The Whirling Blade and the Steaming Cauldron

Thesis by

Robert Edward Bodenheimer, Jr.

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy

California Institute of Technology
Pasadena, California
1995
Submitted June 2, 1995

Acknowledgements

The moral is, it hardly need be shown,
All those who try to go it sole alone,
Too proud to be beholden for relief,
Are absolutely sure to come to grief.
—Robert Frost

With high spirits I undertook to thank the many people who have helped me during my graduate tenure at Caltech. I have had the opportunity and pleasure of meeting many very talented people during this period, but unfortunately I am unable to thank them all. I hope those not mentioned here will not take umbrage — be sure that I have not forgotten.

John Doyle was always quick to point out the "big picture" and helped me not to lose sight of the forest for the trees. His profound understanding of control theory coupled with his research advice has been a constant and reassuring source of support throughout. John is an interesting and unique advisor.

A good deal of my control education was due to Andy Packard. He's always been willing to explain things to me in as much detail as I required. These discussions often occurred over midnight hot dogs with hot peppers when he was here at Caltech. Much of the work in this dissertation was possible because he spent the time to make the theory clear to me. I especially appreciate his effort to be on my committee.

Richard Murray kept me on the right track with his constant encouragement. His keen perception and enthusiasm have been an invaluable asset. While this small encomium is wholly inadequate to express my gratitude, I trust he's aware of it.

I'm beholden to Drs. Pietro Perona and Robert McEliece for serving on my dissertation committee. I thank them both very much.

Special thanks are due to my fellow students at Caltech. Matt Newlin has been an unwavering beacon of insight and enlightenment. Carolyn Beck worked very hard with me and our collaboration has been a true pleasure. John Morris and I have spent hours discussing control and helping one another throughout our graduate careers. Both Carolyn and John's thoughtful criticisms on this manuscript aided me greatly. Mike Kantner has always been eager to help me with the ducted fan, both now and when I was first learning to use it. It pleases me more than I can say to count these people as friends as well as colleagues.

The control group at Caltech is a very open and cooperative one, and I have benefitted tremendously from this atmosphere. The enlightening conversations I've had with Jorge Tierno, Raff D'Andrea, Wei-Min Lu, and Fernando Paganini have certainly clarified my understanding of control and systems theory and I hope that it's reflected in this work. Jorge and Fernando, in particular, have a penchant for rigor and precision which is always helpful and very refreshing.

Heather Sherman gave of her time and energy in a very careful reading of a draft of this dissertation. She also told me what the cow was doing in the powder room during the opera, and which opera it was.

Finally, my heartfelt gratitude goes to Pascale Bendotti. Her drive and determination have been a fountain of inspiration. Most of the results in this work are based on our long collaboration, and I am deeply indebted to her. *Merci, mille fois.*

Abstract

This dissertation applies recent theoretical developments in control to two practical examples. The first example is control of the primary circuit of a pressurized water nuclear reactor. This is an interesting example because the plant is complex and its dynamics vary greatly over the operating range of interest. The second example is a thrust-vectored ducted fan engine, a nonlinear flight control experiment at Caltech.

The main part of this dissertation is the application of linear parameter-dependent control techniques to the examples. The synthesis technique is based on the solution of linear matrix inequalities (LMIs) and produces a controller which achieves specified performance against the worst-case time variation of measurable parameters entering the plant in a linear fractional manner. Thus the plant can have widely varying dynamics over the operating range, a quality possessed by both examples. The controllers designed with these methods perform extremely well and are compared to $\mathcal{H}_\infty$, gain-scheduled, and nonlinear controllers.

Additionally, an in-depth examination of the model of the ducted fan is performed, including system identification. From this work, we proceed to apply various techniques to examine what they can tell us in the context of a practical example. The primary technique is LMI-based model validation.

The contribution this dissertation makes is to show that parameter-dependent control techniques can be applied with great effectiveness to practical applications. Moreover, the trade-off between modelling and controller performance is examined in some detail. Finally, we demonstrate the applicability of recent model validation techniques in practice, and discuss stabilizability issues.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Does the road wind up-hill all the way?
    Yes, to the very end.
Will the day's journey take the whole long day?
    From more to night, my friend
—Christina Rosetti

As technology increases, control systems are being used on more complex systems than ever before. Simultaneously, the control community has worked to develop new theoretical machinery to apply to an increasing range of systems. Unfortunately, in the last twenty years, there has been extremely little transition from modern theoretical developments to their widespread use in industrial applications. In fact, almost all industrial controllers are either switching or proportional-integral-derivative (PID). Industry has shown itself eager to adopt new methodologies in control, but only when the methodologies are made accessible to them. The few, isolated examples where this transition occurs show that modern theoretical developments find use in industry only when they are presented as a set of reliable and accessible software tools. Moreover, the theory must be applied at least to simple applications for industry to accept them.

The intent of this dissertation is to bridge the gap that exists between theory and practice by evaluating several recent techniques for use in control system design. The evaluations are done by testing the techniques' usefulness either on a real application, or a realistic simulator. Our goals are to assess how well the techniques apply (are they straightforward to apply as the theory is presented?), how good the techniques are (do they provide a better control law or provide an insight not easily obtainable otherwise?), and how computationally reliable the techniques are (can we compute what the theory provides for us?). By addressing these issues, we will provide useful information of two sorts: first, experimentalists can compare our application of these techniques to their own problems to determine if they are useful; second, theoreticians can examine the limitations revealed by these techniques as a point to begin their development of further theory.

The specific techniques we explore are a new method for gain-scheduling and a recent technique aiding in the evaluation of robust control models. In all cases, this work represents the first application of these techniques to a physical system. Our primary emphasis is on the gain-scheduling technique.

All of the computational procedures discussed in this work involve solving linear matrix inequalities (LMIs). The solution of an LMI is a convex optimization problem and the solution of such problems is much studied and well understood [15, 36]. All LMIs herein were solved via a software package developed by Gahinet, Nemirovskii, Laub, and Chilali [22], which employs what is probably the most efficient of the known algorithms for solving LMIs — Nemirovskii's projective algorithm. The control community has almost reached the state where it regards any problem that has been reduced to an LMI as solved, but we shall caution against this. As we shall see in Chapters 4 and 6, there is still work needed on numerical LMI solvers.

1.1 Gain Scheduling

Gain scheduling is a venerable technique in control systems, since it has long been recognized that linear control techniques are inadequate to handle many systems of interest, and nonlinear techniques have not yet advanced to the stage where they can adequately control these systems. The idea behind gain scheduling is that one can obtain linear controllers for a variety of different operating points of the system, and then mix them together in an *ad hoc* fashion so that a control law is fashioned over the entire operating regime of the plant. The plant and controller are then exhaustively simulated to determine if the non-local performance of the system is adequate. The disadvantage of this technique is that there are no known methods of scheduling such controllers which provide an *a priori* guarantee on the stability or performance of the closed-loop system. Additionally, large and often unacceptable transients can occur when switching between these controllers. To avoid this transient problem (the issue of "bumpless transfer") the gains of the controllers are usually interpolated. As a result, it is more effort to schedule a set of LQG or $\mathcal{H}_\infty$ controllers, since the controller must be put in an observer-based form for this to succeed. This problem has been investigated for $\mathcal{H}_\infty$ controllers in [26, 38], among others.

Shamma [45, 46] has studied gain scheduling by examining systems called linear parameter-varying (LPV) systems. These are systems of the form

$$\begin{aligned} x(k+1) &= A(\theta(k))\,x(k) + B(\theta(k))\,u(k) \\ y(k) &= C(\theta(k))\,x(k) + D(\theta(k))\,u(k) \end{aligned}$$

where $\theta(k)$ represents a time-varying vector of parameters upon which the entries of the state space matrices depend. For a general LPV system, the values of $\theta(k)$ are not known *a priori*, but are assumed to be measurable. If the time-variation of $\theta(k)$ is known for all k, the system becomes a linear time-varying (LTV) system, while if $\theta(k) = \theta_0$ is constant for all k, the resulting system is linear time-invariant (LTI) system.

Only sufficient conditions for the analysis or synthesis for this type of system are known. Much work has been done on a generalization of LPV systems where θ is no longer a time-varying vector but is rather an operator from, for example, ℓ_2 to ℓ_2. Clearly, conditions derived in the operator framework are sufficient for the LPV framework, but are not necessary. Unfortunately, the control literature is vague about the fact that these are only sufficient conditions for the real systems of interest.

