generates high gains near the Nyquist frequency [22]. To maintain stability, the constant $\gamma = [B^{-}(e^{jw})][B^{-}(e^{-jw})]$ is used to scale the magnitude of the filter output to unity at the prescribed frequency.

5.2.3 Causal Implementation

For strictly proper systems and/or those with nonminimum phase zeros, the resulting inversions will be non-causal. They must be provided with enough "look-ahead" or "preview" steps to form a causal filter. Separating into a causal filter and non-causal previews

$$z^m F_{zpec,caus} = z^m \frac{A[B^-]^*}{\gamma z^{(d+2n_u)} B^+} \tag{5.9}$$

where n_u is the number of unstable zeros and $m = d + n_u$ represent the preview requirement for each respective filter.

A property of the peak filter formulation can be exploited in that L is relative-order 1. This allows the internal model to absorb a preview step. To compensate n-step previews, the notch filter can be up-sampled to

$$H_n = \prod_{k=1}^{p} \frac{1 - 2\beta_k \cos n\omega_k z^{-n} + \beta_k^2 z^{-2n}}{1 - 2\rho_k \cos n\omega_k z^{-n} + \rho_k^2 z^{-2n}} \tag{5.10}$$

Then (5.4) would produce a peak filter with n-relative order. Thus, an n-step preview requirement can be compensated by the up-sampled filter and $L_n z^n$ will remain proper. However, directly applying (5.4) to produce L_m will cause $(m-1)$ aliased peaks to show up in the band at locations

$$\omega_{k,i} = \left| \left(\omega_k \pm \frac{2\pi i}{m} \right) \mod (2\pi) \right| \tag{5.11}$$

where $i \in \mathbb{Z}$ between $0 \leq i \leq \lfloor \frac{m}{2} \rfloor$. The aliased peaks are removed while maintaining the correct filter order by combining two filters such that for an m-step preview requirement

$$L_m = L_1 L_{(m-1)} \tag{5.12}$$

5.3 Experimental Results

In the following experimental results, the rotor is spun to 200Hz to generate the harmonic disturbances. The shaker is also connected to generate a broadband vibration between 50-100Hz. The results of three distinct controller configurations are presented to illustrate the differences among the outer-loop controllers. In each case, the same stabilizing LQG controller is applied first, then one of three outer-loop control configurations (Resonator, Receding-Horizon, and both Resonator and Receding-Horzon) is closed at $t = 20s$. It should also be noted that the particular configuration of the broadband shake-table support system primarily affects the translational channels x_T and y_T and very little broadband disturbance is observed in the rotational output channels x_R and y_R.

Figures 5.3 and 5.4 show the results from only closing the resonator control loop; there is no receding-horizon adaptive control in this configuration. From Fig. 5.3, it is immediately apparent that greatest reduction in output RMS is seen in the rotational channels, 92% RMS reduction in x_R and 88% RMS reduction in y_R. On the other hand, the translational channels x_T and y_T experience far less reduction 37% and 27%, respectively. This is explained by examining the power spectral densities shown in Fig. 5.4. In every channel, the resonator is remarkably effective in eliminating the harmonic components in the output, as each of the 3-4 harmonic peaks in each channel is essentially reduced to the noise floor. However, the resonator cannot address the broadband disturbance component seen prominently in the translational channels. The result is that the overall RMS reduction in these channels is minimal.

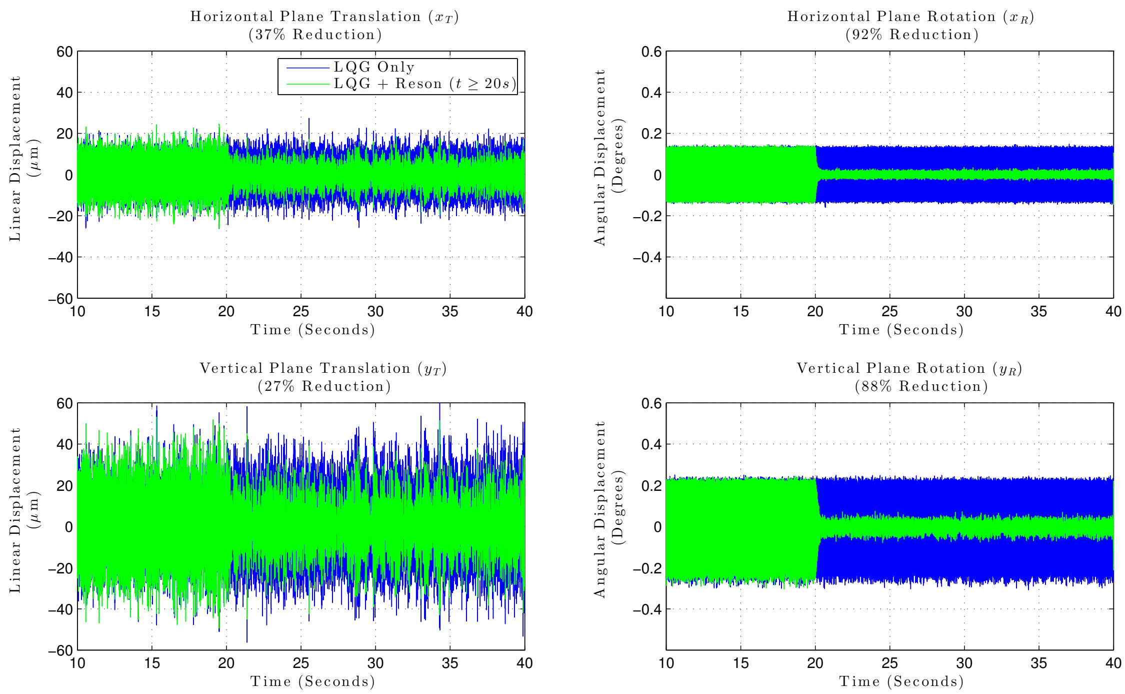

Figure 5.3: Time series plots of output channels with peak resonator control.

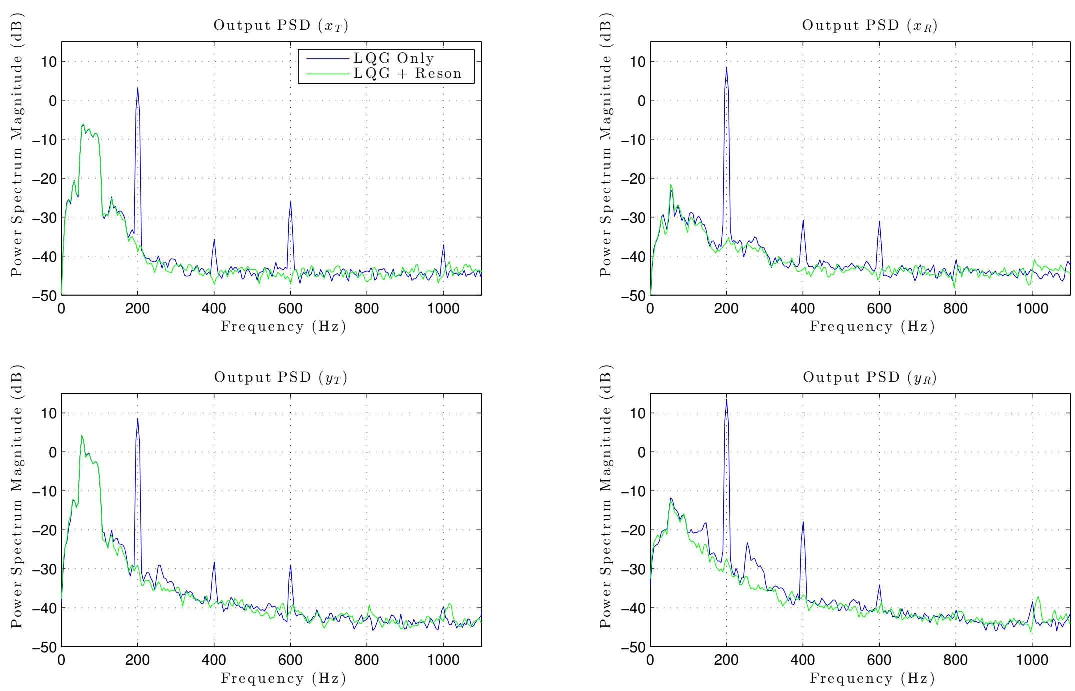

Figure 5.4: Power Spectral Densities of output channels with LQG only and LQG + resonator control for $t \geq 20s$.

The results from closing only the receding-horizon adaptive controller are depicted in Figs. 5.5 and 5.6. The receding-horizon controller, unlike the resonator, attempts to reject all observed disturbances. In this configuration, the RMS reduction is greater in the x_T and y_T channels than with the resonator, but not nearly as substantial in the rotational channels x_R and y_R. The PSDs in Fig. 5.6 indicate that most of the reduction occurs below about 300Hz, which is the cutoff frequency of the frequency weighting filter used with this controller. An examination of the translational channel PSDs indicates a reduction of the broadband disturbance component of about 10dB which explains the greater RMS reduction in the translational channels. Because this adaptive controller is particularly sensitive to the system nonlinearities that occur above about 250Hz, the controller cannot address the higher frequency peaks, and in some cases, tends to amplify the higher frequency components.

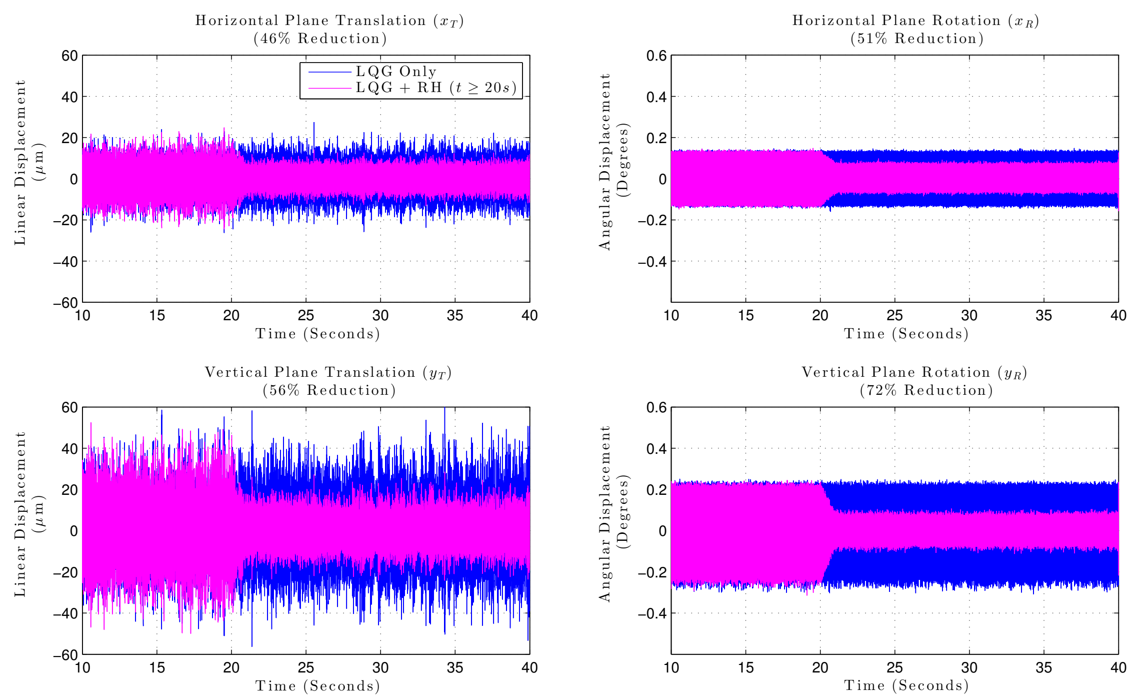

Figure 5.5: Time series plots of output channels with receding-horizon control.

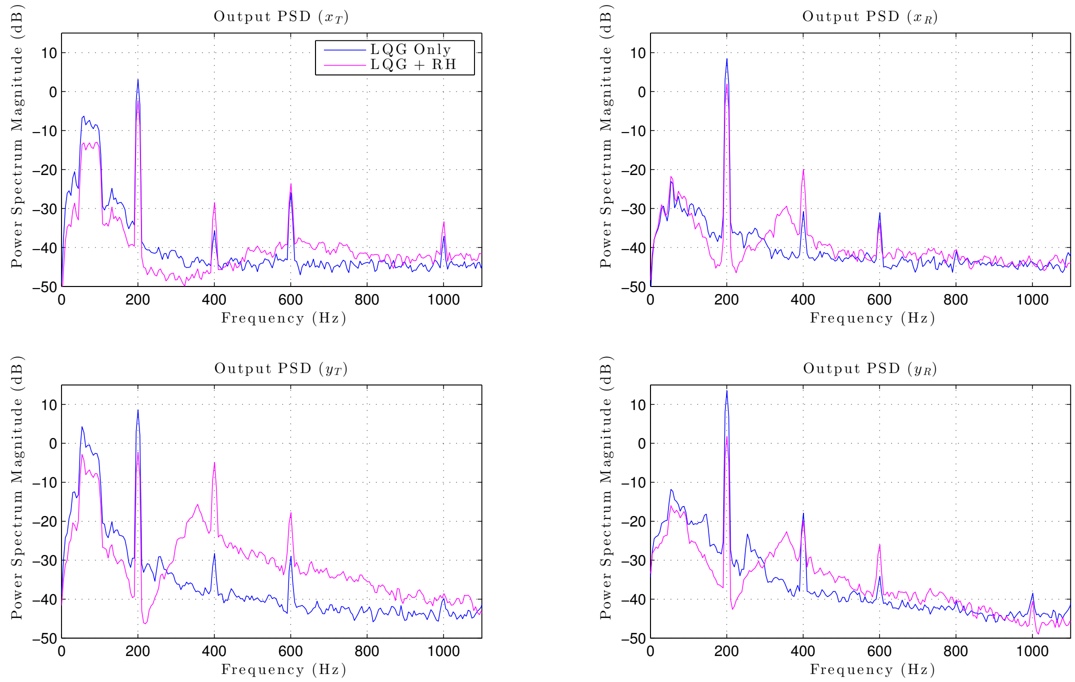

Figure 5.6: Power Spectral Densities of output channels with LQG only and LQG + receding-horizon control for $t \geq 20s$.

In the third configuration, both the resonator and the receding-horizon adaptive controller are engaged. Figures 5.7 and 5.8 show the results from this configuration. It is immediately apparent from the PSDs in Fig. 5.8 that all of the disturbances are effectively rejected. The resonator is able to eliminate the harmonic components just as in Figs. 5.3 and 5.4. However, in this configuration the receding-horizon adaptive controller only observes the 50Hz broadband disturbance component, which it easily attenuates. The time series plots in Fig. 5.7 indicate a reduction in output RMS far superior to either controller alone.