Both Shamma [44] and Megretskii [31] have derived necessary and sufficient conditions for an LPV system in the operator framework to be stable. Lu [30] later developed a theory of stabilization in this framework when the LPV systems can be expressed as a linear fractional transformation (LFT) of a constant matrix and an operator set. That is, the system has a transfer function equal to

$$D + C\Delta\left(I - A\Delta\right)^{-1} B$$

where Δ is a member of a prescribed set of operators (as shown in Chapter 2, a normal LTI system results when Δ is chosen to be the shift operator z^{-1}). He was able to derive necessary and sufficient conditions for the stabilizability and detectability of such systems.

Building on these results, a major breakthrough came in [39] from Packard, who developed machinery which showed that in the case of those LPV systems which can be represented as an LFT the synthesis problem is a convex optimization problem. Packard's theory produces an optimal parameter-dependent controller; i.e., the controller provides optimal performance against the worst-case operator in the prescribed set. A similar result in this area was obtained by Apkarian and Gahinet [1].

The advance this represents is that once a parameter-varying model is obtained, the synthesis is a one step procedure, and provides theoretical performance and stability guarantees. Also, the entire problem of bumpless transfer is avoided implicitly. The disadvantage of this procedure is that it is potentially quite conservative, since the worst-case operator may have little relation to a parameter variation that can physically occur. Nonetheless, these results do represent an upper bound on achievable performance, and assessing how well this upper bound works on physical examples is the goal of this work.

1.2 Model Validation

A robust control model is a model for a system which contains not only a model of the system dynamics, but an uncertainty and noise description as well. The models are used extensively in the $\mathcal{H}_\infty/\mu$ design framework [40]. The "model validation" problem was originally formulated by Smith and Doyle [48] to provide a connection between a robust control model and data measured from a physical system, and in this work we will use the term in the sense that they define it. It seeks to answer the question "Does the robust control model account for the measurements from the physical system?" A computational framework to apply the model validation problem was developed by Newlin and Smith [37]. Although the way we will apply the model validation problem in Chapter 6 is not exactly in the context in which it was formulated, it is a straightforward corollary to the standard framework. It represents the first application of these types of model validation techniques to a physical system.

The problem of approximating a system by one of lower order is another intensely studied problem in control (see [7] and the references therein). For general LPV systems in an LFT framework systematic techniques are quite recent. The techniques are essentially corollaries of the gain scheduling machinery developed by Packard and mentioned previously [39]. The work presented here represents the first attempt at an experimental application of these results.

1.3 Overview of the Dissertation

In Chapter 2 the theory needed to understand the LPV synthesis procedure will be detailed. We will also review LFTs and stability, stabilizability, and detectability for LPV systems. In addition, enough background is provided to show how the parameter-varying model of a general system can be constructed.

Chapter 3 introduces the first application — the primary circuit of a nuclear power plant — and details the LPV design procedure for this application. This application was selected because the dynamics of it vary enough that a linear controller is inadequate to provide reasonable performance over the operating range. An interesting point this application possesses is that the construction of the parameter-varying model must be done from identified, rather than first principles, models. We evaluate the resulting LPV controller in terms of its performance on a realistic nonlinear simulator.

Chapters 4 through 6 are devoted to the second application, a thrust-vectored ducted fan engine. This application is a nonlinear flight control experiment built at Caltech. In Chapter 4 the application and its model are described. Both an $\mathcal{H}_\infty$ and LPV controller are designed for it. The construction of the parameter-varying model is from a first principles model. Chapter 5 describes the trajectories the ducted fan is tested on, describes the performance measures used to compare the performance of different controllers, and presents an evaluation of the controllers based on these measures. In Chapter 6 we explore the use of model validation when applied to the ducted fan, and additionally discuss stabilizability issues of the model. As a precursor to this, a model of the ducted fan is identified from input-output data.

Finally, we conclude with an evaluation of these techniques, and the issues we feel merit more investigation. Our intent is to present the two applications independently. Thus some repetition of material is unavoidable, although we attempt to hold it to a minimum.

Chapter 2

Review of Linear Parameter-Varying Synthesis

Er that I ferther in the tale pace,
Me thinketh it acordaunt to resoun
To tell yow al the condicioun
Of each of hem, so as it semed to me
—Geoffrey Chaucer

In this chapter much of the theoretical machinery applied in this work is reviewed. We assume the reader has some familiarity with the fundamentals of feedback system analysis. In particular, we assume knowledge of the small-gain theorem and of memoryless operators. Willems [49] is an excellent reference for such material. The basic concepts of μ-analysis and μ-synthesis are also assumed; they can be found in Packard and Doyle [40]. Complete and rigorous explanations of this work can be found in [19, 40, 39].

The notation is standard. $\mathbb{R}$ and $\mathbb{C}$ denote the fields of real and complex numbers, respectively. $\mathbb{R}^k$ and $\mathbb{C}^k$ denote the real and complex k-dimensional vector spaces. $\mathbb{R}^{n\times m}$ and $\mathbb{C}^{n\times m}$ are the rings of real and complex $n \times m$ matrices. If $M \in \mathbb{C}^{n\times m}$, the maximum singular value of M is denoted by $\overline{\sigma}(M)$; both M^* and M' denote the complex conjugate transpose. The Hilbert space of square summable sequences is denoted by ℓ_2; $\mathcal{L}(\ell_2)$ represents the set of all linear time-varying operators on ℓ_2. The shift operator on ℓ_2 is denoted by z^{-1}.

2.1 Linear Fractional Transformations

The background machinery for almost all results in this work is that of linear fractional transformations (LFTs). These were first introduced into the control community by Redheffer [42], but did not gain wide acceptance until the work of Safonov [43] and Doyle [18]. With LFTs, we can easily describe sets of systems as an operation between an operator and a matrix. All linear interconnections of systems and LFTs and more importantly, any rational function, can be represented as an LFT [50, 15].

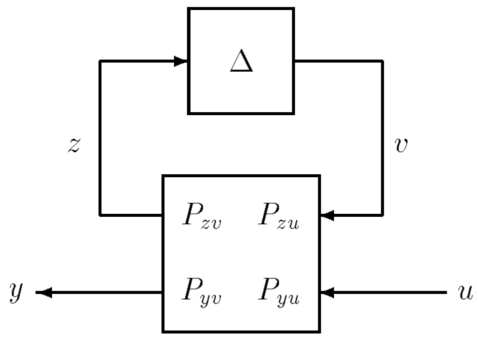

Figure 2.1: Block diagram representing a general LFT.

Our notation for LFTs will be as follows. Let $\mathbf{\Delta}$ denote the set

$$\left\{\text{diag}\left[\delta_1 I_{n_1}, \ldots, \delta_s I_{n_s}, \Delta_{s+1}, \ldots, \Delta_{s+F}\right] : \delta_i \in \mathcal{L}(\ell_2),\ \Delta_{s_j} \in \mathcal{L}(\ell_2)^{m_j \times k_j}\right\}.$$

For each $\Delta \in \mathbf{\Delta}$, consider the loop equations

$$\begin{aligned} z &= P_{zv} v + P_{zu} u \\ y &= P_{yv} v + P_{yu} u \\ v &= \Delta z. \end{aligned}$$

This set of equations is shown pictorially as a block diagram in Figure 2.1, and we let

$$P = \begin{bmatrix} P_{zv} & P_{zu} \\ P_{yv} & P_{yu} \end{bmatrix}.$$

Usually P will be considered as a matrix. It has interpretations as a system and an operator too, though, and we will move between these interpretations without making an explicit mention of the transition. It should be clear from the context what interpretation P has. For a general matrix P with many elements, we may explicitly denote a partition of P instead of writing what P_{zv}, P_{zu}, P_{yv}, and P_{yu} are; this partition will be denoted by lines, for example,

$$P = \left[\begin{array}{cc|c} A & B & C \\ \hline D & E & F \end{array}\right].$$

Here $P_{zv} = [A\ B]$, $P_{zu} = C$, $P_{yv} = [D\ E]$, and $P_{yu} = F$.

We can eliminate z and v from the loop equations, solving for y in terms of P, Δ, and u, which gives us

$$y = \left(P_{yu} + P_{yv}\Delta\left(I - P_{zv}\Delta\right)^{-1} P_{zu}\right) u \tag{2.1}$$

subject to the condition, which we shall henceforth assume, that $I - P_{zv}\Delta$ is invertible. The LFT is said to be *well-posed* when this is the case. Equation 2.1 will be denoted as

$$y = \left(\Delta \star P\right) u. \tag{2.2}$$

Another common notation is $y = F_u(\Delta, P)u$.