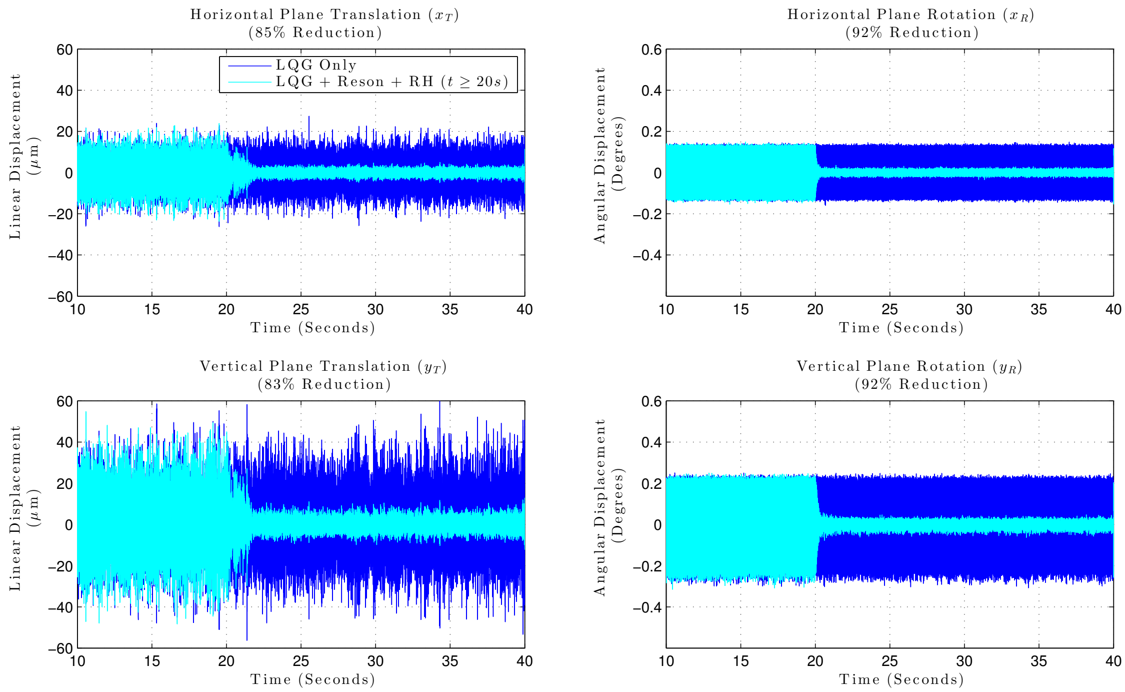

Figure 5.7: Time series plots of output channels with both peak resonator and receding-horizon control.

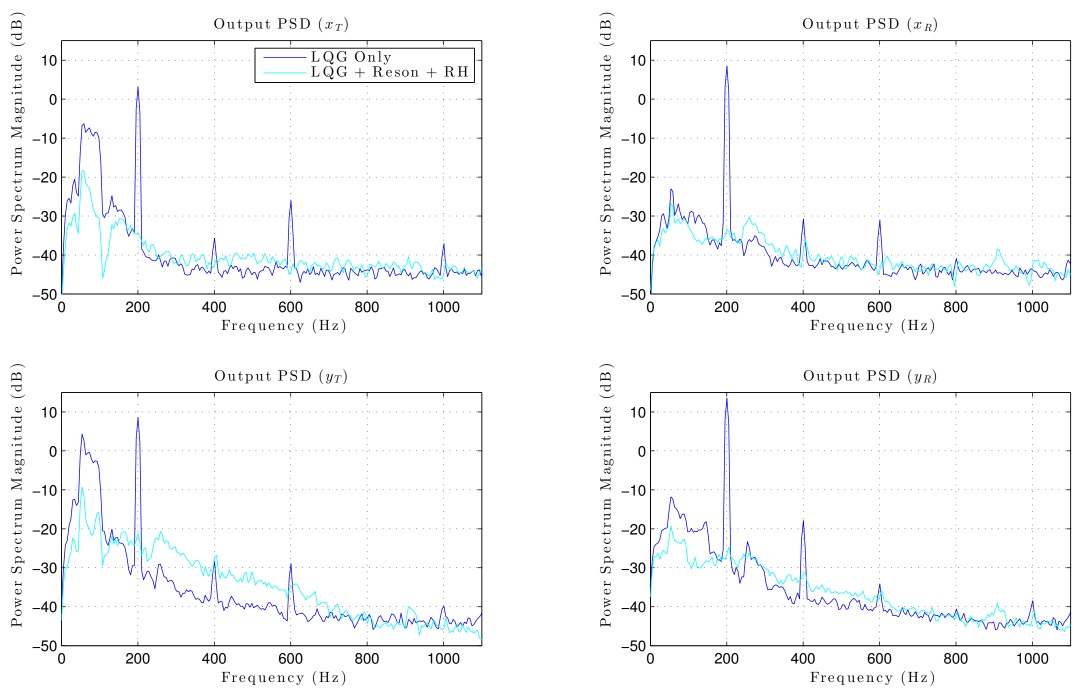

Figure 5.8: Power Spectral Densities of output channels with LQG only and LQG + resonator + receding-horizon control for $t \geq 20s$.

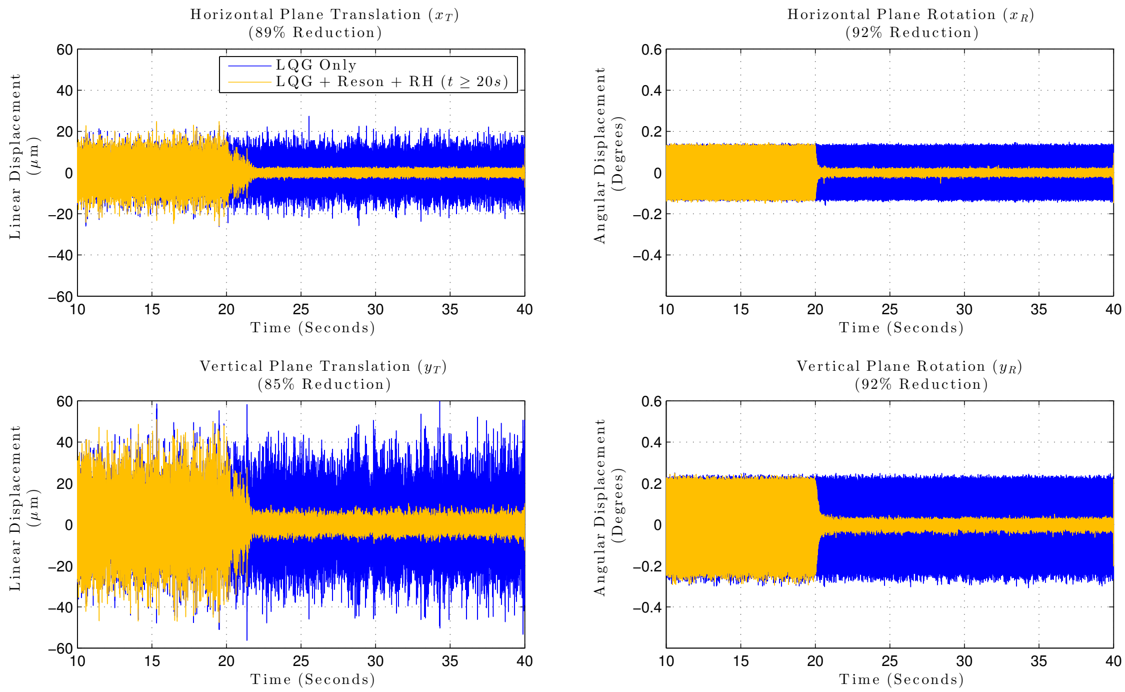

Figure 5.9: Time series plots of output channels with both peak resonator and receding-horizon control (aggressive tuning).

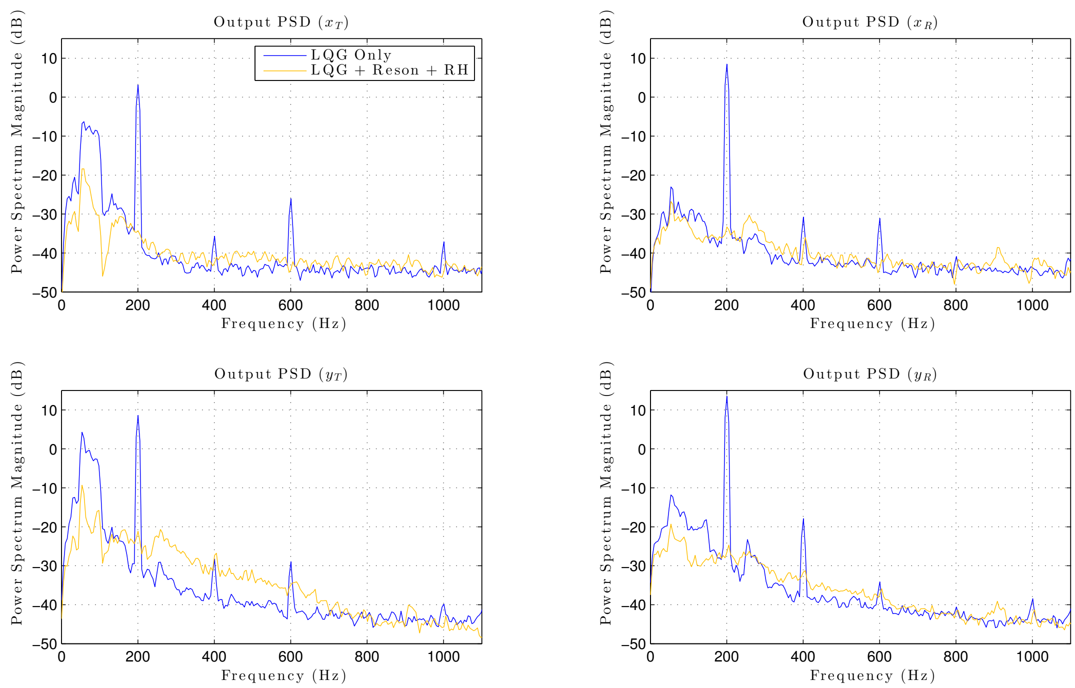

Figure 5.10: Power Spectral Densities of output channels with LQG only and LQG + resonator + receding-horizon control (aggressive tuning) for $t \geq 20s$.

An interesting phenomenon occurred when attempting to tune the receding-horizon control parameters R_u, R_f when applied in combination with the peak resonator and when applied by itself. It was discovered that the receding-horizon parameters can be tuned more aggressively (less control command penalty, smaller R_u, R_f) when applied in combination with the peak resonator. Figures 5.9 and 5.10 depict this case, in which the receding-horizon control parameters were tuned more aggressively than in Figs. 5.5 and 5.7 which utilized the exact same receding-horizon control parameters. The result was improved performance over the combination configuration in Fig. 5.7 in the x_T and y_T channels with a negligible reduction in performance in the x_R channel.

This result alludes to the sensitivity of the receding-horizon controller to plant modeling error ΔG. When the harmonic peaks are not first attenuated by the peak resonator as in the LQG + receding-horizon case, the shaft is spinning at 200Hz and is subject to the mass *unbalance* which causes it to wobble at the primary frequency and its higher harmonics. This implies a larger ΔG because the plant model was obtained with no shaft rotation. In the case of the LQG + resonator + receding-horizon combination control, the resonator first attenuates the harmonic disturbances, so that the receding-horizon controller observes essentially a non-rotating shaft with an additional broadband disturbance. In this case, the plant model is much closer to the true plant and the ΔG is smaller. With a smaller ΔG, the controller can be tuned more aggressively without risk of instability. This phenomenon highlights the sensitivity of the receding-horizon adaptive controller to varying amounts of plant modeling error, which is consistent with the findings in previous chapters.

For the steady-state conditions examined, the commands generated by each controller demonstrate the independent nature of the control structure in Fig. 5.1. Figure 5.11 shows that each controller generates a control command appropriate for reducing one portion of the overall disturbance. That is, the peak resonator only rejects the harmonic components while the receding-horizon controller aims only to reduce the broadband disturbance component. While only the y_T channel commands are depicted in Fig. 5.11, this trend is typical among all control channels. This is a desirable characteristic to observe when using multiple controllers because it indicates that the controllers do not *fight* each other, or attempt to reject the same disturbance components, which can potentially cause instability or saturation.

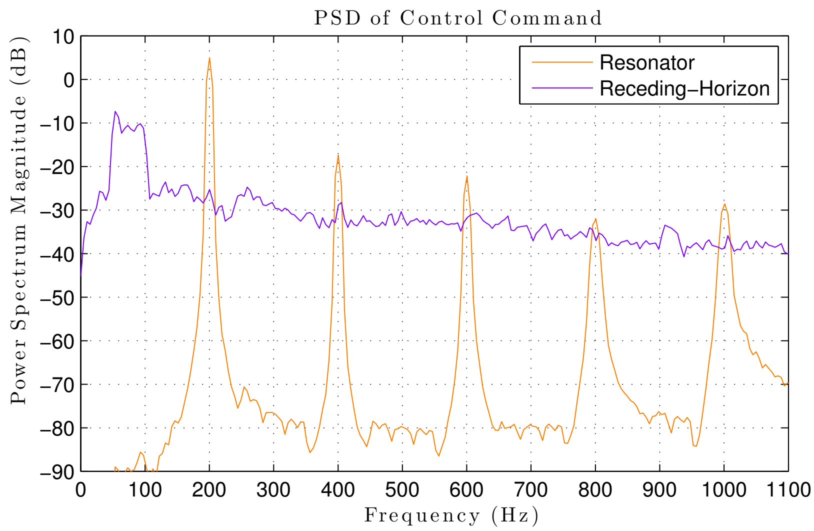

Figure 5.11: PSDs of control commands generated by peak resonator and receding-horizon controller for y_T channel in steady-state.

The RMS reduction of the output error is tabulated in Table 5.1 for each channel and for each controller configuration. It is clear that the configuration in which both the peak resonator and the receding-horizon controller are applied in parallel produces the greatest RMS reduction in output error across all four channels. One may observe that the combination configuration produced only a marginal improvement on the x_R channel when compared to that of the the LQG + resonator case. This is due to the fact that the x_R channel experienced essentially no effect from the broadband shaker, as evidenced by the *LQG only* PSD of the x_R channel in any of Figs. 5.4, 5.6, or 5.8. In this channel, the resonator substantially reduced the harmonic peaks, and thus the receding-horizon adaptive controller had essentially no further disturbance to attenuate. This reiterates the independence of the two controllers, as it was not the case that the two controllers aimed to reduce the same disturbance.

Channel	LQG + Resonator	LQG + RH	LQG + Resonator + RH	LQG + Reson + RH (Aggressive Tuning)
x_T	37.04%	45.75%	84.65%	88.79%
x_R	91.57%	51.48%	92.14%	91.90%
y_T	27.16%	56.06%	82.90%	84.56%
y_R	88.05%	71.58%	92.18%	92.18%

Table 5.1: Percent reduction in RMS value of various controller configurations.

To summarize the results of this study, a peak resonator based on the Internal Model Principle was applied in conjunction with a receding-horizon adaptive controller to reject both harmonic and broadband disturbances in an AMB-rotor system. The controllers were applied in a plug-in architecture to minimally effect the stability of the underlying stabilizing controller while also not affecting the complementary control scheme. The results presented show superior rejection when applying a combination of the two controllers over feedback (LQG) control alone as well as the over each controller applied independently. Further, it is shown that when a combination control architecture is applied, the receding-horizon controller can be tuned more aggressively because the resonator eliminates the harmonic disturbances which otherwise increase the plant modeling error ΔG associated with the receding-horizon controller.

CHAPTER 6

Online System Identification

6.1 Background

A key component of the Receding-Horizon Adaptive Control scheme is having an accurate plant model to generate an estimate $\hat{w}(t)$ of the disturbance. In this research as well as in [5, 6, 10], this model has been generated by a one-time, offline system identification procedure. A white noise sequence of sufficient length ($50,000$ data points at 5kHz) is used to generate input / output data that is processed by MATLAB's *n4sid* subspace algorithm. The result is a model (transfer function) that is essentially the best linear representation of the plant in question, and is henceforth referred to as the *truth* model. This model is subsequently used by the adaptive controller in an experimental run. The limitation of this type of offline identification process is apparent when making the transition to a real-world application. The reason this method is not feasible in practice is twofold:

- This method takes place with no disturbance present. This is axiomatically impractical, as disturbances in real systems are ever present and cannot be paused or temporarily turned 'off' to generate a good plant model.

- This method excites the system with white noise sequences to levels that would be detrimental to the performance of the adaptive controller if executed in realtime. While it is true that the system needs to be excited across all frequencies to obtain an accurate plant model, the overall goal of the adaptive controller is to *reduce* output error rather than amplify it.