LFT examples

We will now present some very simple examples for constructing LFTs, which form the basis of the parameter-varying models of Chapters 3 and 4. As a first example, consider $y = (a\delta + b)u$. For this example $\Delta = \delta \in \mathbb{C}$ and

$$P = \left[\begin{array}{c|c} 0 & 1 \\ \hline a & b \end{array}\right].$$

Suppose we have a two-dimensional y

$$y = \begin{bmatrix} \alpha(\delta) \\ \beta(\delta) \end{bmatrix} u$$

where

$$\begin{aligned} \alpha(\delta) &= a\delta + c \\ \beta(\delta) &= b\delta + d. \end{aligned}$$

Then a representation for P is

$$P = \left[\begin{array}{cc|c} 0 & 0 & 1 \\ 0 & 0 & 1 \\ \hline a & 0 & c \\ 0 & b & d \end{array}\right].$$

This representation of P is not "minimal" in the sense of having the smallest size of Δ. A minimal representation is

$$P = \left[\begin{array}{c|c} 0 & 1 \\ \hline a & c \\ b & d \end{array}\right].$$

Suppose $y = (a\delta^2 + b\delta + c)u$. Then $\Delta = \delta I_2$ and a minimal representation for P is

$$P = \left[\begin{array}{cc|c} 0 & 0 & 1 \\ 1 & 0 & 0 \\ \hline b & a & c \end{array}\right].$$

Finally, consider the transfer function $G(z) = D + C(zI_n - A)^{-1}B$, for a given state space representation A, B, C, and D with n states, p inputs, and m outputs. Then $y = G(z)u$, and it can be represented as an LFT with $\Delta = z^{-1}I_n$ and P given by

$$P = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right].$$

Notice that the poles of G are precisely the points where $I_n - Az^{-1}$ is not invertible.

2.2 Stability, Stabilizability, and Detectability

A detailed discussion of the results in this section may be found in [30]. For a given set $\mathbf{\Delta}$, we want to address the problem of when the interconnection $\Delta \star P$ is stable. See [41] and the references therein for more detail. Let

$$\mathbb{B}_\Delta = \{\Delta \in \mathbf{\Delta} : \|\Delta\| \leq 1\}.$$

$\Delta \star P$ is stable when $(I - P_{zv}\Delta)$ is causally invertible as an operator on ℓ_2 for every $\Delta \in \mathbb{B}_\Delta$.

Assume now that one δ_i is the shift operator, and denote by $\mathcal{T}$ the commuting matrix sub-algebra

$$\mathcal{T} = \left\{T \in \mathbb{C}^{n\times n} : \det(T) \neq 0,\, T\Delta = \Delta T,\, \forall \Delta \in \mathbf{\Delta}\right\}. \tag{2.3}$$

For a given system

$$P = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right],$$

$\Delta \star P$ is stable if and only if there exists $X > 0$, $X \in \mathcal{T}$, such that

$$AXA^* - X < 0.$$

We say P is stabilizable if there exists $F \in \mathbb{C}^{n\times p}$ such that $(A + BF)$ is stable; P is detectable if there exists $L \in \mathbb{C}^{m\times n}$ such that $(A + LC)$ is stable. These conditions are met, respectively, if and only if there exist $X > 0$, $Y > 0$, with $X, Y \in \mathcal{T}$, such that

$$AXA^* - X - BB^* \;<\; 0 \tag{2.4}$$

$$A^*YA - Y - C^*C \;<\; 0. \tag{2.5}$$

2.3 $\mathcal{H}_\infty$ Synthesis

In this section we assume all systems are LTI. P will refer to the *generalized* plant, that is, what is normally called the plant, plus any weighting functions. Consider the standard feedback system shown in Figure 2.2. The vector signal d of exogenous inputs contains all disturbances, noises, and commands; e is the vector signal of quantities we wish to minimize; u and y are the controls and measurements, respectively.

Roughly speaking, our goal is to find a controller K which minimizes the transfer function from d to e, denoted $P \star K$, in the sense of making the maximal energy captured by $P \star K$ small. This problem was elegantly solved by Doyle *et al.* in the classic paper [19].

The actual synthesis procedure is sub-optimal in the sense that a controller K is found such that

$$\|P \star K\|_\infty < \gamma \tag{2.6}$$

for some pre-specified γ.

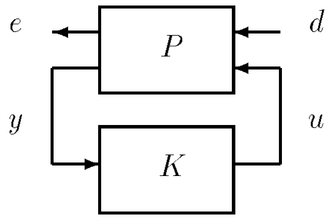

Figure 2.2: Feedback interconnection of P and K.

There exists some optimal γ_o such that for all $\gamma < \gamma_o$ no stabilizing controller can be found for which Equation 2.6 holds. Optimal $\mathcal{H}_\infty$ controllers do not have many desirable qualities [25] and the standard practice is to approximate the optimal controller with a sub-optimal one for some desired tolerance. This procedure of minimizing the value of γ to a prescribed tolerance is known as γ-*iteration*. We may refer to a controller as being an $\mathcal{H}_\infty$ optimal controller: what we really mean is a sub-optimal controller to some tolerance.

Finally, note that in robust performance terminology, $\mathcal{H}_\infty$ synthesis is a one block technique, and this block is a performance block. When there is uncertainty present in the model, we will still perform $\mathcal{H}_\infty$ synthesis by collecting all the uncertainty blocks into a diagonal structure with the performance block and covering this structure with one full block. This approach is conservative.

2.4 LPV Synthesis

In this section a brief overview of the LPV synthesis theory is presented. A complete and rigorous explanation of the synthesis technique can be found in [39].

Consider the general time-varying system shown in Figure 2.3, where $x(k), e(k), y(k), d(k)$, and $u(k)$ are the state, error, measurement, disturbance, and input vectors, respectively. We assume the time-variation of the plant can be represented as an LFT of a parameter set and a constant matrix. Thus $P(k)$ is given by

$$P(k) = P_{22} + P_{21}\Delta(k)\left(I - P_{11}\Delta(k)\right)^{-1} P_{12} \tag{2.7}$$

where

$$\Delta(k) = \begin{bmatrix} \delta_1(k)I_{n_1} & & \\ & \ddots & \\ & & \delta_m(k)I_{n_m} \end{bmatrix} \tag{2.8}$$

with $\delta_k \in \mathbb{C}$, $\|\delta_i(k)\| \leq 1$ for all k, and n_i is the dimension of the identity matrix associated with δ_i. We assume that each δ_i can be measured on-line.

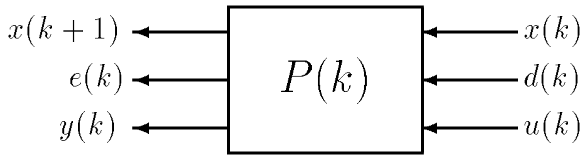

Figure 2.3: Time-varying system.

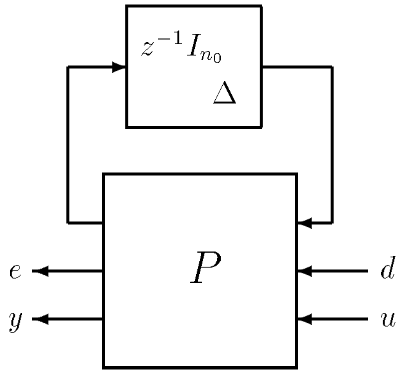

Figure 2.4: Parameter-dependent plant. The $z^{-1}I_{n_0}$ term represents the states of P, and the Δ represents the time variation of Equation 2.8.

Note that any system with rational time-varying entries can be represented in this framework, and many others can be arbitrarily closely approximated. This type of system is known as a parameter-dependent LFT system. The representation of P as an LFT is shown in Figure 2.4.

The controller we will design for this plant will also be parameter-dependent, depending on the same δ_i's as the plant; these copies are collectively denoted by $\widehat{\Delta}$. K thus has the form shown in Figure 2.5. P can be augmented to collect all the time-varying parameters and states together; K can then be treated as a simple matrix. This is depicted in Figure 2.6, where R is the augmented form of P, and K is a matrix. The problem then appears as a robust control problem with a special structure on the plant and parameters. The design objective is to find a controller K such that the interconnection is stable and the $\ell_2 \rightarrow \ell_2$ induced norm from d to e is small for all allowable parameter variations $\Delta(k)$ (see Equation 2.8). Combining the gain from d to e with the gain of $R \star K$ (that is, treating the gain from d to e as a "performance block," another block in the $\boldsymbol{\Delta}$ structure) gives us a small-gain condition. Since the small-gain theorem can be quite conservative, we can reduce the conservatism by introducing scaling matrices from the set $\mathcal{T}$ which commutes with the set of parameter variations.