Therefore, the goals for the online identification procedure are as follows:

1. Be able to generate a plant model $\hat{G}(z)$ in the presence of a disturbance.

2. Do so without significantly increasing the amplitude of the output error.

3. Achieve the above goals in the shortest amount of time possible.

It should be noted that the above objectives are in contention with one another. To generate an accurate plant model, the output rms must be somewhat increased due to the dither signal, and a sufficient amount of data is required with more data generally producing better results. The subsequent discussion quantifies the accuracy of the plant models generated using variations on the main identification algorithm (Chapter 6.1.1) and analyzes each method.

6.1.1 Overview of Online System Identification Method

The algorithm for generating an accurate plant model relies on a pair of adaptive lattice filters in series that together process the noisy output measurement $y(t)$. Figure 6.1 below shows the block diagram that achieves the desired plant model.

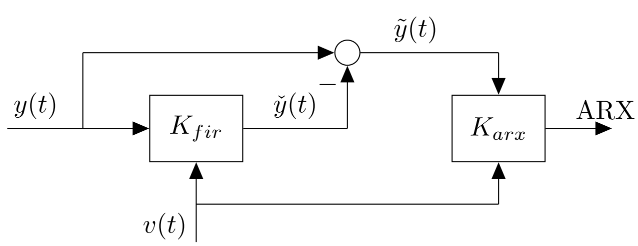

Figure 6.1: Block diagram for online system identification

A small, white dither signal $v(t)$ is injected into the plant such that the rms value of the output error signal is not significantly increased. The first adaptive filter K_{fir} accepts both the dither signal and the total output error $y(t)$ (which includes components due to the dither as well as the disturbance) and projects out the portion of the output that is uncorrelated with the dither signal denoted $\check{y}(t)$. This $\check{y}(t)$ signal is then subtracted from $y(t)$, yielding $\tilde{y}(t)$ which is interpreted as the portion of the output due only to the dither signal. The second adaptive lattice filter K_{arx} then takes $v(t)$ and $\tilde{y}(t)$ and generates the ARX coefficients of the plant. The plant transfer function is then computed and updated in the RH adaptive controller at any arbitrary time specified by the user.

When applied to the receding-horizon control loop, this method allows the plant model $\hat{G}(z)$ to be updated in real time, as opposed to a fixed plant model used indefinitely by the controller. The overall receding-horizon control scheme including the online ID addition is depicted in fig. 6.2 below:

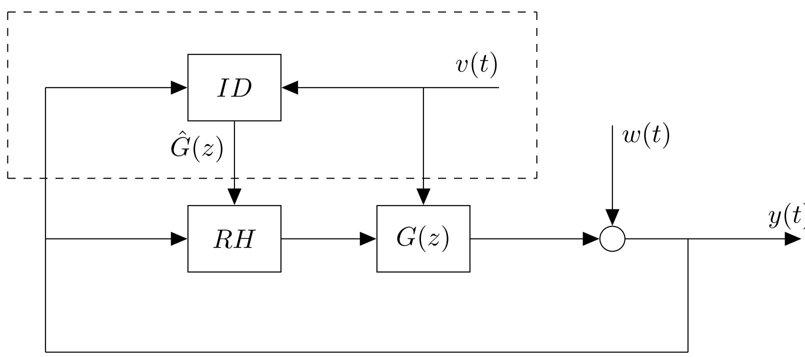

Figure 6.2: Block diagram for overall adaptive control scheme using online ID method

The dashed region denotes the addition of the online ID algorithm.

6.1.2 Variations on the primary method

The method described above generates ARX coefficients of the plant after being passed through two adaptive lattice filters. While this method has been shown to generate quite accurate plant models, other variations on this method may be able to generate plant models of similar accuracy in a shorter period of time.

For example, sections 6.2 and 6.3 analyze tests that conclude optimal filter orders for K_{fir} and K_{arx}. Section 6.4 discusses the effect of the data length required by K_{arx} to produce an accurate $\hat{G}(z)$. Section 6.6 explores the possibility of bypassing the K_{arx} filter altogether. This is made possible by a feature of the K_{fir} filter to compute an FIR realization of the adaptive lattice filter at any arbitrary time.

In order to analyze the effectiveness of this algorithm and study its nuances, an accurate plant model that can be deemed a *truth* model is required. This model should be the best possible linear representation of the true plant, and as such, is generated by the standard offline method discussed in section 6.1.

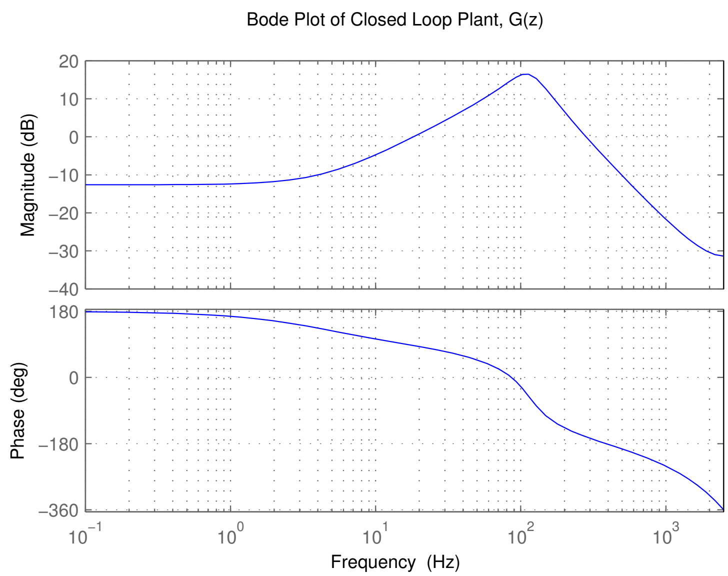

Figure 6.3: Bode plot of *truth* model $\hat{G}_{true}$

Figure 6.3 shows the identified truth model of the closed-loop vertical channel from the laser beam steering experiment described in Chapter 3. This model was generated in the same way as in Chapter 3, wherein the open-loop system was first identified with no disturbance present using MATLABs *n4sid* subspace algorithm. A classical LTI controller was then designed to maximize the disturbance rejection bandwidth with the constraint of minimizing higher frequency amplification. The resulting closed-loop plant was then re-identified using the same process as for the open-loop plant. The subsequent results of applying the online ID algorithm are compared to this truth model $\hat{G}_{true}$ as a measure of the effectiveness of the method.

The output disturbance sequence $w(t)$ (as in Fig. 6.2) used throughout the online system identification research is broadband (175Hz - 275Hz) and is depicted in the figure below:

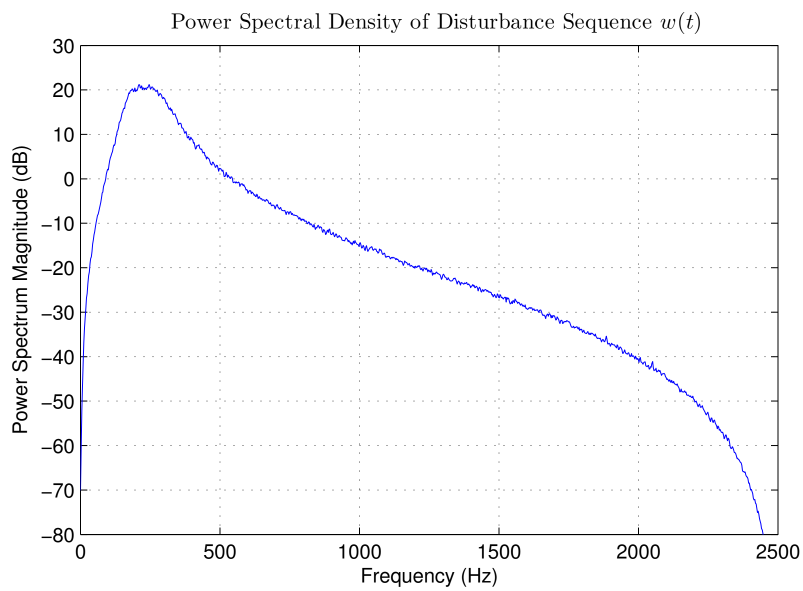

Figure 6.4: Power Spectral Density of Jitter $w(t)$

Additionally, unless otherwise noted, the amount of the white noise dither sequence $v(t)$ added to the plant in all of the subsequent simulation results is such that the output rms of $y(t)$ is increased by roughly 7%, which satisfies the second of the three goals of the online ID procedure listed in section 6.1.

6.2 Effect of n_{fir}

The goal of the first filter K_{fir} is to project out the portion of the output that is uncorrelated with the dither signal. The resulting FIR realization of the steady-state lattice filter has gains that approximate the system's impulse response. Thus, it is important to choose an appropriate filter order for K_{fir}.

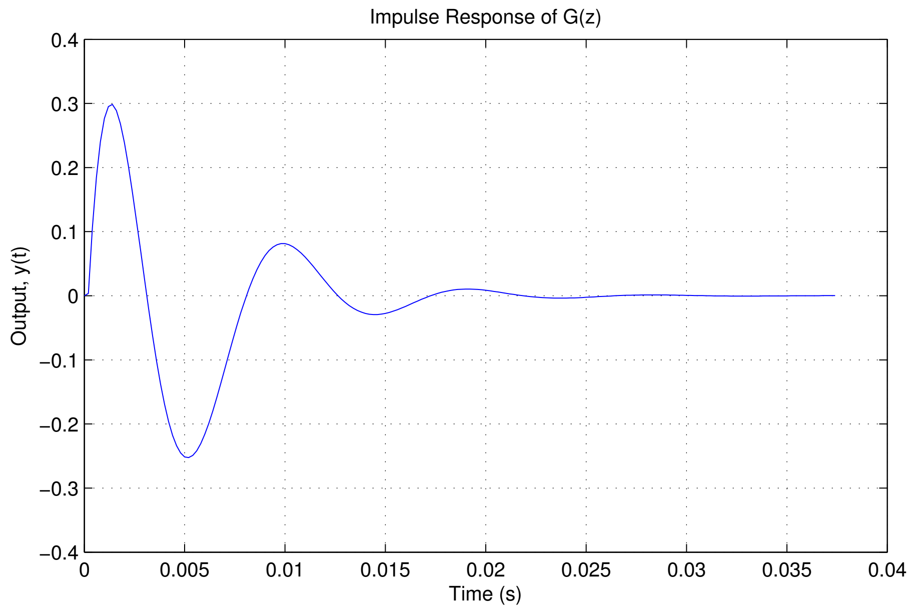

Figure 6.5: Simulated impulse response of *Truth* Model $G(z)$

Too small a filter order would not adequately capture the entire impulse response, but too large a filter order would result in the FIR filter incorrectly attempting to identify many terms that are essentially zero. This would generate erroneous data and introduce error. The simulated impulse response in fig. 6.5 shows a 1% settling time of about $0.025s$, or 125 samples at 5kHz. Thus, the filter order for K_{fir} was chosen to be $n_{fir} = 125$. Figure 6.6 below shows the effect of using too small or too large a filter order for K_{fir} for one particular dither / jitter combination.

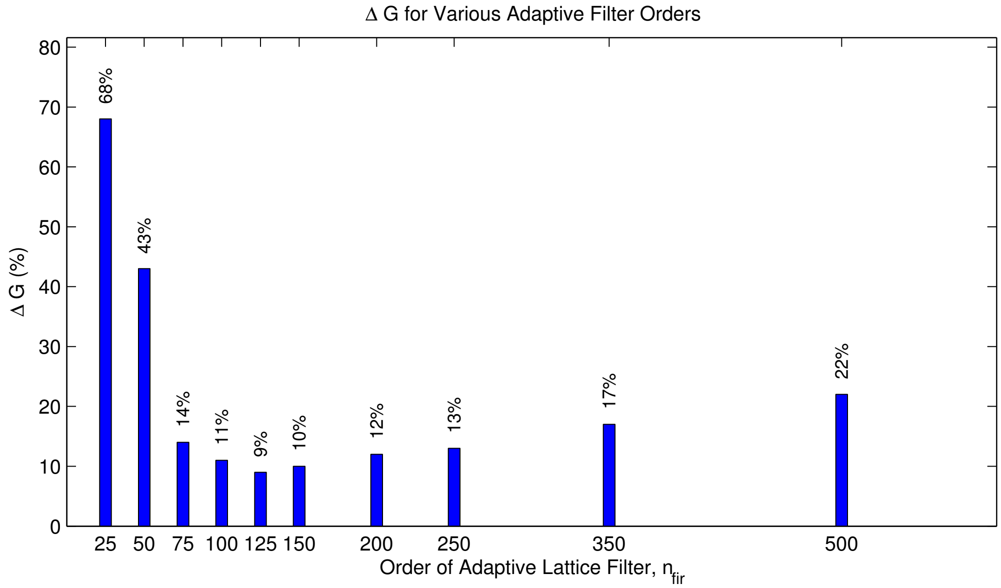

Figure 6.6: $\Delta\hat{G}$ for various orders of K_{fir}. $\Delta\hat{G} = \dfrac{||\hat{G}_{true} - \hat{G}_{arx}||_2}{||\hat{G}_{true}||_2}$

$\hat{G}_{arx}$ used in the computation of $\Delta\hat{G}$ in fig. 6.6 is the plant model given by the second lattice filter K_{arx}, with the model order fixed at $n_{arx} = 8$. The truth model is referred to as $\hat{G}_{true}$. This plot iillustrates the effect of the n_{fir} parameter and thus justifies the choice for the n_{fir} value; it is not indicative of the best possible performance in terms of the smallest $\Delta\hat{G}$.

While a smaller value of n_{fir} corresponds to faster convergence of the filter and thus a shorter time to obtain a plant model, it is clear from Fig. 6.6 that a minimum filter order is required to maximize the accuracy of the plant model. To visualize the rate of convergence of the K_{fir} filter, the dither signal is first passed through $\hat{G}_{true}$ in a parallel simulation, and the output is denoted y_g. This is the output due only to the dither sequence and should ideally be equal to $\tilde{y}$ in steady-state.

Plotting the error between y_g and $\tilde{y}$ reveals how long the filter takes to reach steady state:

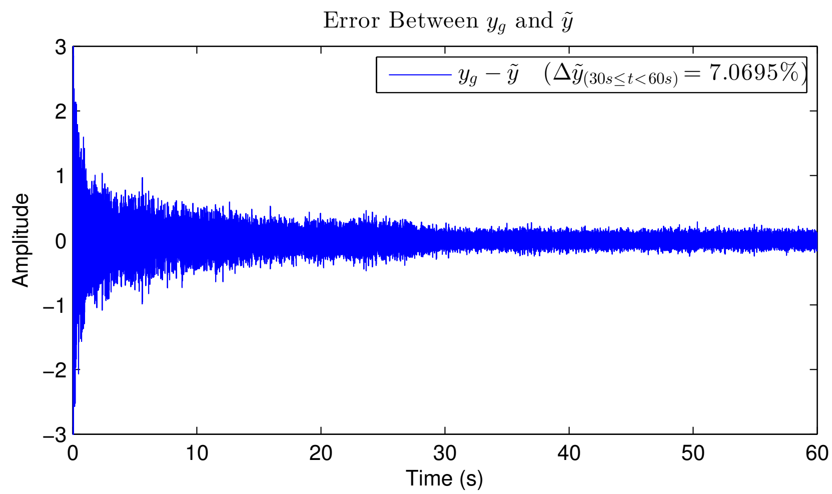

Figure 6.7: Error between y_g and $\tilde{y}$ shows K_{fir} reaches steady-state gains in approximately $30s$.