The resulting condition is then the state-space upper bound (SSUB) of [40]. This condition now becomes (compare Lemma 3.1 of [39] and Theo-

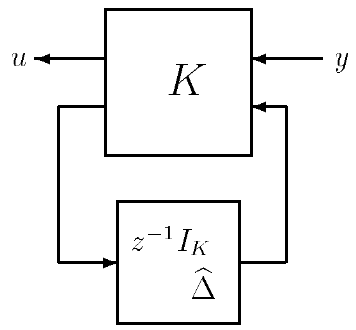

Figure 2.5: Parameter-dependent controller; $z^{-1}I_{r_0}$ represents the states of the controller and $\widehat{\Delta}$ the time variations.

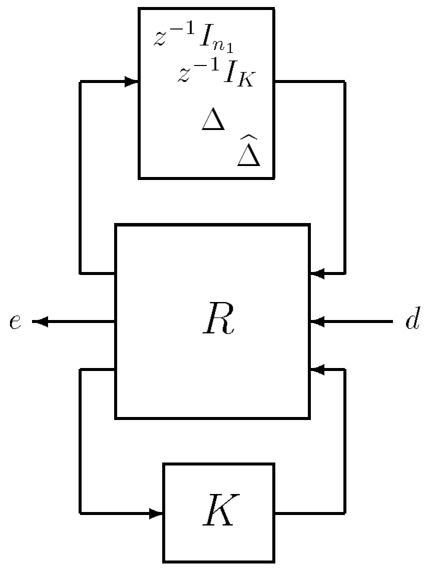

Figure 2.6: Parameter-dependent closed-loop system.

rem 10.4 of [40]):

Theorem 1 *Let R be given as above, along with an uncertainty structure $\mathbf{\Delta}$. If there is a $T \in \mathcal{T}$ and a stabilizing, finite-dimensional, time-invariant K such that*

$$\left\| \begin{bmatrix} T & 0 \\ 0 & I \end{bmatrix} (R \star K) \begin{bmatrix} T^{-1} & 0 \\ 0 & I \end{bmatrix} \right\|_\infty < 1 \tag{2.9}$$

then there is a γ, $0 \le \gamma < 1$, such that for all parameter sequences $\delta_i(k)$ with $\|\delta_i\|_\infty \le 1$, the system in Figure 2.6 is internally exponentially stable, and for zero initial conditions, if $d \in \ell_2$, then $\|e\|_2 \le \gamma \|d\|_2$.

Pictorially, this theorem is shown in Figure 2.7. A natural question arising from this theorem is when does such a K exist for any value of γ, not just $\gamma < 1$? It is a simple corollary of results in [30] that such a K will exist when R is stabilizable and detectable with respect to the block structure

$$\mathbf{\Delta} = \operatorname{diag}\left(z^{-1} I_{n_1}, z^{-1} I_K, \Delta, \widehat{\Delta}\right).$$

The important fact resulting from Theorem 1 is that the synthesis of D and K to meet the objective can be cast as a computationally tractable convex optimization problem involving 3 LMIs. These LMIs have the following form

$$U_\perp^T \left(E \begin{bmatrix} X & 0 \\ 0 & I \end{bmatrix} E^T - \begin{bmatrix} X & 0 \\ 0 & I \end{bmatrix} \right) U_\perp < 0,$$

$$V_\perp \left(E^T \begin{bmatrix} Y & 0 \\ 0 & I \end{bmatrix} E - \begin{bmatrix} Y & 0 \\ 0 & I \end{bmatrix} \right) V_\perp^T < 0,$$

$$\begin{bmatrix} X & I \\ I & Y \end{bmatrix} \ge 0,$$

where $U_\perp$, $V_\perp$, and E are obtained from the system realization, and X and Y are structured positive definite matrices. Interested readers may find the exact LMIs in Theorem 6.3 of Packard [39]. E, U, and V have a scaling γ absorbed into them, and thus the synthesis procedure is a γ-iteration, as $\mathcal{H}_\infty$ is. Once a desired γ level has been reached, a controller K can be obtained by linear algebraic operations on X and Y.

A few points are important in understanding the ramifications of employing the state-space upper bound (SSUB). Most importantly, this technique results in a controller optimal with respect to a time-varying perturbation with memory (the sequence $\Delta(k)$ of Equation 2.8, becomes a time-varying operator with memory, rather than a sequence of complex numbers). The relationship between such an operator and a parameter useful in gain-scheduling

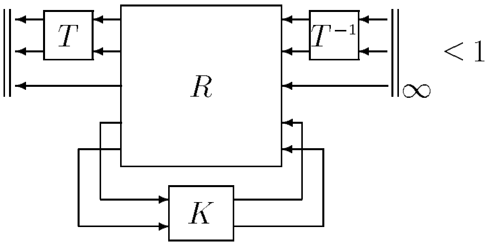

Figure 2.7: Diagram of Theorem 1.

is tenuous, at best. Depending on the problem, this technique could conceivably yield controllers so conservative as to have extremely poor performance. Nonetheless, if a controller with acceptable performance can be designed with this technique, then it will have at least the same level of performance for all variations of the operating point (the operating point is a fixed value of Δ). Additionally, a time-varying operator with memory does not in general have a frequency spectrum, so there is no way to "filter" it to achieve a closer relationship to an operating parameter. Moreover, it is interesting to contrast this technique with μ-synthesis where instead of the SSUB the frequency-domain upper bound is usually employed; this difference reflects the different assumptions about the type of perturbations.

If Δ is a constant value and is "wrapped into" the plant, the resulting model becomes a linear model around the operating point of the Δ. Similarly, we can do this for controllers, and the LPV controller becomes a linear controller. We will refer to the linear controller obtained by holding Δ at a constant value as the LPV controller *locked* at the value of Δ. We are interested in looking at controllers locked in various positions because by comparing them with the full LPV controller we hope to gain better insight into the nature of LPV control.

Chapter 3

Application: Pressurized Water Reactor

I—
really I can do little
as little now as then,
about the infernal fires—
I cannot blow out a match.
—Robert Lowell

In this chapter we present one of the applications discussed in this dissertation — the primary circuit of a nuclear power plant. Our goal is control of this system over the operating range from 50% to 100% of its output power. We show the difficulties a linear controller encounters because of the changing dynamics and how LPV control solves them. For the general control designer, the main points of interest are the construction of the parameter-varying model from identified, rather than analytic, models, and the use of a design weight which varies as the operating point changes.

In Section 3.1 the application and problem are motivated and introduced. Section 3.2 is devoted to a precise description of the problem statement. Section 3.3 describes the identification and modelling of the plant. Section 3.4 presents the design of $\mathcal{H}_\infty$ controllers around two operating points. The main results of the chapter are presented in Sections 3.5 and 3.6 which describe the synthesis and evaluation of LPV controllers.

The material in this chapter arose from joint work with Pascale Bendotti [10, 13]. In particular, she is responsible for the system identification and $\mathcal{H}_\infty$ control design, presented in [11, 12, 10].

3.1 Introduction

Motivation

In France and certain other countries the major contribution to electricity production is provided by nuclear power. When this is the case, the nuclear power plant must provide electricity as it is needed and the plant becomes a time-varying system with dynamics changing slowly as the internal power changes. Nonetheless, large transients can occur, for example, when the plant shuts down. Most nuclear power plants are pressurized water reactors (PWR). The dynamics of a PWR change enough over its operating range that a linear controller cannot guarantee performance over the entire range, especially when operating conditions change suddenly.

If a fixed linear controller is not capable of maintaining performance over the entire operating range, then a possible approach to control a PWR is to design a parameter-dependent controller with the output power as the parameter. One advantage such a controller would have over a standard gain-scheduled controller is that performance and stability could be guaranteed over the operating range of the plant, and large transients in switching are avoided. An additional advantage of LPV synthesis is that the controller is designed in one step, rather than by designing several controllers and then scheduling them. The potential drawback of LPV synthesis is that the technique is conservative. This conservatism may be so great that the controller performs quite poorly.

3.2 Problem Statement

The main objective in controlling a PWR is to provide the commanded power while respecting certain physical constraints. Consider the application depicted in Figure 3.1. This is the primary circuit, and our goal is to control this part of the reactor. The pressurized water in the primary circuit transmits the heat generated by the nuclear reaction to the steam generator. In the steam generator, water of the secondary circuit turns into hot steam, which drives a turbo-alternator to generate electricity. The rate of the reaction is regulated by the control rods. The rods capture neutrons, slowing down the nuclear reaction; withdrawing the rods increases the reaction. The PWR has two independent sets of rods which are used as controls.