Figure 6.7 shows that for a filter order of $n_{fir} = 125$, the adaptive filter takes about $30s$ to reach steady-state. The steady-state percent error between y_g and $\tilde{y}$ is given by:

$$\Delta\tilde{y} = \frac{||y_g - \tilde{y}||_2}{||y_g||_2} \times 100 \tag{6.1}$$

which is computed for $t > 30s$. It should also be noted that this error in steady-state is zero when there is no disturbance present, as expected. A preliminary look at the efficacy of this online ID method is depicted in Fig. 6.8 below:

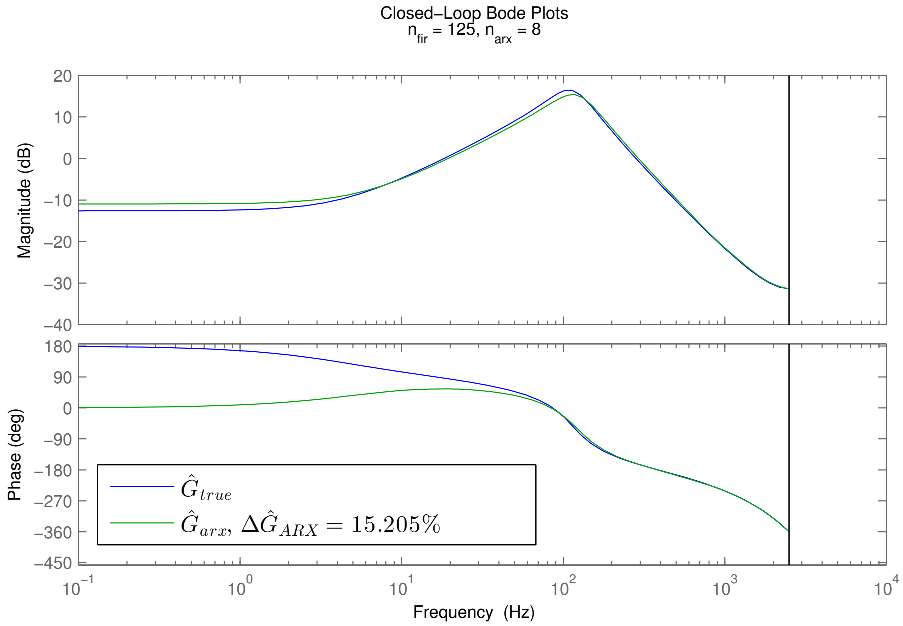

Figure 6.8: Bode plots for $\hat{G}_{true}$ and $\hat{G}_{arx}$

The percent modeling error is quantified as:

$$\Delta\hat{G}_{arx} = \frac{||\hat{G}_{true} - \hat{G}_{arx}||_2}{||\hat{G}_{true}||_2} \times 100 \tag{6.2}$$

and represents how close the online ID method comes to the truth model. Preliminary experimental tests have shown that a plant similar to this ($\Delta G \approx 15\%$) may be adequate as a $\hat{G}$ to use in the full receding-horizon adaptive jitter control experiment, but the performance invariably suffers and consistent stability is questionable. Because the performance of the receding-horizon adaptive controller is dependent on an accurate plant model, it is critical to identify a plant model as close as possible to the true plant.

Another visualization of the accuracy of this plant is to compare the impulse response of $\hat{G}_{arx}$ to the simulated impulse response of the true plant, as shown in Fig.6.9.

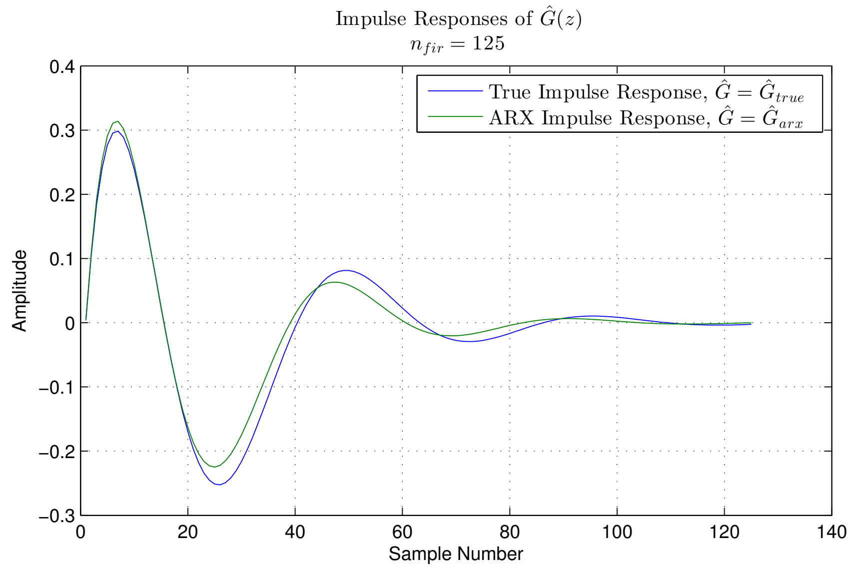

Figure 6.9: Simulated impulse responses for $\hat{G}_{true}$ and $\hat{G}_{arx}$

Again, it is clear that there is a notable discrepancy between the impulse response of the true plant and that of $\hat{G}_{arx}$. With an appropriate K_{fir} filter order chosen ($n_{fir} = 125$) based on Fig. 6.6, these preliminary results show promise in the online ID algorithm. However, a more careful analysis of the K_{arx} filter order (n_{arx}) is necessary to more fully understand the capabilities of the algorithm.

6.3 Effect of n_{arx}

While the model of the true plant $\hat{G}_{true}$ has 3 states, it was shown that a ΔG of about 15% resulted from choosing an n_{arx} value of 8. The following plot indicates that a minimum order of $n_{arx} \approx 24$ is required to achieve the smallest possible $\Delta\hat{G}$.

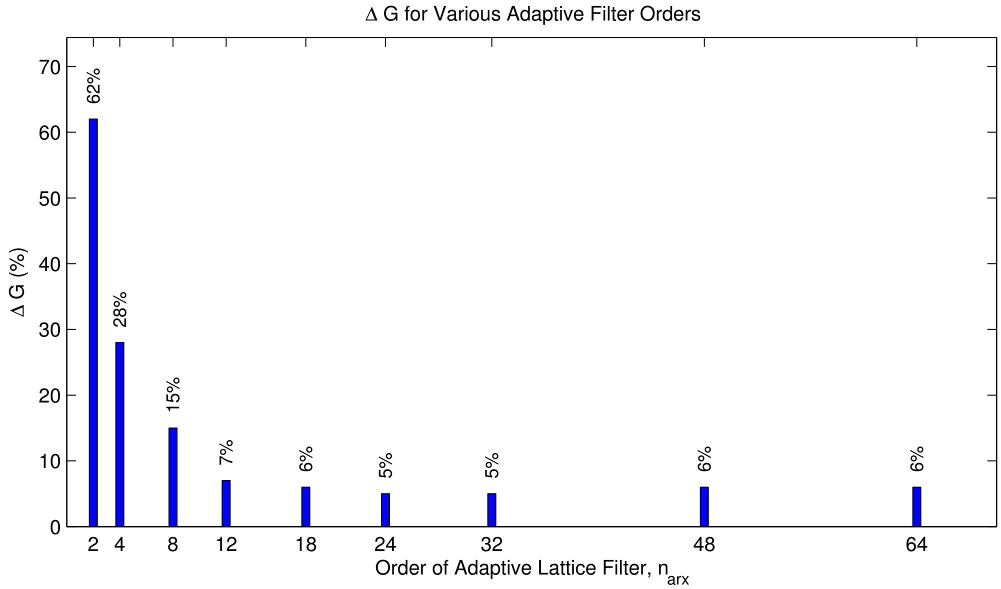

Figure 6.10: $\Delta\hat{G}$ for various orders of K_{arx}. $\Delta\hat{G} = \dfrac{||\hat{G}_{true} - \hat{G}_{arx}||_2}{||\hat{G}_{true}||_2}$. $K_{fir} = 125$.

The reason a seemingly overparameterized filter generates better results is likely a result of insufficient data lengths. Even with 150,000 data points being used for the ARX identification, the dither and jitter sequences are not completely uncorrelated (this can only occur as data lengths approach infinity). The result is that the $\tilde{y}$ sequence used by the ARX filter contains traces of the broadband jitter sequence, which K_{arx} attempts to identify, but can only do so with sufficiently large filter order.

Figure 6.11 below shows that, with $n_{arx} = 24$, the plant model becomes significantly closer to the true plant, with $\Delta G = 5.46\%$ computed in the same manner as in (6.2).

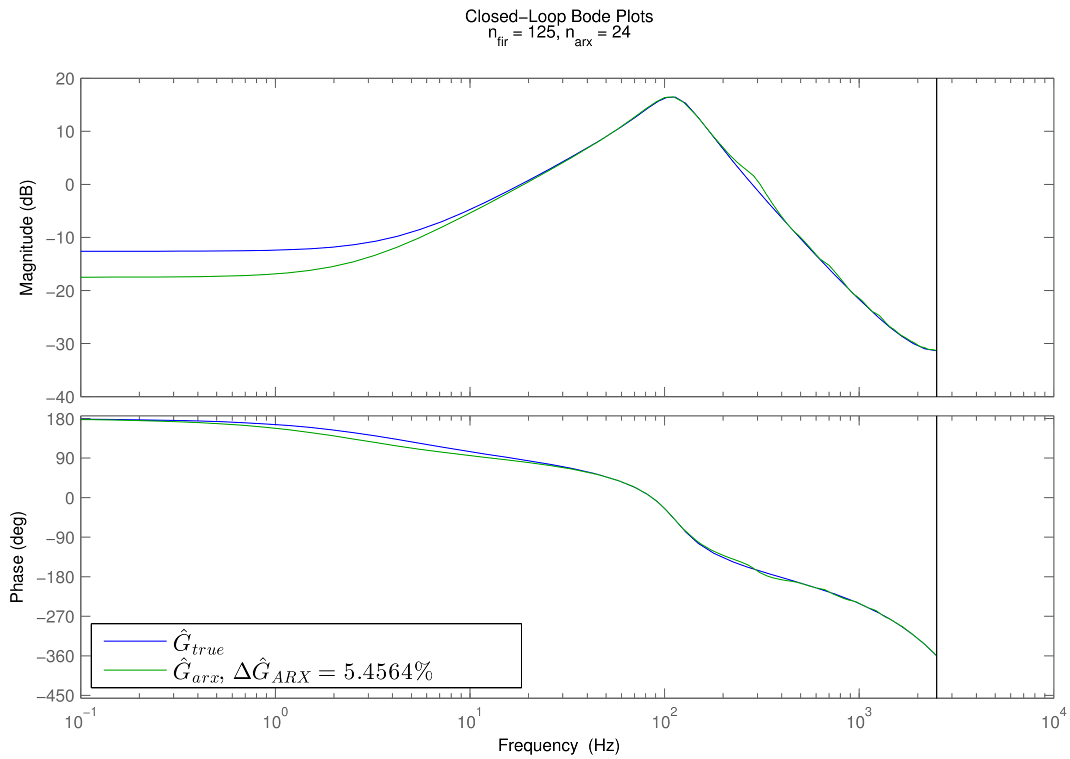

Figure 6.11: Bode plots for $\hat{G}_{true}$ and $\hat{G}_{arx}$ with optimal n_{arx}

6.4 Data sequence length required by K_{arx}

It was shown that the minimum order for K_{fir} was about $n_{fir} = 125$. This filter order requires about $30s$ for the filter to converge to steady-state, effectively imposing a minimum time constraint on the first filter. The analysis in section 6.3 indicated an optimal filter order of $n_{arx} \approx 24$, however, those results were generated by allowing an additional $30s$ (150,000 samples) of steady-state data ($u(t), \tilde{y}(t)$) to be passed into K_{arx} filter prior to generating the ARX coefficients. Figure 6.12 below shows that K_{arx} can produce an accurate plant model with significantly less data.

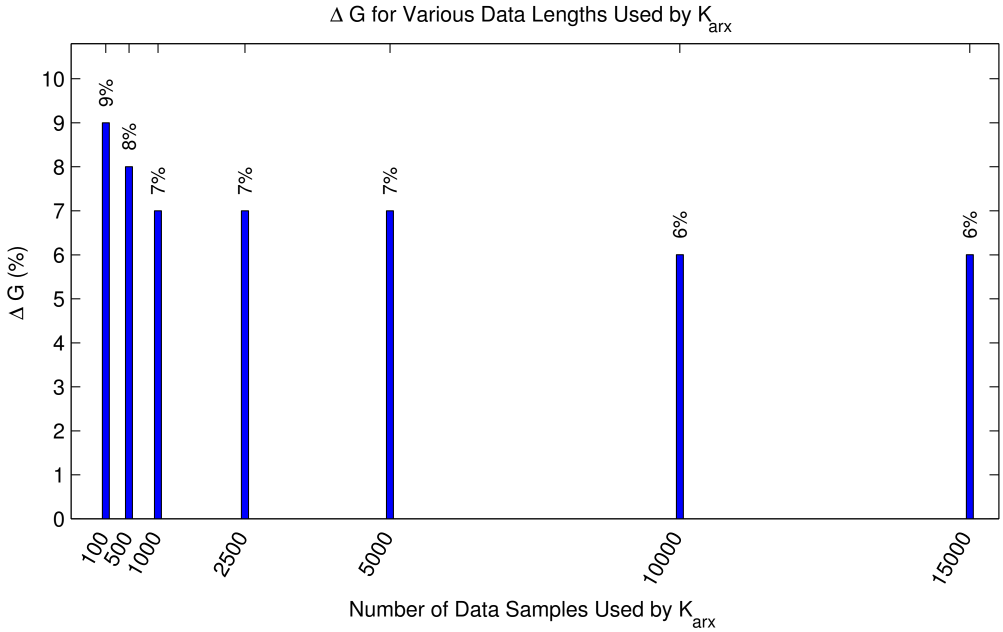

Figure 6.12: $\Delta\hat{G}$ for Various Data Sequence Lengths

Recall from Fig. 6.10 that the smallest $\Delta\hat{G}$ (5%) was obtained using $n_{fir} = 125$ and $n_{arx} = 24$, with the K_{arx} filter using 150,000 data points (30s) before producing ARX coefficients. Figure 6.12 indicates that nearly the same performance can be achieved using only 15,000 data points (3s). Additionally, if some accuracy is sacrificed, a $\Delta\hat{G}$ as small as 7% can be achieved using just 1000 data points (0.2s).

Below are Bode plots of plant models constructed from the ARX coefficients generated using different lengths of data.

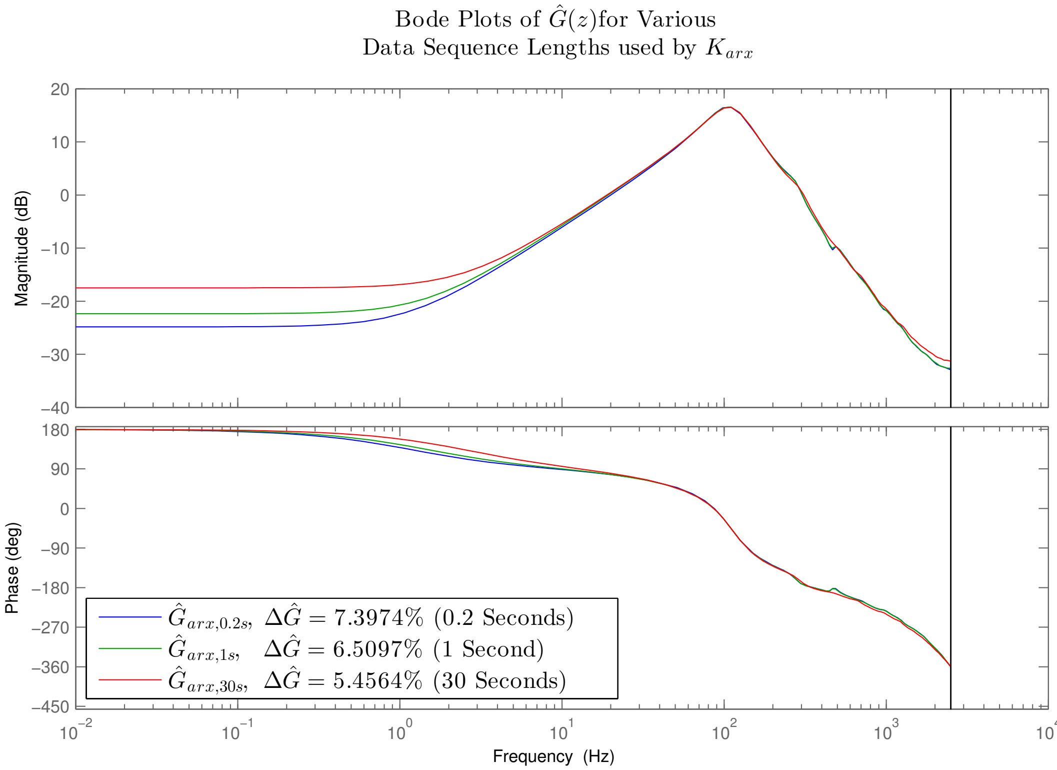

Figure 6.13: Bode plots of $\hat{G}_{arx}$ for various data sequence lengths

Figure 6.13 indicates that the primary difference between using 30 seconds of data vs. 0.2 or 1 second occurs at higher frequencies (> 400 Hz), where the magnitude is relatively small. The main dynamics of the plant (30 Hz $< \omega <$ 350 Hz) seem to be identified almost exactly the same when using 30 seconds or just 0.2 seconds of data.

Note, however, that the results in this section only hold for a starting time of $t = 30s$. That is, the 0.2, 1, and 30 seconds of data used by the ARX filter *began* at $t = 30s$, after K_{fir} had already reached steady-state. The result is not the same when the window of ARX data begins at a value of $t < 30s$. This is because the FIR gains, and thus $\tilde{y}(t)$, have not yet reached steady-state. In this case, the amount of data required by K_{arx} is slightly longer, but still significantly shorter than the original test length of $30s$.

6.5 Data sequence length required by K_{fir}

While the required order of K_{fir} was established in section 6.2, the amount of data points required by K_{fir} was not previously studied. Figure 6.7 shows that, for a given set of parameters, the error between y_g and $\tilde{y}$ reaches steady state at approximately $30s$. While this gives some indication of the amount of data required by K_{fir}, figure 6.14 below reveals a more detailed analysis of $\Delta\hat{G}$ vs. the number of data points used by K_{fir}.

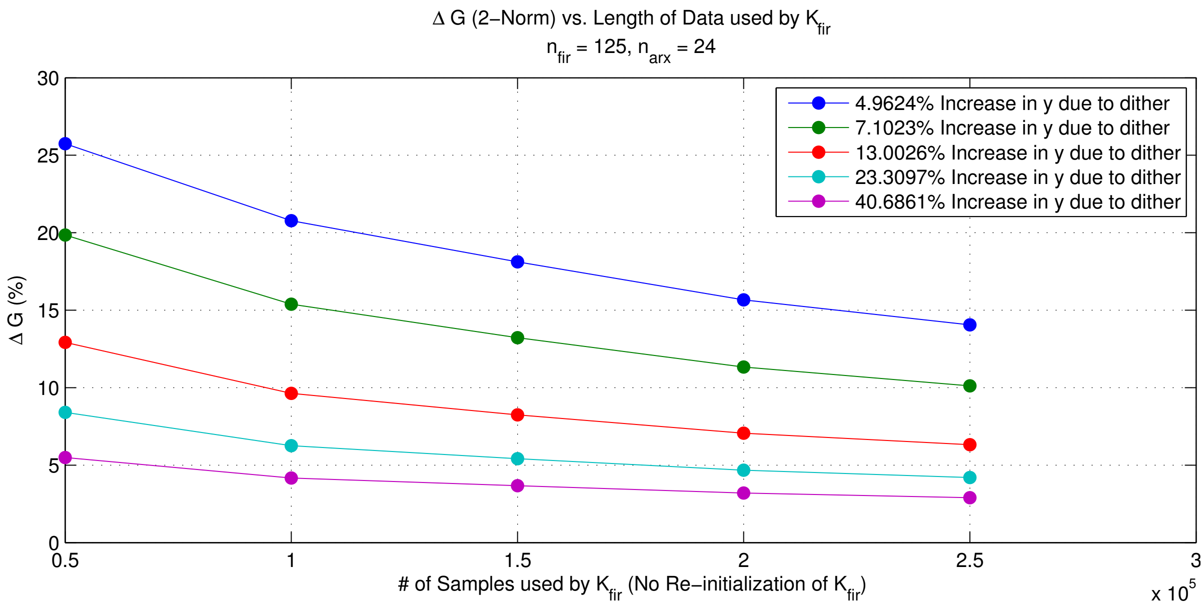

Figure 6.14: $\Delta\hat{G}$ for various data sequence lengths used by K_{arx} (no reinitialization)

Each line in Fig. 6.14 corresponds to a specified amount of dither relative to jitter. The percentages in the legend reflect how much the dither increases the rms value of the output while the disturbance is present. It is shown, as expected, that with a larger amount of dither injected into the plant, the plant model can be identified more accurately than with a shorter dither sequence. Additionally, as more data is used, the plant model generated by K_{arx} becomes a closer representation of the true plant. This is logical, as the accuracy of $\tilde{y}$ gets better with time, ultimately settling to steady-state at approximately $t = 30s$, as seen in Fig. 6.7. One question that arises is then whether the improvement in the plant model is truly due to K_{arx} using significantly *more* data, or simply that the data $\tilde{y}$ used by K_{arx} is closer to the true y_g as more time passes.

Note that in Fig. 6.14, the horizontal axis represents the total amount of data points used by K_{arx}, starting from $t = 0$, since K_{arx} has a forgetting factor of $\lambda_{arx} = 1$. On the other hand, in Fig. 6.15 below, the horizontal axis represents the time at which K_{arx} is reinitialized. Then, each point represents a plant model generated by K_{arx} using the subsequent 50,000 data points.

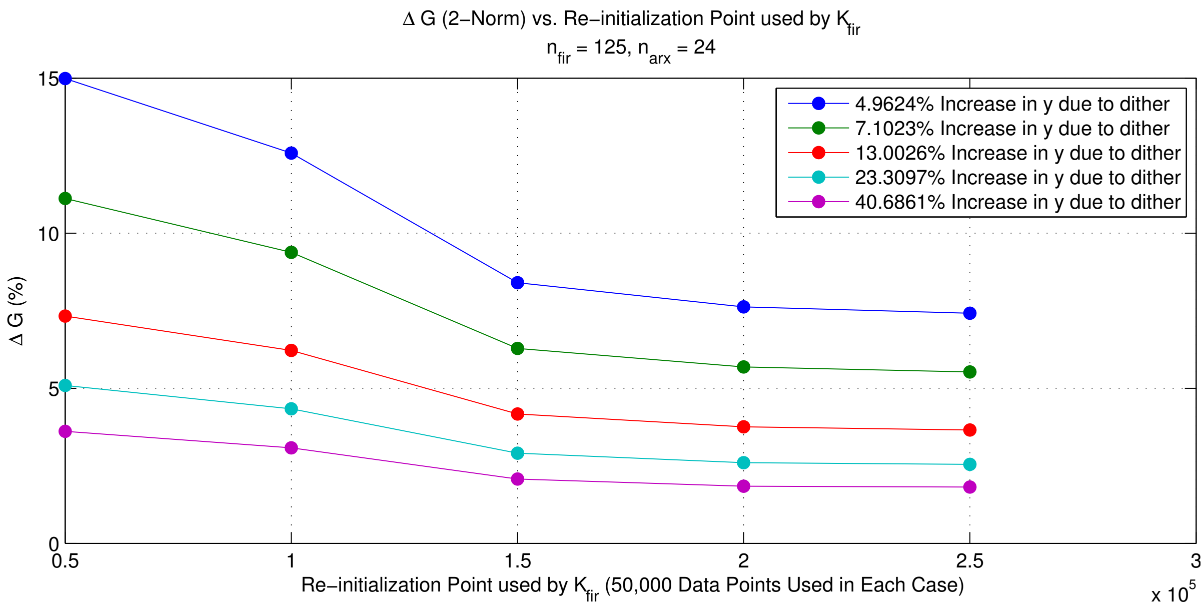

Figure 6.15: $\Delta\hat{G}$ for various re-initialization points used by K_{arx} (re-initialization at times according to horizontal axis)

The trend in Fig. 6.15 is similar to that of Fig. 6.14 in that the model error is reduced as more time passes. Perhaps more importantly, however, is the conclusion that a shorter, but more accurate $\tilde{y}$ sequence (Fig. 6.15) used by K_{arx}, produces a better plant model than using a longer sequence of data (Fig. 6.14), of which some is fairly inaccurate (the transient portion of $\tilde{y}$, see Fig. 6.7). For example, Fig. 6.15 shows that at a 7.1023% increase in y due to the dither (green plot), using $10s$ of data between $20s < t < 30s$ produces a $\Delta\hat{G}$ of 9.38%. On the other hand, the corresponding plot in Fig. 6.14 shows that using all 150,000 data points from $0 < t < 30s$ produces a plant that yields a larger $\Delta\hat{G}$ of 13.2%.

The results in section 6.4 are also consistent when applied to this test. That is, while Fig. 6.15 uses a fixed number of data points (50,000), the trend is similar when fewer data points are used. Specifically, for $t > 30s$, the results are nearly identical when using just 10,000 points (2 seconds) rather than 50,000 points (10 seconds).

While there are several plant models (50) used to generate Fig. 6.14 and Fig. 6.15, not all of them yield stable results when using them in an actual simulation of the experiment. When applying the plant models in Fig. 6.14 and Fig. 6.15 on an adaptive control simulation using the same disturbance, several cases where the $\Delta\hat{G}$ exceeded 10% resulted in instability. While this stability threshold is not a theoretical result, it offers a sense of the acceptable accuracy required by the adaptive controller for stable operation.

6.6 Bypassing K_{arx}

The original online ID method consists of two adaptive lattice filters (K_{fir} and K_{arx}) and it was shown that, with properly chosen filter orders, $\Delta\hat{G}$ could be made very small. It was also shown that the amount of time required to generate an accurate plant model could be significantly reduced by feeding K_{arx} fewer data points. Another variation is to eliminate the need to for K_{arx} by taking advantage of the ability of K_{fir} to compute the FIR realization of the lattice filter at any specified time. The corresponding set of FIR gains produced by the K_{fir} filter is theoretically identical to the impulse response of the system $G(z)$, however, in practice it is corrupted by the disturbance. By then performing an SVD (singular value decomposition) on a Hankel matrix computed from the FIR gains (an approximation of the impulse response), the Controllability and Observability matrices can be constructed and the state-space matrices of the plant G can be generated.

Figure 6.16 illustrates a group of results for when the FIR gains are generated using only 10,000 data points, or 2 seconds of data. Figure 6.16 (bottom) shows the $n_{fir} = 125$ gains generated by K_{fir} at $t = 2s$ superimposed on the simulated impulse response of the true plant $\hat{G}_{true}$. As expected, the FIR gains are significantly different from the impulse response, as the FIR gains have not yet converged (convergence occurs at approximately $t = 30s$ according to Fig. 6.7). The percent difference between the simulated impulse response h_{sim} and the FIR gains are computed as:

$$\Delta_{\text{FIR}} = \frac{||\text{FIR gains} - h_{sim}||_2}{||h_{sim}||_2} \times 100 \tag{6.3}$$

and is equal to about 12.3% as indicated.

What is surprising about this case is the accuracy of the plant model when a 3-state model is identified. Figure 6.16 (top left) indicates that the $\Delta\hat{G}$ computed as in equation (6.2) is only about 2.57%, when compared to $\hat{G}_{true}$. Extracting such an accurate plant model with only 2 seconds of data (as opposed to requiring 30 or more seconds as in previous sections) is indeed noteworthy, however, a closer look at the conditions required for this to occur reveal that it may be impractical in practice.

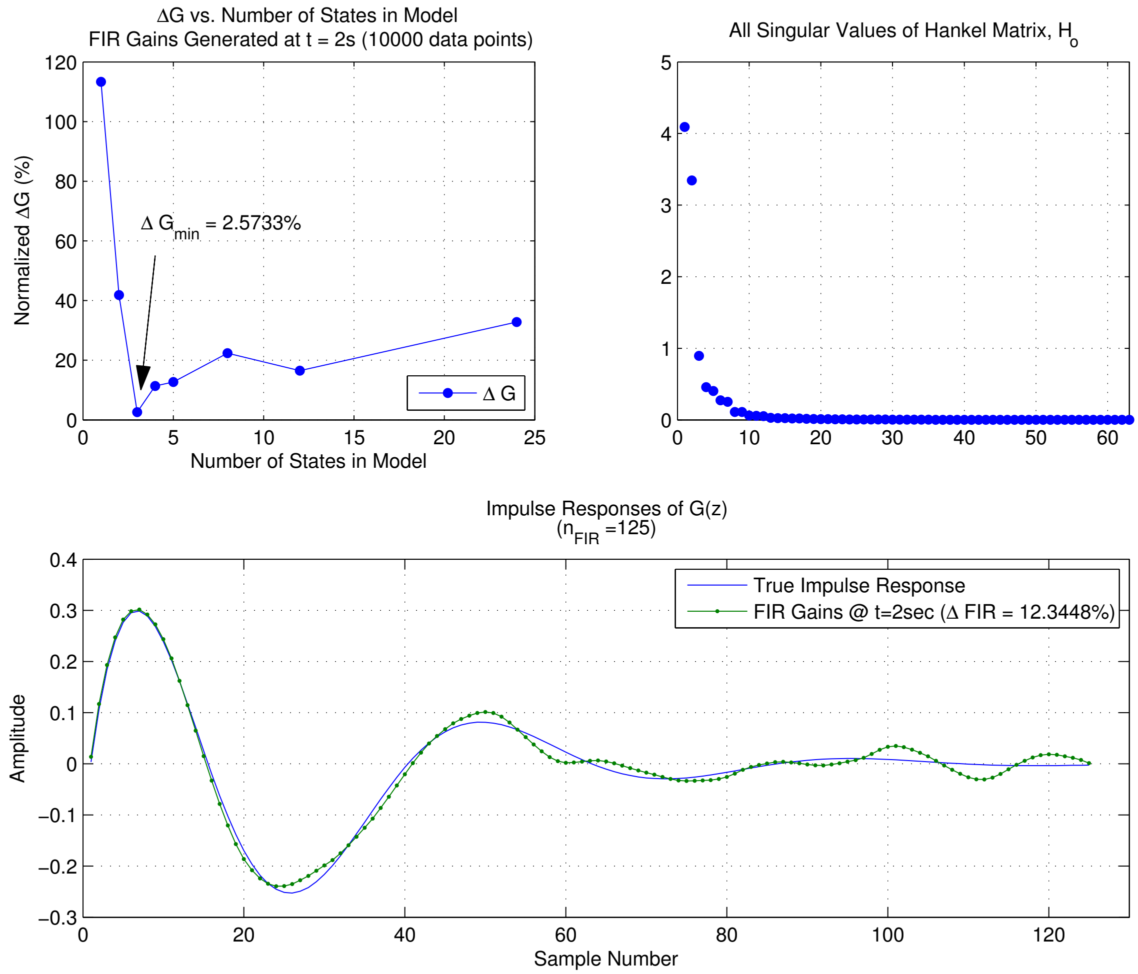

Figure 6.16: Results for when FIR gains are generated at $t = 2s$.

The truth model is a 3-state system, implying that the Hankel matrix should have 3 significant non-zero singular values and rest should essentially be zero. When looking at Fig. 6.16 (top right), there are 3 significant singular values, however, there are still several more that are small, but non-zero. The non-zero singular values likely correspond to portions of the broadband disturbance that have not been fully projected out of the measurement. This is a consequence of constructing the Hankel matrix with FIR gains that have not reached steady-state. Coincidently, however, the 3 significant singular values in Fig. 6.16 (top right), are in fact very close to the corresponding singular values when the Hankel matrix is constructed from the steady-state FIR gains at $t = 30s$ (Fig. 6.17

(top right)). The result is that, when the order of the plant model constructed from the Hankel matrix is exactly equal to 3, the plant model is quite accurate, again indicated by Fig. 6.16 (top left). The trouble is if the order of the plant is not known *exactly*, then the $\Delta\hat{G}$ becomes large very quickly. For example, in Fig. 6.16 (top left), if a 4th order plant model was generated from the Hankel matrix, the $\Delta\hat{G}$ is about 12%, which exceeds the empirical threshold for stability when actually used by the adaptive controller as discussed in section 6.5. Because the accuracy of the plant model hinges on knowing the exact order of the true plant, this method of identification using just 2 seconds of data is impractical.