The PWR has an inner control loop which holds the pressure in the primary circuit constant. Thus for a steam flow increase in the secondary circuit, the temperature in the primary circuit will decrease. From a control standpoint, the required power corresponds to a specific steam flow that may be viewed as a measurable disturbance. Hence, one natural control objective is to track a temperature reference derived from the steam flow. Because of the way in which the control rods enter the reactor, the rate of reaction is always higher at the bottom of the reactor. The *axial offset* is defined as the difference in power generated between the top and bottom of the PWR. Safety specifications require minimizing the axial offset; this also increases the lifetime of the fuel and reduces operating costs. To achieve such objectives two control inputs are available, the rates of motion of the control rods, denoted u_1 and u_2. The positions of the control rods are denoted v_1 and v_2, respectively. The positions are, of course, measurable. Due to the physics of the reactor, u_2 has

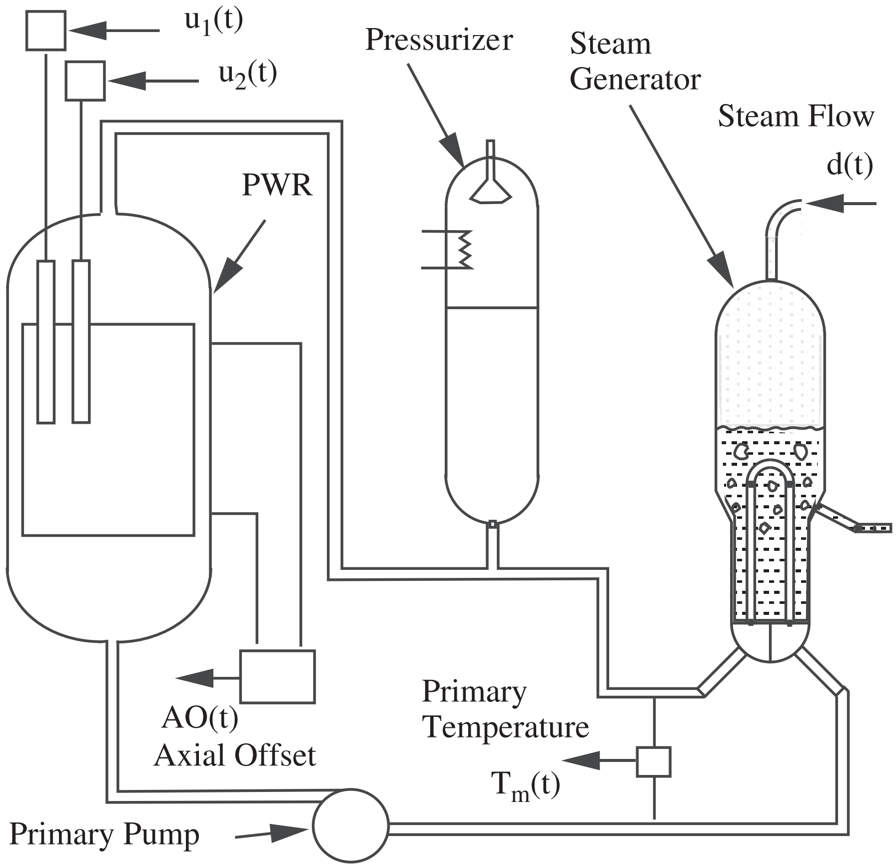

Figure 3.1: Primary circuit and steam generator.

more authority than u_1 at low power and using it results in a smaller axial offset. At high power, however, u_2 has almost no authority, so all control must come from u_1.

Due to the complexity of the physical plant, performance specifications cannot be uniquely or easily derived. Indeed, investigations into the best performance specifications are currently underway at Electricite de France (EDF). Below are the first attempts to automatically control the axial offset, a specification we will also use here. Nonetheless, we do not have precise specifications the controllers must meet.

3.3 System Identification and Modelling

The first step of any design procedure is to obtain a model. Here we review the identification process for the PWR. Our goal is to obtain a reasonably low order model for the plant.

Identification Experiments

The identification experiments were carried out using a realistic nonlinear simulator developed at EDF. The simulator is based on various finite element models of the PWR.

The system possesses nonlinearities of two types. The first depends on the operating condition and hence is strongly related to the commanded power. No *a priori* knowledge can be used in the identification process for this type of nonlinearity, so the experimental data are obtained around different operating points and the resulting model is a linearization at the operating point. The second nonlinearity is on the input magnitude of v_2. This control becomes ineffective when the commanded power tends to its maximum. This maximal value is usually referred to as the nominal power of the plant, $\mathcal{P}_n$. The static characteristic of the input effectiveness is actually known, so its inversion allows identification close to the nominal power, where the nonlinear effect is maximal.

MIMO State-space Description

Consider the system depicted in Figure 3.2, where T_m, AO, P_I, d, v_1 and v_2 are the temperature, the axial offset, the power, the steam flow and the vertical positions of the rods, respectively.

The physical system is described by a linear time-invariant (LTI) system around an operating point given by the following:

$$x(t+1) = Ax(t) + Bv(t) + \Gamma d(t) \tag{3.1}$$

$$y_s(t) = Cx(t) + Dv(t) \tag{3.2}$$

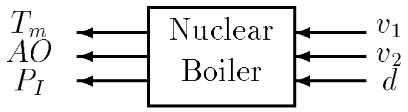

Figure 3.2: Input/Output diagram of the primary circuit.

with

$$y_s(t) = \begin{pmatrix} T_m(t) \\ AO(t) \\ P_I(t) \end{pmatrix}$$

and

$$v(t) = \begin{pmatrix} v_1(t) \\ v_2(t) \end{pmatrix}$$

where $x(t)$, $y_s(t)$, $v(t)$, and $d(t)$ represent the state, the output, the input, and the disturbance at time t, respectively. The parameters of the state-space realization consist of the elements of the A, B, C, D, and Γ matrices.

Since the number of parameters rises quadratically with the state dimension, there are a large number of them in a state-space realization. To meaningfully reduce them, specific realizations are used where some parameters are fixed at either zero or one, for example, the well known MIMO canonical forms. Unfortunately, these realizations still contain too many parameters to be uniquely identified.

Re-parameterizing the realizations using physical considerations can overcome this problem. Preliminary identification of several SISO and MISO transfer functions are performed providing insight into an appropriate re-parameterization (*cf.* references in [11]). Indeed, the primary temperature and power are mainly related to the control inputs by a second and first order system, respectively. Furthermore, the inputs affect the plant dynamics in an identical manner, although the gains are different. The axial offset is almost a linear combination of the inputs: thus no states are needed for it. These insights provide an appropriate identification-oriented state-space realization structure. Hence, only the temperature and the power have dynamics. The effect of the disturbance has a larger delay than the effect of the control and hence the dimension of the state must reflect this. More precisely, 3 delay values for each of temperature and power (instead of 2 and 1, respectively, in the disturbance-free case) are required to appropriately predict the input-output behavior. This leads to a sixth order state-space realization defined as follows:

$$A = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ a_1^{11} & a_2^{11} & a_3^{11} & 0 & a_2^{12} & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ a_1^{21} & 0 & 0 & a_1^{22} & a_2^{22} & a_3^{22} \end{pmatrix}$$

$$B = \begin{pmatrix} b_1^1 & b_2^1 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \ \Gamma = \begin{pmatrix} 0 \\ 0 \\ \gamma_3 \\ 0 \\ 0 \\ \gamma_6 \end{pmatrix} \tag{3.3}$$

$$C = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & c_2^2 & 0 & c_4^2 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}, \ D = \begin{pmatrix} 0 & 0 \\ d_2^1 & d_2^2 \\ d_3^1 & d_3^2 \end{pmatrix}.$$

The state matrix A can be partitioned as follows:

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

where A_{11} and A_{22} denote the third order systems for the temperature and the power; A_{12} and A_{21} represent the cross-coupling matrices which contain only one non-zero term, appropriately located. The input matrix B accounts for the actual delayed effect of the control inputs on primary temperature and power, so only the first row of B has non-zero elements. Similarly, the delay between the control inputs and the disturbance can be taken into account using only the third and sixth elements of the disturbance input matrix Γ. The second row of the output matrix C adds memory to the axial offset. Finally, the elements of D correspond to the direct terms appearing in the axial offset and the power. This results in a specific identification-oriented realization with 18 parameters, instead of the 28 parameters in the standard canonical form.

MIMO Identification

Similarly, the overall system (3.1)-(3.2) can also be modelled by the transfer function:

$$y_s(t) = G_v(\theta, q)v(t) + G_d(\theta, q)d(t) \tag{3.4}$$

where q denotes the standard forward shift operator (the corresponding z operator will be omitted for simplicity), and θ represents the vector of free parameters to be identified.

Given a description (3.4) properly parameterized by the specific form (3.3) and the input–output data v, y_s and d, the prediction error ε is computed as follows:

$$\varepsilon(t) = y_s(t) - G_v(\theta, q)v(t) - G_d(\theta, q)d(t).$$

The identification method consists in determining the parameter estimates by minimizing the following quadratic criterion:

$$\widehat{\theta} = \arg \min_{\theta} \det \left[\frac{1}{N} \sum_{i=1}^{N} \varepsilon(t)\varepsilon^T(t) \right]$$

using an iterative Gauss-Newton algorithm [29]. See the discussion in Section 6.1 for related details.

Confirmation and Results

Finally, we confirm the identified model is accurate by checking how well it predicts the behavior of the physical system when simulated with a different data set to see how well it matches the actual output of the physical system. This procedure was successful over a large operating range due to the static inversion performed at the plant input. In particular, the specific form used for the parameterization was verified.

The time-domain responses of the identified model obtained around $0.5\mathcal{P}_n$ (dashed) are plotted against the experimental data (solid) in Figure 3.3. The inputs used to generate this data — v_1 (solid), v_2 (dashed) and d (dotted) — are plotted in the right lower diagram in Figure 3.3. The step-responses of the identified models obtained around $0.5\mathcal{P}_n$, $0.9\mathcal{P}_n$ and $0.99\mathcal{P}_n$, called $\widehat{G}_0$, $\widehat{G}_1$, and $\widehat{G}_2$, respectively, are shown in Figure 3.4.

Model Reduction

Since our design methods yield controllers with state dimension equal to that of the open-loop interconnection structure, often we can reduce the order of the controller by reducing the plant model before controller synthesis. In the present case we are reducing the sixth order models. A balanced realization technique [32], including specified model reduction weightings, is used [21].

In particular, dynamic behavior at high frequencies can be considered uncertainty. Therefore, the measurements are weighted with low-pass filters to attenuate the high frequency dynamics. Finally, the reduced order model is obtained by truncating weakly controllable and observable states. The resulting MIMO reduced order plant model is first order, i.e., only the dominant mode is retained.

The nominal reduced order plant model is the design model G_0 while the nominal plant model is $\widehat{G}_0$. Figure 3.4 shows the step-responses of $\widehat{G}_0$ in

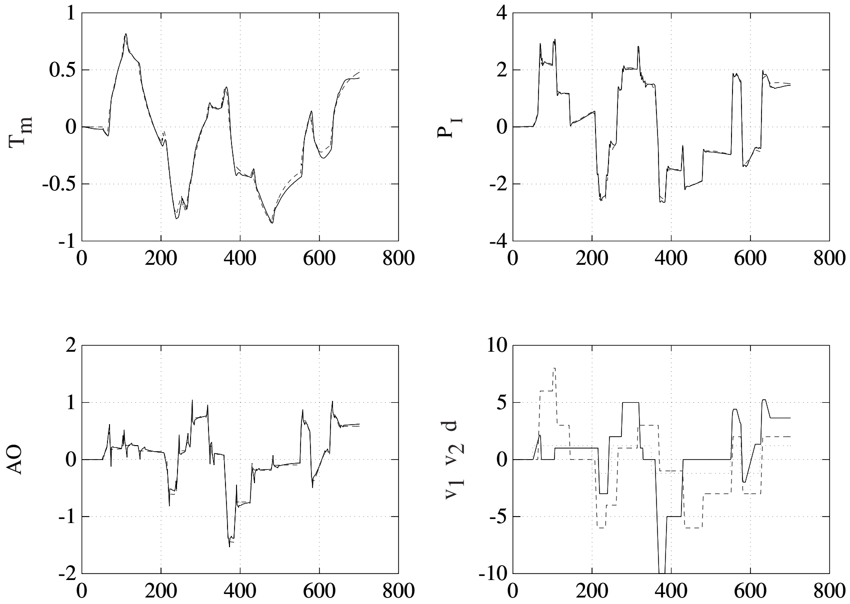

Figure 3.3: Experimental data versus time-domain responses of $\widehat{G}_0$.

solid lines against those corresponding to G_0 in dotted lines. For purposes of comparison, those corresponding to $\widehat{G}_1$ (dashed) and $\widehat{G}_2$ (mixed) are plotted on the same graphs.

Due to the lower order approximation, model inaccuracy is unavoidable. Since the high frequency dynamics are no longer modelled, there is a significant difference between the identified model and the reduced order model. Figure 3.5 shows the Bode plots corresponding to the multiplicative-errors relating the design model G_0 to $\widehat{G}_0$ (solid), $\widehat{G}_1$ (dotted) and $\widehat{G}_2$ (dashed), respectively. Please note, the figure also contains weights which will be explained later.

Parameter Dependence

Our model in the form of Equation 2.7 will be derived using the first-order models of the previous section. To derive the parameter dependence, each term of the three first-order models is compared; those which vary are individually fitted with a rational function of δ, $-1 \leq \delta \leq 1$, using a least-squares technique. For the PWR, first order LFTs of the form $e + f\delta(1 - g\delta)^{-1}h$ fit the parameters extremely well, as shown in Figure 3.6. In this figure, $0.5\mathcal{P}_n$ corresponds to $\delta = -1$, $0.9\mathcal{P}_n$ corresponds to $\delta = 0.6$, and $0.99\mathcal{P}_n$ to $\delta = 0.998$ (these are the asterisks in the figure). The resulting model with δ-dependence,

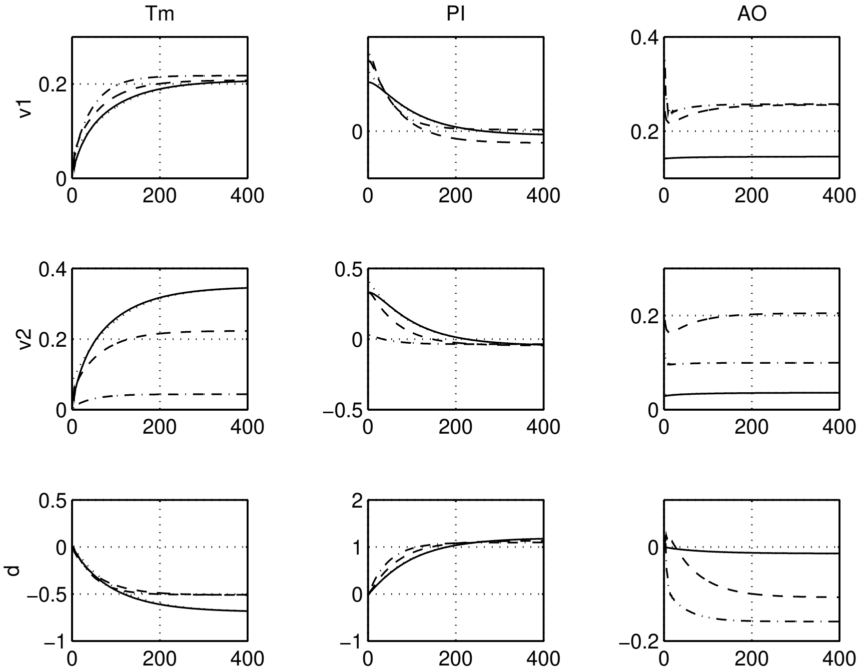

Figure 3.4: Step-responses of $\widehat{G}_0$, G_0, $\widehat{G}_1$ and $\widehat{G}_2$.

$P(\delta)$, becomes

$$\left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right] = \left[\begin{array}{c|ccc} a(\delta) & b_1 & b_{v_1}(\delta) & b_{v_2}(\delta) \\ \hline c_1 & d_{11} & d_{12} & d_{Tm_2}(\delta) \\ c_2 & d_{12} & d_{22} & \kappa b_{v_2}(\delta) \\ c_{AO}(\delta) & 0 & d_{AO_1}(\delta) & d_{33} \end{array}\right]. \tag{3.5}$$

The inputs for this model are the steam disturbance d, v_1, and v_2; the outputs are the mean temperature T_m, the power $\mathcal{P}$, and the axial offset AO, respectively (see Figure 3.1). Placing this model in the form of Equation 2.7 results in a system shown in Figure 2.4, where $n_0 = 1$ and $\Delta = [\delta I_6]$.

From Figure 3.6, notice the system matrix $a(\delta)$ is inversely proportional to the operating power and the time constant changes by a factor of 2 over the operating range. Also, the variation of b_{v_2} and d_{P_2} differs only by a constant, κ, which is used to reduce the size of the final Δ-block. More importantly, the effectiveness of u_2 decreases as the power increases, and is almost zero at full power. The gain in the axial offset channel increases as power increases, making it more difficult to control at high power. In particular, the effect of u_1 on the axial offset (d_{AO_1}) increases, while the effect of u_2 decreases. This makes it practically impossible to require any performance on axial offset at high power.