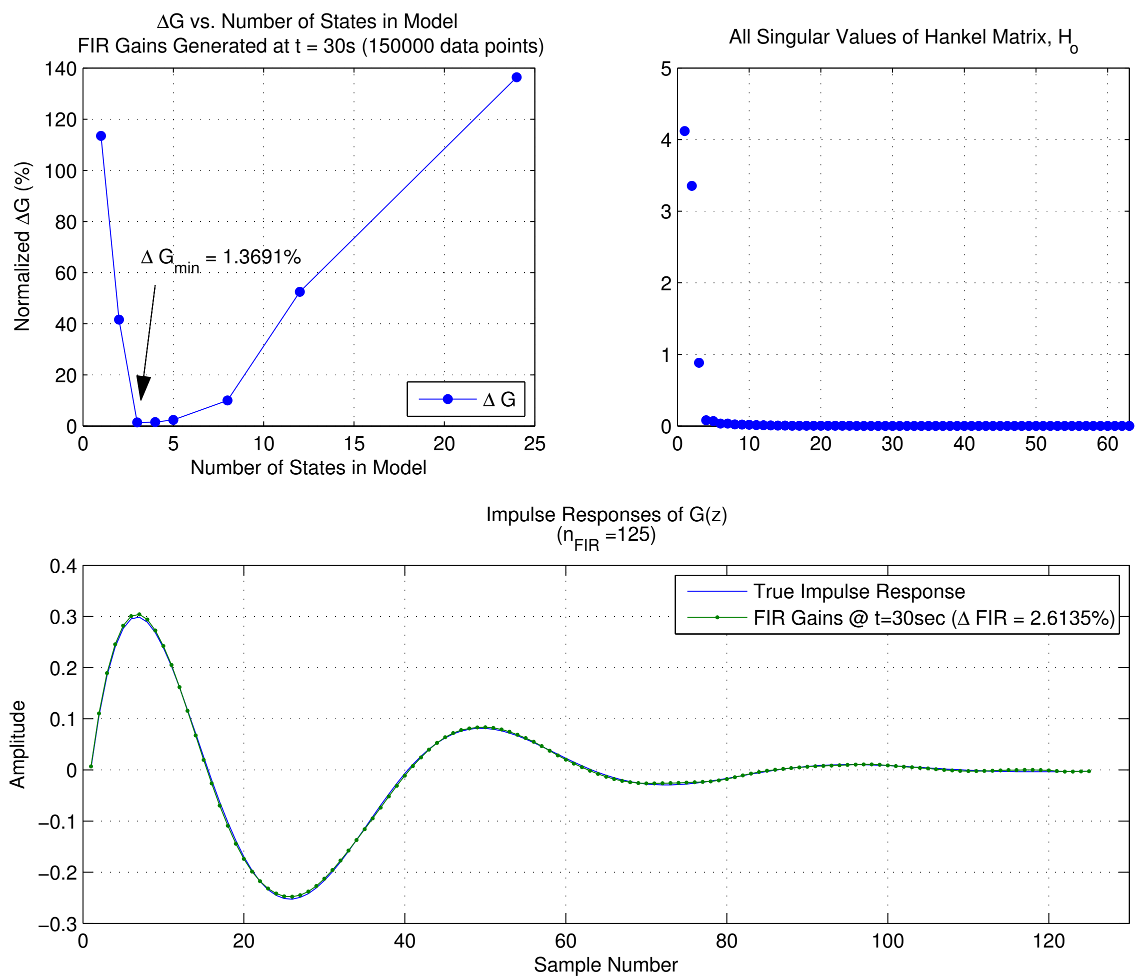

Figure 6.17: Results for when FIR gains are generated at $t = 30s$.

On the other hand, when the same method of identification is used when the FIR gains are allowed to reach steady-state, the result changes. This is the case in Fig. 6.17, in which the FIR gains are generated at $t = 30s$, and have reached steady-state. The bottom plot in Fig. 6.17 shows that the FIR gains are a much better approximation of the true system's impulse response ($\Delta_{\text{FIR}} = 2.61\%$). The result is that there are truly 3 significant Hankel singular values (Fig. 6.17, top right), and the rest are essentially zero. The 3 significant singular values correspond to the system states and, similar to when using 2 seconds of data (Fig. 6.16), if the order of the plant model generated is exactly equal to 3, the plant model is very good ($\Delta\hat{G} = 1.369\%$).

However, contrary to the case where only 2 seconds of data is used, Fig. 6.17 (top left) indicates that the plant model is still quite accurate when choosing a 4th order, or even a 5th order model. In fact, the $\Delta\hat{G}$ in this case doesn't exceed 10%, until a 8th order model is generated. The reason for this is that the plant model states corresponding to the next few, essentially zero, singular values contribute very little to the model, and so the model is not significantly affected. The caveat to this is that if too many states corresponding to erroneous, near-zero singular values are identified in the system, the accuracy of the overall plant suffers. This is consistent with Fig. 6.17 (top left).

To illustrate that there is a range of cases in between these two, Fig. 6.18 depicts a case where the FIR gains were generated at $t = 10s$. The gains have nearly converged, as indicated in the bottom plot of Fig. 6.18, but not completely. In this case, the additional Hankel singular values are closer to zero than in Fig. 6.16 (top right), but are still not as close to zero as in Fig. 6.17. The upper left plot in Fig. 6.18 shows that, similar to Fig. 6.17, there is a range of acceptable plant orders to choose from in order to keep the overall $\Delta\hat{G}$ below 10%. Since the FIR gains have not completely converged, this range is smaller (between 3rd order and 5th order), when compared to the case where the FIR gains have converged (between 3rd order and 8th order).

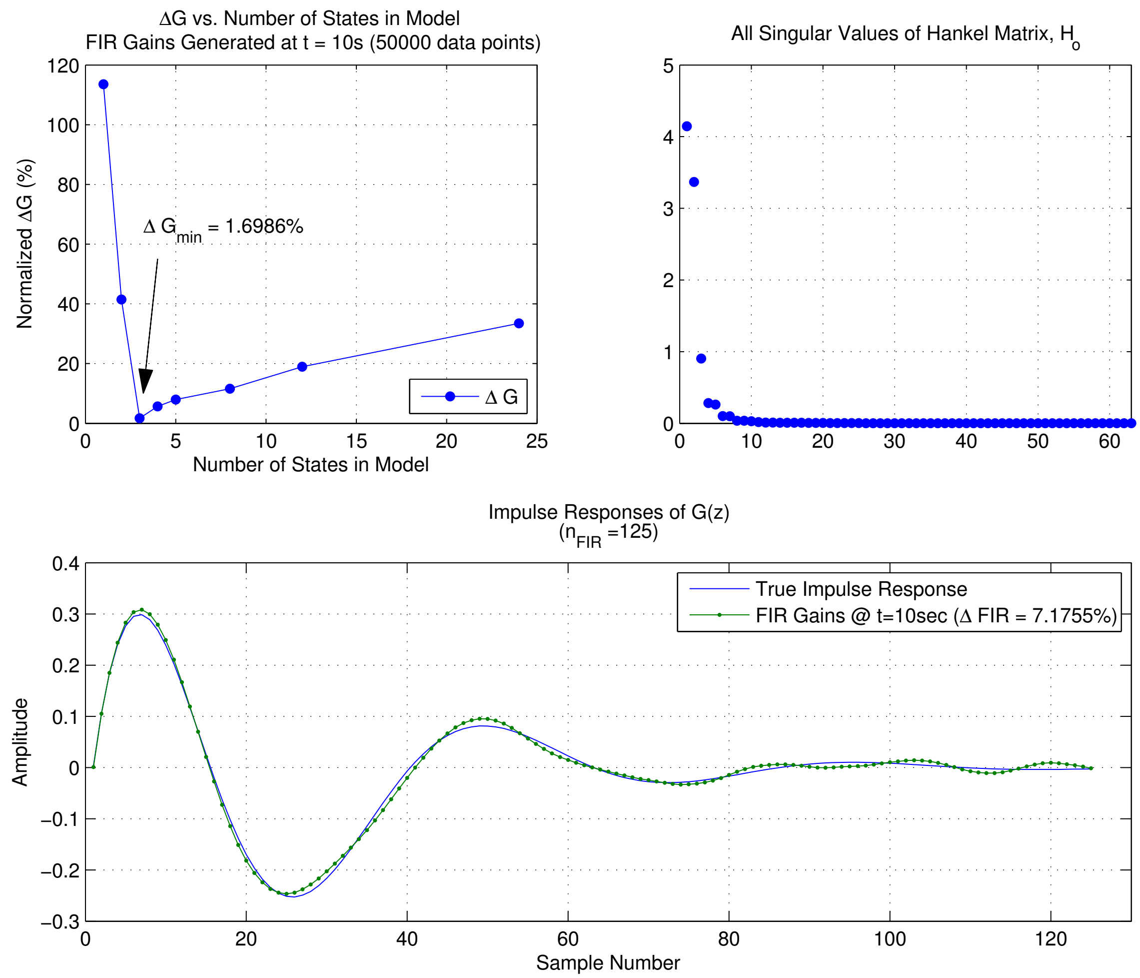

Figure 6.18: Results for when FIR gains are generated at $t = 10s$.

The conclusion to be drawn from this analysis is that, this system identification method of using the FIR gains of K_{fir} to build a Hankel matrix, generate Controllability and Observability matrices, and compute state-space matrices of the plant can be practical in some cases. If the FIR gains are allowed to reach steady-state, and there is confidence that those gains closely resemble the true impulse response of the system, *and* if the approximate order of the true plant is known beforehand (in the best case between a 3rd order and 8th order system), then this method is feasible in practice and can save additional computation required by generating the ARX coefficients in K_{arx}. If the FIR gains are not allowed to reach steady-state and computed much too early, the case in

Fig. 6.16 argues almost undeniably that this method is not feasible in practice because knowledge of the exact order of the plant is required a priori.

One modification to this ID method may, however, make the process feasible in practice. Rather than comparing a plant model to the true plant and generating a ΔG, which is not practical since the true plant G is unknown, the output error could be used instead. That is, the true output y or the estimated $\tilde{y}$, both of which can be stored in memory, could be compared to several simulated outputs $\hat{y}$ generated by passing the same input dither through the many $\hat{G}$'s of varying order. Then the $\hat{G}$ that results in the smallest Δy or $\Delta\tilde{y}$ would presumably be the most accurate plant model. Δy and $\Delta\tilde{y}$ are computed as:

$$\Delta y = \frac{||y_{sim} - y||_2}{||y||_2} \times 100 \tag{6.4}$$

$$\Delta\tilde{y} = \frac{||y_{sim} - \tilde{y}||_2}{||\tilde{y}||_2} \times 100 \tag{6.5}$$

where y_{sim} is the simulated output generated by passing the input dither through the particular plant model $\hat{G}$. It should be noted that this simulated output is due only to the dither sequence, since the disturbance sequence in unknown.

Figures 6.19 - 6.21 below each show the Δy and $\Delta\tilde{y}$ for the various plant models $\hat{G}$. The grey plots indicating ΔG as a function of states in the model is the exact same data as in Figs. 6.16 - 6.18. The green and orange plots indicate Δy and $\Delta\tilde{y}$ and are overlaid on the ΔG data. If the number of states in the model corresponding to the minimum value for either the Δy and $\Delta\tilde{y}$ plots coincide with the number of states in the model corresponding to the minimum ΔG value, then the correct number of states could be easily chosen. However, Figs. 6.19 - 6.21 all indicate that this is not the case. In every situation, the $\hat{G}$ that yields the smallest ΔG contains 3 states. However, the minimum value for each of the Δy and $\Delta\tilde{y}$ plots occurs when using the 5-state model.

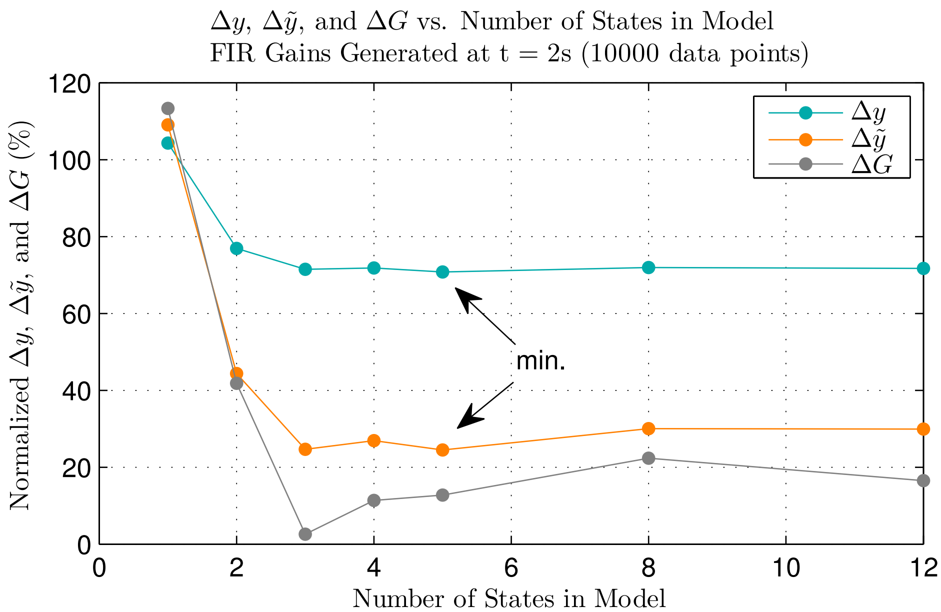

Figure 6.19: Δy and $\Delta\hat{y}$ for when FIR gains are generated at $t = 2s$.

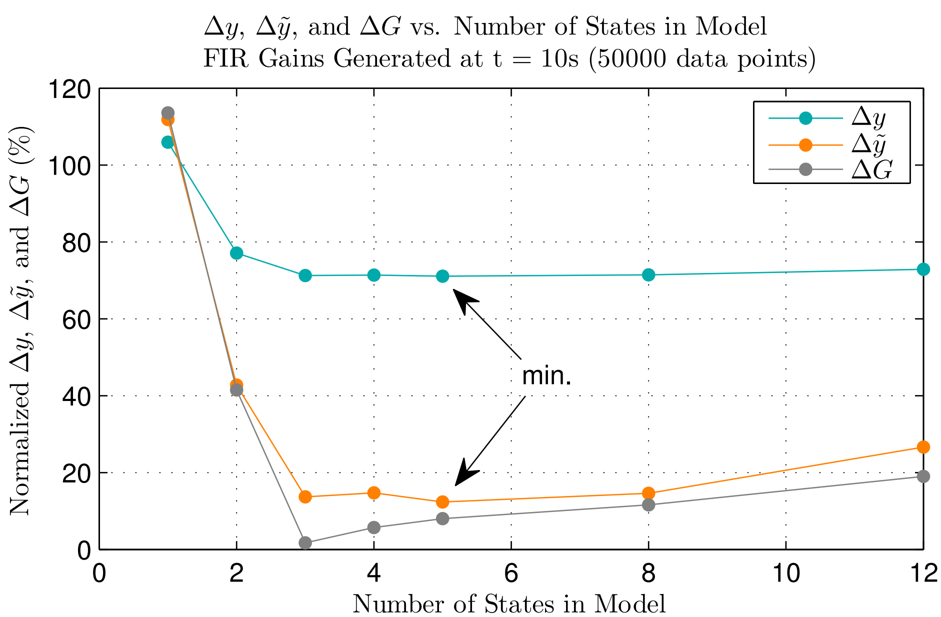

Figure 6.20: Δy and $\Delta\hat{y}$ for when FIR gains are generated at $t = 10s$.