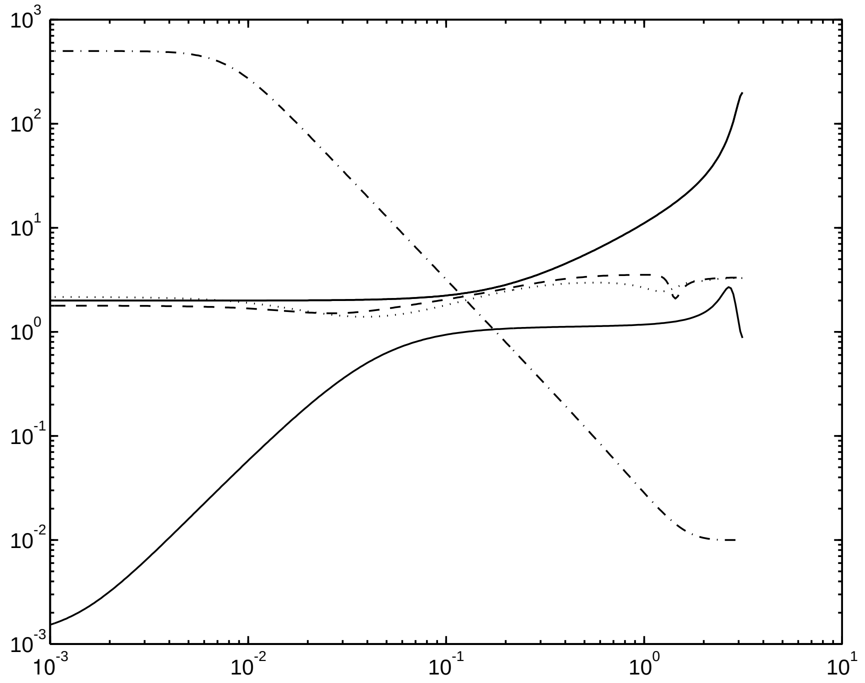

Figure 3.5: Uncertainty-performance weights and relative errors relating G_0 to $\widehat{G}_0$, $\widehat{G}_1$, $\widehat{G}_2$, respectively. $\widehat{G}_0$ is in solid lines, $\widehat{G}_1$ is in dotted lines, and $\widehat{G}_2$ is in dashed lines. W_m is also shown in solid lines and W_{T_m} is mixed lines.

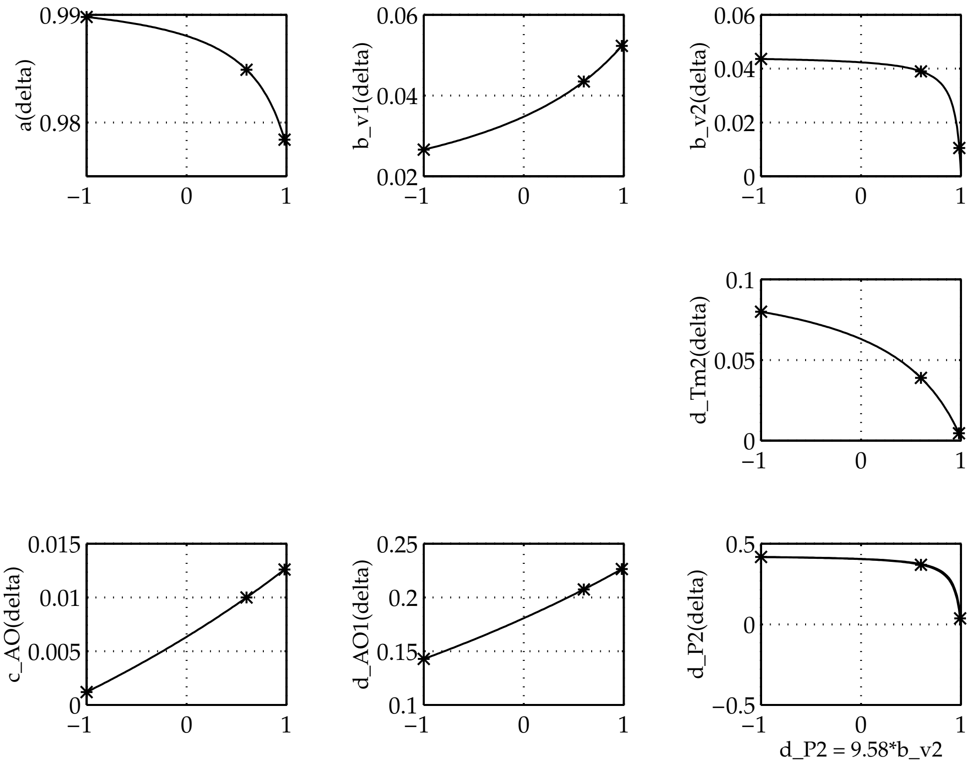

Figure 3.6: Parameter variations versus δ for the model of Equation 3.5. A '*' shows an actual value, and the line shows the LFT fit.

3.4 $\mathcal{H}_\infty$ Controller Design

The first controllers are designed using an $\mathcal{H}_\infty$ methodology; this gives us a first approximation for the weights used in LPV synthesis. Recall that $\mathcal{H}_\infty$ synthesis is aimed at disturbance rejection. A tracking problem such as the PWR can be cast as disturbance rejection by rejecting the low frequency components of the error between plant output and the reference. As the synthesis is in continuous time, the weighting functions are specified in continuous time as well. Then the discrete time $\mathcal{H}_\infty$ controller is obtained using the bilinear transformation. The LPV design is carried out in discrete-time.

Uncertainty Description

As the controllers must stabilize the actual plant, our design methodology must account for the discrepancy between model and reality. We employ the standard approach of designing a controller stabilizing the nominal model in the presence of modelling errors.

A multiplicative-error is used to provide a description of the plant mismatch as well as a characterization of robust stability.

Consider the identified plant model (3.4) and rewrite it as follows:

$$\begin{pmatrix} y_1(t) \\ y_2(t) \end{pmatrix} = \underbrace{\begin{pmatrix} G_{1v} & G_{1d} \\ G_{2v} & G_{2d} \end{pmatrix}}_{G_0} \begin{pmatrix} v(t) \\ d(t) \end{pmatrix} \tag{3.6}$$

where y_1 denotes the controlled outputs and y_2 is an auxiliary output:

$$y_1(t) = \begin{pmatrix} T_m(t) \\ AO(t) \end{pmatrix} \text{ and } y_2(t) = P_I(t).$$

The plant model description corresponding to the identified model $\widehat{G}_0$ is obtained by replacing G with $\widehat{G}$ in (3.6).

Given a nominal model G_{1v} as well as the weighting function W_m, the multiplicative model set is defined as:

$$\Sigma(G_{1v}, W_m) = \{G_{1v}(I + \Delta_m W_m) \mid \Delta_m \text{ stable}, \|\Delta_m\|_\infty \leq 1\}$$

where $\|\cdot\|_\infty = \max_\omega \overline{\sigma}(\cdot)$.

A typical robust stability test for multiplicative perturbations is to find a stabilizing controller K which minimizes

$$\|W_m \, T_i\|_\infty \tag{3.7}$$

where T_i is the plant input complementary sensitivity function and W_m is the multiplicative uncertainty weight specifying the amount of uncertainty in the model as a function of frequency.

In the present case, the uncertainty weight is of the form $W_m = w_m\ I_2$, where w_m is a stable minimum-phase scalar valued function and has a large magnitude in the frequency range where the modelling error is too large; w_m is chosen as follows:

- In the frequency range where known dynamics have been neglected,

$$|w_m| \geq \|G_{1v}^{-1}(\widehat{G}_{1v} - G_{1v})\|_2 \tag{3.8}$$

  where $\|\cdot\|_2 = \overline{\sigma}(\cdot)$ denotes the maximum singular value of a matrix, and $\widehat{G}_{1v}$ is the identified transfer function from v to y_1.

- Outside the frequency range of the experiment, $|w_m|$ is large to account for unmodelled dynamics.

- In the frequency range where the model is accurate, $|w_m|$ is chosen to account for nonlinearities in the physical plant.

Figure 3.5 on page 26 shows W_m (solid) and the relative modelling errors (3.8) relating G_0 to $\widehat{G}_0$ (solid), $\widehat{G}_1$ (dotted) and $\widehat{G}_2$ (dashed), respectively.

Performance Specifications

The synthesis structure is shown in Figure 3.7. Recall from Section 2.3 that our two block problem is treated as an $\mathcal{H}_\infty$ problem by covering the uncertainty structure with one full block and not exploiting the structure of the problem. This is conservative, and the diagram is drawn this way (as an implicit two block problem) to remind us of that. The design model includes the actuator dynamics, modelled by two integrators. Also, the vertical positions of the control rods are measured. In this figure, d, n, r, u, and y refer to the disturbance, noise, reference, control, and measurement signals, respectively.