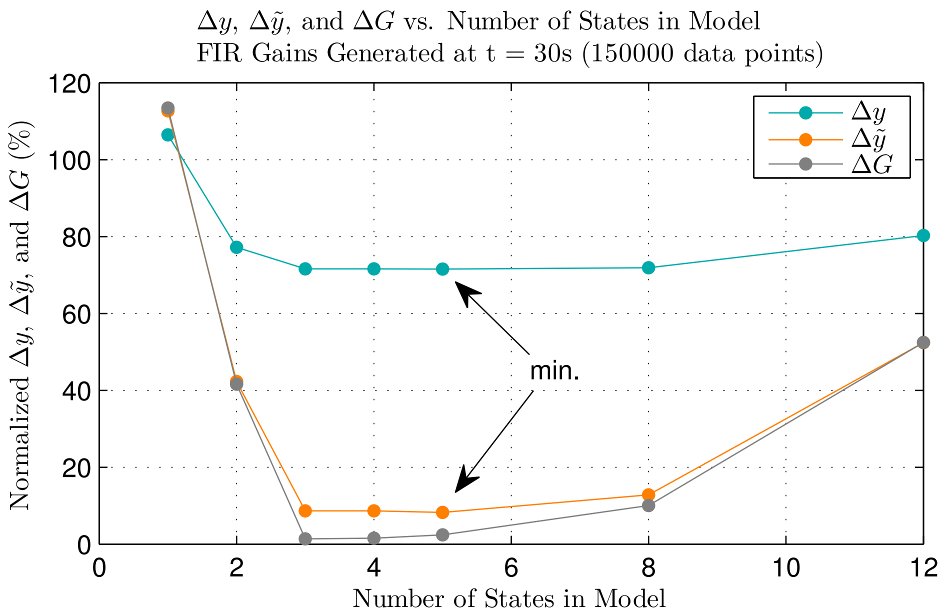

Figure 6.21: Δy and $\Delta\hat{y}$ for when FIR gains are generated at $t = 30s$.

The reason for the Δy discrepancy is because the total output y is used in the computation, whereas the the y_{sim} is the simulated output due only to the dither signal. When the simulated output contains the effects of both the dither and the disturbance, the Δy plot is nearly identical to the ΔG plot. However, noting again that the disturbance is considered to be unknown, this would be unrealistic in practice.

The discrepancy in the $\Delta\tilde{y}$ plot is because the dither and disturbance sequences are not fully uncorrelated. It was suspected that the discrepancy was due to the $\tilde{y}$ sequence being generated when the FIR filter was not in steady-state, however, this was not the case as the same discrepancy was present even when the $\Delta\tilde{y}$ values were computed using only a window of data well after steady-state conditions were reached. Thus, even in steady-state, the $\tilde{y}$ sequence contains traces of the disturbance, which negatively affects the $\Delta\tilde{y}$ values.

The result in Fig. 6.21 indicates that, even if the minimum value of the $\Delta\tilde{y}$ plot lead the user to select the 5-state plant, in this case, the 5-state plant might still be a good choice since its associated ΔG is far below the empirically determined 10% threshold for stability. Still, the results presented here indicate that, in general, the choice of the number of states in the model is rather difficult and that this modification to the method of bypassing K_{arx} may be impractical.

6.7 Summary and Case Study

The results in the prior sections have demonstrated that, for a practical application of this online ID method, the requirements are:

- The order of K_{fir} should be sufficiently large to capture the impulse response of the system. Over-parameterization of K_{fir} yields reduced accuracy of the plant model.
- The order of K_{arx} should be larger than the order of the plant, to account for correlation between dither and jitter. Over-parameterization of K_{arx} did not indicate significant detrimental effects.
- K_{fir} should be allowed to reach steady-state before the data is used by K_{arx}. In other words, K_{arx} should be reinitialized once K_{fir} reaches steady-state.
- Once K_{arx} is reinitialized, it requires substantially less time before accurate ARX coefficients are produced.

The following analysis applies the previous results to identifying a 4th-order plant representing the sum of two channels of a laser beam steering experiment. Each channel of the beam steering experiment (horizontal and vertical) was designed to operate in an open-loop configuration. Then on each channel a simple classical LTI controller was applied and designed such that the closed-loop system displayed a lightly-damped mode at a specified frequency. Once both systems were summed, the resulting 4th-order plant displayed 2 resonant peaks, and represented an overall more challenging plant to identify than in the previous sections.

The plant G in this study is:

$$G(z) = \frac{0.001542z^3 + 0.02829z^2 - 0.06025z + 0.03084}{z^4 - 3.937z^3 + 5.846z^2 - 3.879z + 0.9703} \tag{6.6}$$

and has the impulse response shown in Fig. 6.22. Recalling that the K_{fir} filter needs a filter order that corresponds to the number of samples required to capture the impulse response of the system, Fig. 6.22 indicates that the order of K_{fir} should be roughly 500-700.

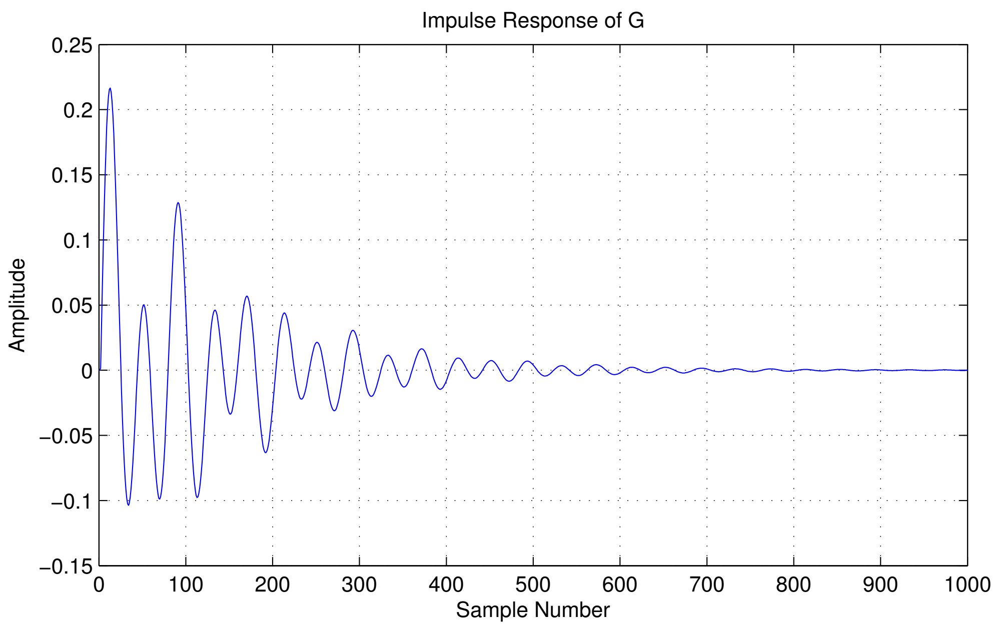

Figure 6.22: Impulse response of $G(z)$.

A more detailed analysis of the K_{fir} filter order is represented in Fig. 6.23, which suggests that the minimum ΔG is achieved once the order reaches about 400. This implies that capturing the entire impulse response is not necessary, and maximum performance can be achieved using the majority of the impulse response sequence.

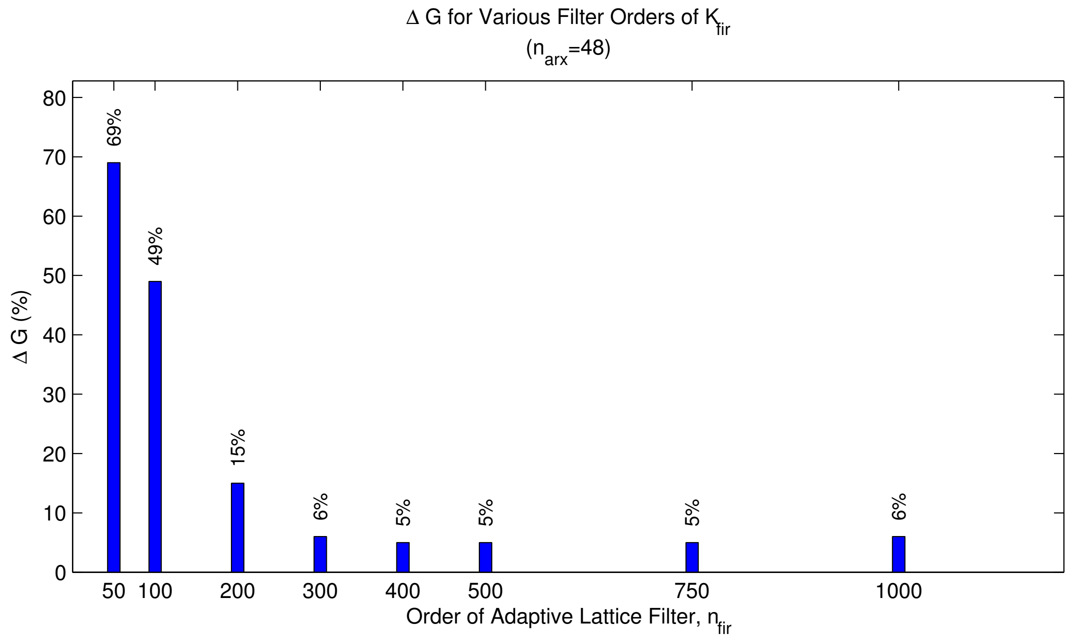

Figure 6.23: ΔG vs. n_{fir} for fixed value of n_{arx}.

Consistent with the study in the previous sections, the order of the K_{arx} filter required in this analysis is substantially higher than true order of the plant. Figure 6.24 indicates that the K_{arx} filter should have at least an order of 48 to realize maximum performance. Again, this is largely due to the inevitable correlation between the dither and the disturbance sequences. Because the contribution of $\tilde{y}$ sequence due to the disturbance is not completely zero, the K_{arx} filter attempts to identify those disturbance components and, given the bandwidth of the disturbance, can only do so with a sufficiently larger filter order.

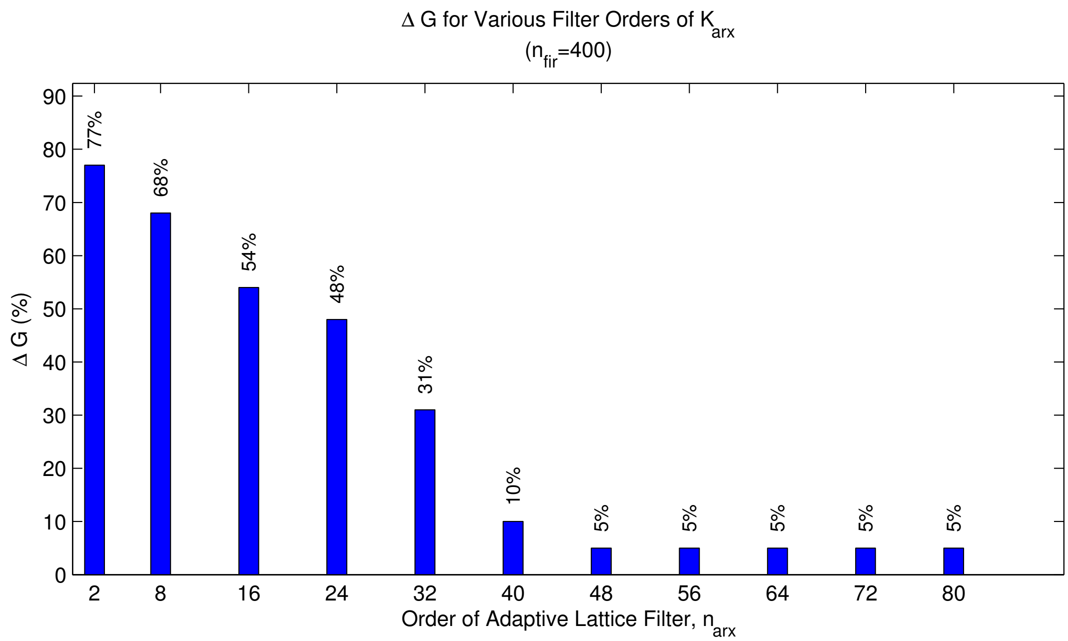

Figure 6.24: ΔG vs. n_{arx} for fixed value of n_{fir}.

Figure 6.25 shows the Bode plots for the true plant, as well as two versions of the ARX identification. In both of the identified plants, the FIR filter was allowed to reach steady-state, at which point the K_{arx} filter was reinitialized; this took 250,000 data points, or 50 seconds. Once the K_{arx} filter was reinitialized, ARX coefficients were generated using 1 additional second of data and 10 additional seconds of data. The plants generated by the ARX coefficients are denoted by $G_{arx,1}$ and $G_{arx,10}$, respectively. It should be noted that using more than 10 seconds of data before generating ARX coefficients did not significantly improve the accuracy of the plant model.

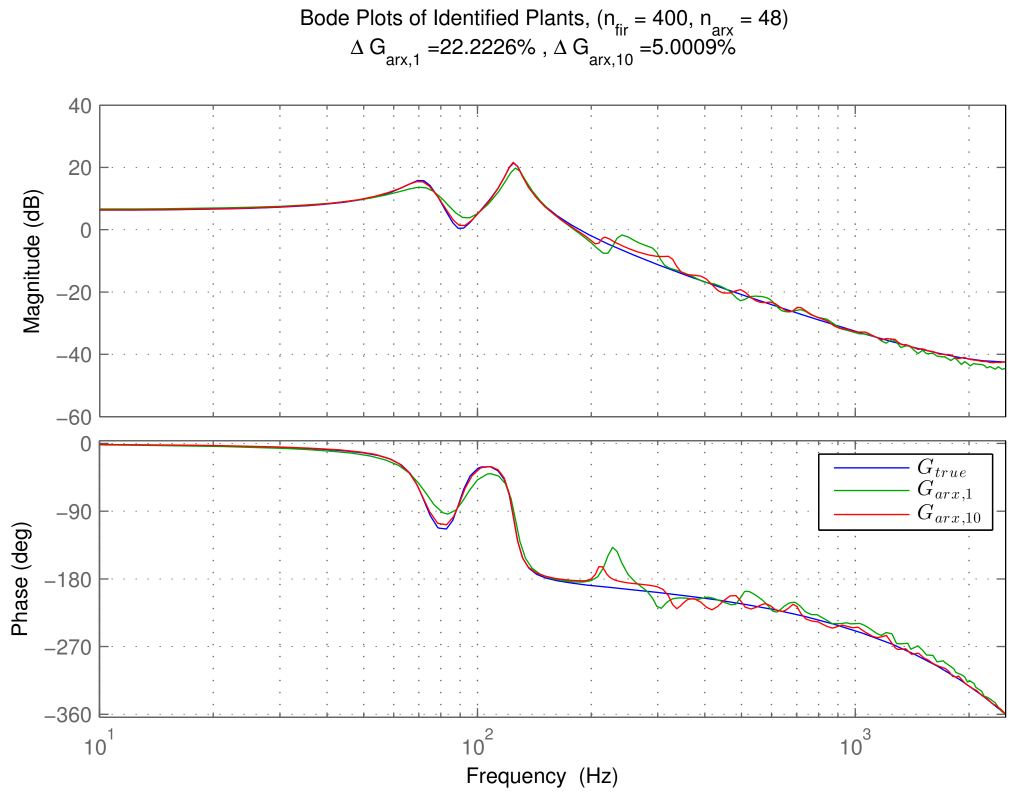

Figure 6.25: Bode plots of true plant and two identified ARX plant models.

It can also be seen from Fig. 6.25 that the primary dynamics (peaks at ≈70Hz and 110Hz) are better captured by $G_{arx,10}$, in both magnitude and phase. In both ARX models, however, there is some model error between 200Hz and 300Hz. It is more pronounced in the $G_{arx,1}$, model, yet it is still apparent in both ARX models. This is not surprising, as the PSD of the disturbance model shown in Fig. 6.26, indicates that nearly all of the power in the disturbance sequence resides in this frequency band. This is another indication that the dither and disturbance sequences are not fully uncorrelated.

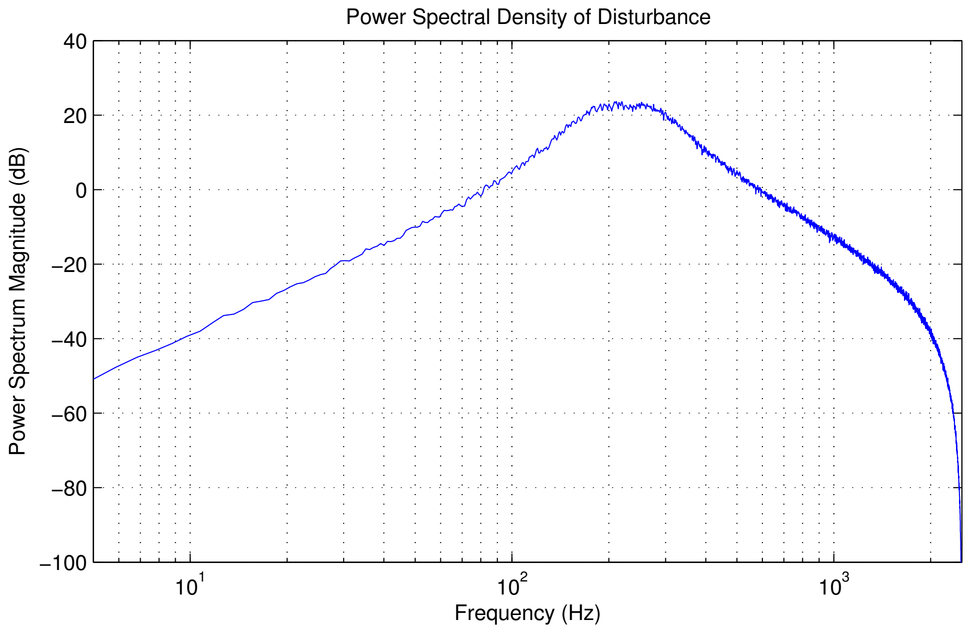

Figure 6.26: Power spectral density of disturbance sequence.

A balanced truncation of the high-order ARX model to a 4th-order plant model yields slightly better plant models, as indicated by Fig. 6.27, but the primary dynamics are captured equally well. Figure 6.28 shows the result from performing an offline batch identification with the n4sid algorithm in the MATLAB System Identification Toolbox [23]. This is included to show that the ARX model generated online is nearly as accurate as the state-of-the-art subspace identification which is much slower and must be performed offline.

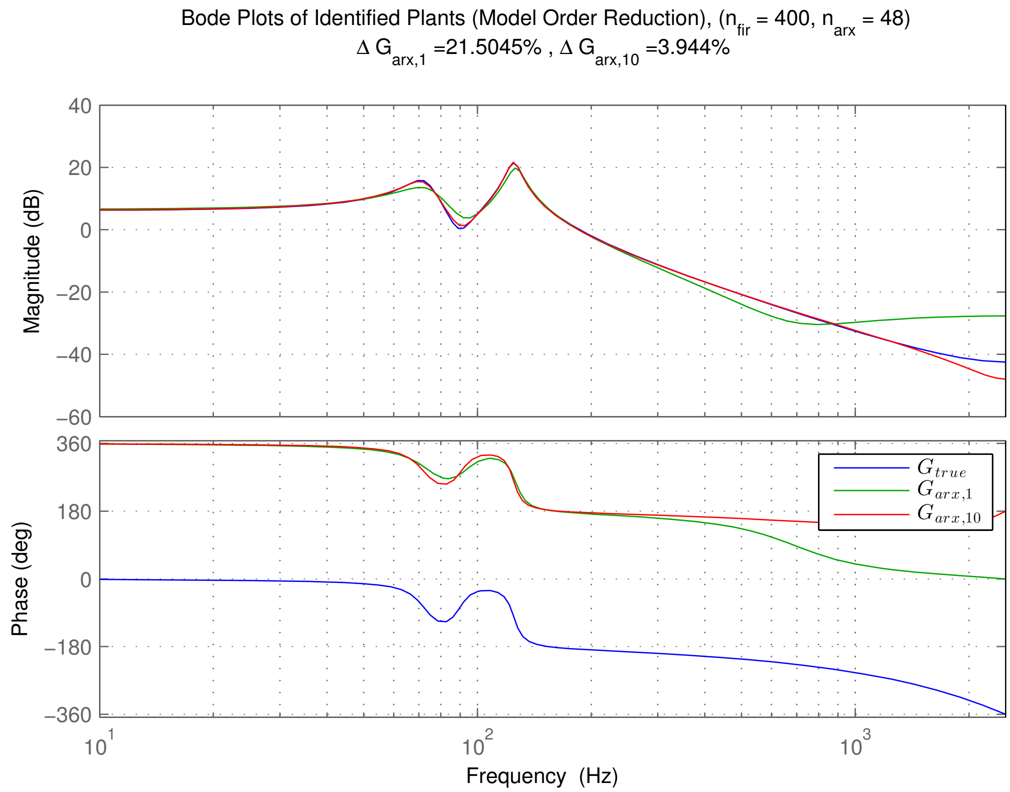

Figure 6.27: Bode plots of true plant and two identified ARX plant models with model order reduction (4th order plants).

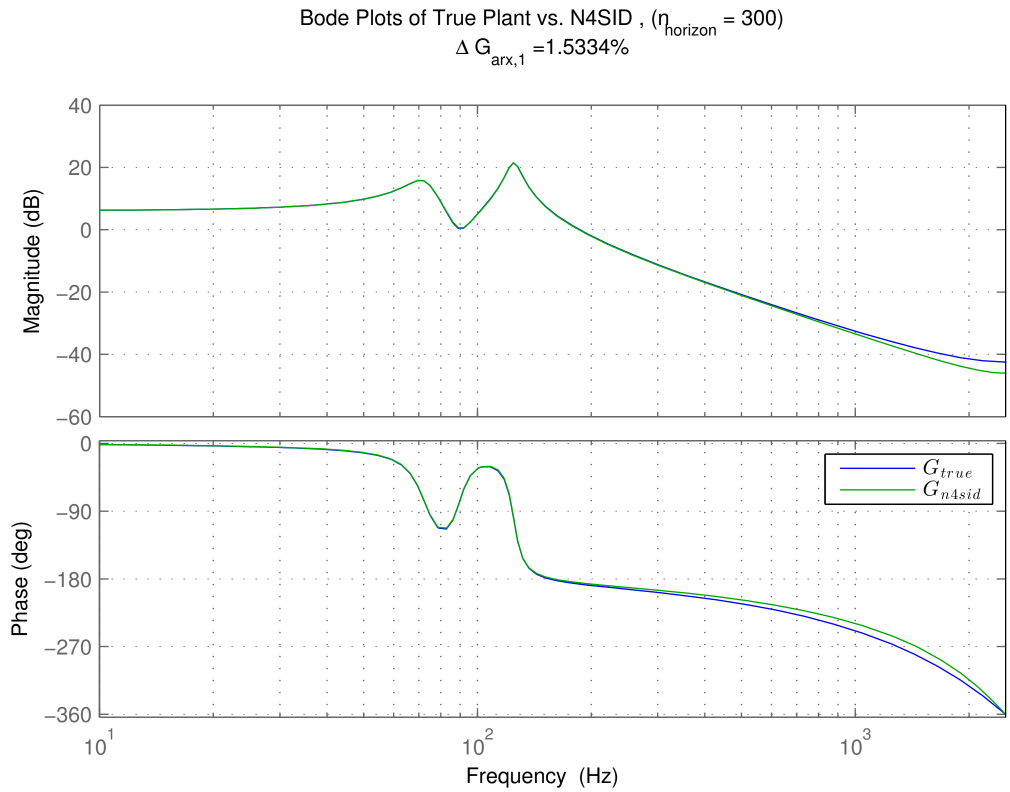

Figure 6.28: Bode plots of true plant and plant identified with N4SID

Finally, it should be noted that the amount of dither used in this study increased the rms of the output by approximately 10%. This serves as a reasonable increase in output error, and would be practical in most cases if applied in practice. As the previous study in section 6.5 indicates, model accuracy can be significantly increased if a larger relative amount of dither is used to drive the plant. However, considering the constraint that the rms output due to the dither should be as small as possible, a 10% increase in output rms for a 5% ΔG (Fig. 6.25) is a reasonable tradeoff.

CHAPTER 7

Concluding Remarks

This research has analyzed a novel receding-horizon adaptive controller designed to increase the disturbance rejection bandwidth far beyond that of any classical LTI controller. In addition to a larger disturbance rejection bandwidth, the receding-horizon performance index contains several tuning parameters not employed by the previous generation minimum-variance adaptive controller. Specifically, the ability of the receding-horizon performance index to penalize the adaptive control command yields a far more robust controller, capable of being applied in experimental settings where plant modeling error and nonlinearities are significant. Further, the structure of the receding-horizon controller allows for frequency weighting to be incorporated in a natural and intuitive manner, which is unique to this controller. This ability to tailor the bandwidth of the adaptive control command further contributes to the robustness of the controller.

The receding-horizon adaptive controller was successfully implemented on two experiments, each of which highlighted different controller attributes. In the laser beam steering experiment, two independent SISO controllers worked to reject a complex disturbance. The effects of the adaptive filter order as well as the control penalty parameters were studied in detail. In the magnetic bearing experiment, the controller's ability operate in a MIMO configuration was tested and the added challenge of a time-varying disturbance was introduced. The effect of the forgetting factor λ was studied in the context of broadband vs. time-varying disturbances. In a third experimental application, the receding-horizon adaptive controller was paired with a peak resonator in a parallel architecture. Independence of the controllers was demonstrated and the results concluded that the combination controller was able to reduce a complex disturbance to levels lower than either controller was capable of alone.

Because the adaptive controller requires a plant model for which it is controlling, a significant effort was devoted to the development and analysis of a novel online system identification method. This method would eliminate the need for a one-time offline system identification, which is generally regarded as impractical in application due to the perpetual nature of the disturbances that affect the particular system. Using a pair of adaptive filters, the identification algorithm first filters out the portion of the output due to a small input dither sequence. Then the second filter processes the input dither and the supposed output due only to that dither and generates the ARX coefficients of the plant. All of this is performed in real-time and in the presence of a complex disturbance, making the identification particularly challenging. Several parameters and alternative methods associated with this online identification algorithm were analyzed and simulated, and the results were summarized and successfully applied to a more complex, 4th-order plant.

References

[1] V. A. Skormin, M. A. Tascillo, and T. E. Busch, "Adaptive jitter rejection technique applicable to airborne laser communication systems," *Optical Engineering*, vol. 34, no. 15, pp. 1263–1268, May 1995.

[2] R. M. Glaese, E. H. Anderson, and P. C. Janzen, "Active suppression of acoustically induced jitter for the airborne laser," *Proc. SPIE 4034, Laser Weapons Technology*, pp. 151–3421, Apr. 2000.

[3] M. A. McEver, D. G. Cole, and R. L. Clark, "Adaptive feedback control of optical jitter using Q-parameterization," *Optical Engineering*, vol. 43, no. 4, pp. 904–910, Apr. 2004.

[4] N. O. Pérez Arancibia, N. Chen, S. Gibson, and T.-C. Tsao, "Variable-order adaptive control of a microelectromechanical steering mirror for suppression of laser beam jitter," *Optical Engineering*, vol. 45, no. 10, pp. 104 206–1–12, Oct. 2006.

[5] P. K. Orzechowski, N. Y. Chen, J. S. Gibson, and T.-C. Tsao, "Optimal suppression of laser beam jitter by high-order rls adaptive control," *IEEE Transactions on Control Systems Technology*, vol. 16, no. 2, pp. 255–267, Mar. 2008.

[6] N. O. Perez Arancibia, J. S. Gibson, and Tsu-Chin Tsao, "Frequency-weighted minimum-variance adaptive control of laser beam jitter," *IEEE/ASME Transactions on Mechatronics*, vol. 14, pp. 337–348, June 2009.

[7] Steve Gibson, T.-C. Tsao, and Dan Herrick et al., "Adaptive jitter control for tracker line of sight stabilization," in *Optics+Photonics*. San Diego, CA: SPIE, August 2010.

[8] T.-C. Tsao, S. Gibson, K.-C. Chiu, and S.-J. Chen, "Robustness enhanced adaptive control and optical head focusing experiment," in *IFAC World Congress*. Milan, Italy: IFAC, August 2011.

[9] J. Tesch and J. S. Gibson", "Receding-horizon adaptive control of aero-optical wavefronts," in *Conference Name*. City, State: Org, Month Year.

[10] N. Tsuchiya, J. S. Gibson, T.-C. Tsao, and M. Verhaegen, "Control of jitter in a laser beam experiment by receding-horizon adaptive control," in *IFAC Symposium on Mechatronic Systems*. Hangzhou, China: IFAC, April 2013.

[11] S. B. Jiang and J. S. Gibson, "An unwindowed multichannel lattice filter with orthogonal channels," *IEEE Transactions on Signal Processing*, vol. 43, no. 12, pp. 2831–2842, December 1995.

[12] S.-B. Jiang and J. S. Gibson, "An unwindowed multichannel lattice filter with orthogonal channels," *IEEE Transactions on Signal Processing*, vol. 43, no. 12, pp. 2831–2842, Dec. 1995.

[13] N. O. Pérez Arancibia, N. Chen, J. S. Gibson, and T.-C. Tsao, "Variable-order adaptive control of a MEMS steering mirror for suppression of laser beam jitter," *Optical Engineering*, vol. 45, no. 10, pp. 1–12, October 2006.

[14] I. Arredondo, J. Jugo, and V. Etxebarria, "Modeling of a flexible rotor maglev system." American Control Conference, June 2006, pp. 406–411.

[15] G. Yang, Y. Xu, Z. Shi, and H. Gu, "Characteristic Analysis of Rotor Dynamics and Experiments of Active Magnetic Bearing for HTR-10GT." *Nuclear Engineering and Design*, no. 237, pp. 1363–1371, 2007.

[16] F. Matsumura, M. Fujita, and C. Oida., "A Design of Robust Servo Controllers for an Unbalance Vibration in Magnetic Bearing Systems." *Magnetic Bearings*, pp. 319–326, Springer Berlin Heidelberg,1989.

[17] X. Zhang, T. Shinshi, L. Li, K. Choi, and A. Shimokohbe, "Precision control of radial magnetic bearing," in *Proceedings of the 10th International Conference on Precision Engineering*, 2001, pp. 714–718.

[18] Y. Wang, K. Chu, and T. C. Tsao, "An analysis and synthesis of internal model principle type controllers," in *American Control Conference, 2009. ACC '09.*, June 2009, pp. 488–493.

[19] C. Kang and T. C. Tsao, "Control of magnetic bearings with plug-in time-varying harmonic resonators," in *American Control Conference, 2014. ACC '14.*, June 2014, pp. 4237–4242.

[20] B. Francis and W. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.

[21] M. Tomizuka, "Zero phase error tracking algorithm for digital control," *Trans. of ASME, Journal of Dynamic Systems, Measurement, and Control*, vol. 109, pp. 65–68, 1987.

[22] J. A. Butterworth, L. Y. Pao, and D. Y. Abramovitch, "Analysis and comparison of three discrete-time feedforward model-inverse control techniques for nonminimum-phase systems," *Mechatronics*, vol. 22, no. 5, pp. 577–587, August 2012.

[23] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Englewood Cliffs, New Jersey: Prentice Hall, 1999.