To have nominal performance we attempt to find a K which minimizes

$$\|W_p\ S_o\|_\infty \tag{3.9}$$

where S_o denotes the plant output sensitivity function and W_p denotes the diagonal weighting matrix reflecting the performance specifications. W_p is given as follows:

$$W_p = \begin{pmatrix} W_{T_m}(j\omega) & & & \\ & W_{P_I} & & \\ & & W_{AO} & \\ & & & W_{pos} \end{pmatrix} \tag{3.10}$$

which weights the performance on temperature, power, axial offset, and vertical position of the control rods. To insure low steady-state error in tracking

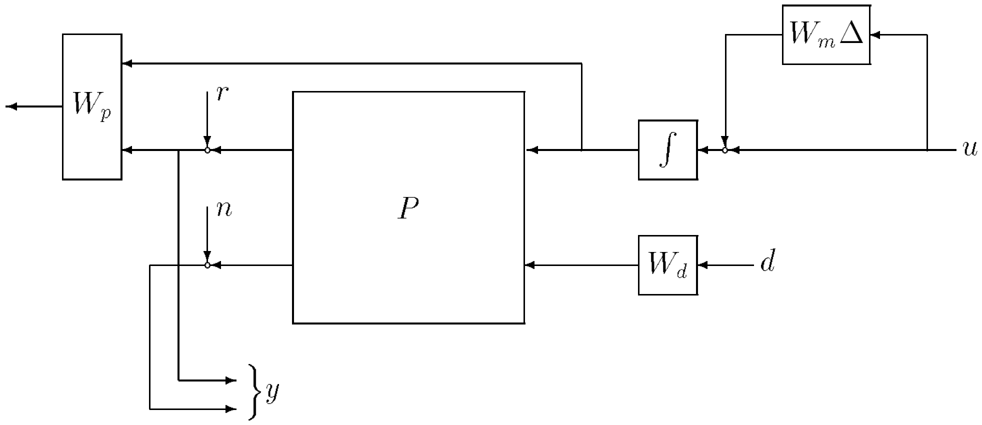

Figure 3.7: Synthesis structure for the PWR.

and to reject step disturbances, W_{T_m} resembles an integrator. Figure 3.5 on page 26 shows the performance weight on W_{T_m} in mixed lines. The weight on the power, W_{P_I} is a constant. As a second objective, the control strategy should minimize the effect of the control on the axial offset. A constant weight W_{AO} is introduced on the axial offset. This causes the use of u_2 to be preferred over u_1 since it has more authority at low power and results in lower axial offset. Because the system has fewer degrees of freedom than performance measures, it is only possible to minimize the axial offset, not to reject it.

To limit the magnitude of the positions of the control rods, a constant weight W_{pos} is used. As an aside, it is interesting to note that if W_{pos} is omitted, both the $\mathcal{H}_\infty$ and LPV methods will produce controllers which give a dramatically lower axial offset. They do this by moving the control rods in opposition to one another, which clearly will produce lower power generation at the top and bottom of the vessel. Unfortunately, this type of motion is not physically possible on this reactor. We eliminate it by placing a penalty on the movement, which works since when the rods move in opposition to one another, they must move more to achieve the same affect on the temperature. Penalizing their movement causes them to move together, at the expense of the axial offset.

Synthesis

Once the weights have been selected, the design process is simply an iteration on improving the weights to get a satisfactory controller, using the process detailed in Section 2.3. Two controllers were designed, for the operating points at $0.5\mathcal{P}_n$ and $0.99\mathcal{P}_n$. The weights were different to optimize the performance at the different operating conditions. These controllers will be termed "H50" and "H99" in the sequel.

Robustness Analysis

The $\mathcal{H}_\infty$ design for H50 is analyzed with respect to structured uncertainty using μ [40]. The upper and lower bounds for μ are calculated on the 7x7 closed-loop response of $F_l(P, K)$ using the following structure:

$$\mathbf{\Delta} = \left\{ \begin{bmatrix} \Delta_m & \\ & \Delta_p \end{bmatrix} : \Delta_m \in \mathbb{C}^{2\mathrm{x}2}, \Delta_p \in \mathbb{C}^{5\mathrm{x}5} \right\}$$

where Δ_m and Δ_p are the uncertainty and performance blocks, respectively.

The bounds for μ with respect to this block structure are plotted in Figure 3.8 in solid lines (they lie on top of one another) along with the maximum singular value in dotted lines. Furthermore, the maximum singular values for robust stability (dashed) and nominal performance (mixed) as defined in (3.7) and (3.9) are shown in the same plot.

3.5 LPV Synthesis

Once the parameterized model $P(\delta)$ is obtained, the controller design becomes similar to the $\mathcal{H}_\infty$ design of the previous section. The synthesis structure used is the same as for the $\mathcal{H}_\infty$ synthesis, shown in Figure 3.7, with uncertainty and performance weights included. The values these weights take for the LPV design is shown in Figure 3.9.

In general, for a system with widely varying dynamics, the same performance requirements over the entire operating range may not be desirable or provide adequate performance. For LPV synthesis, the solution to this problem is to incorporate parameter variations (depending on δ) into the weights. This may be inadvisable, however, because including a δ-dependence in the weight increases the size of the time-varying perturbation block, which may degrade performance. Thus the applicability of this technique must be determined on a problem-by-problem basis.

For the PWR, the same performance requirements over the entire operating range are not desirable.. At low power, the axial offset can be minimized much better than at high, as previously noted. Thus the weight on the axial offset will contain a weight depending on δ, which requires higher performance at low power than at high power. This weight is shown, as a function of δ, in Figure 3.10

Two controllers were designed. The first is called "LPV #1" and is an LPV controller with the weights of Figure 3.9. The second is called "LPV #2" and uses the same weights as LPV #1, except W_{AO} was allowed to depend on the operating power, according to Figure 3.10.

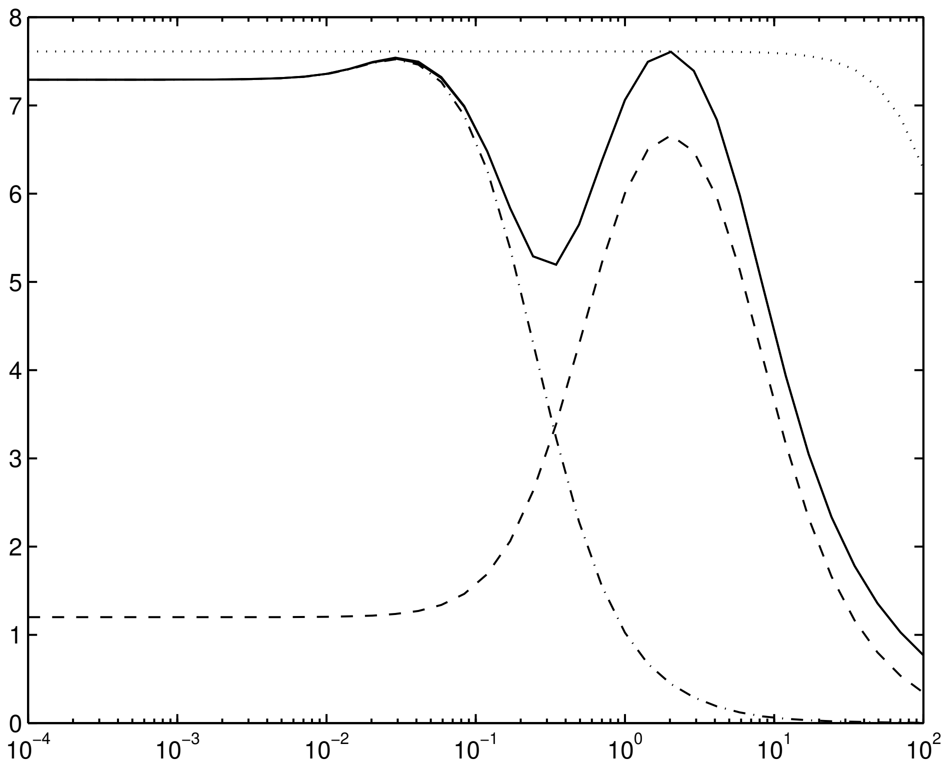

Figure 3.8: μ bounds and maximum singular values for robust stability-nominal performance. The upper and lower bounds for μ are plotted in solid lines, the closed-loop maximum singular value is shown in dotted lines, and the maximum singular values for robust stability and nominal performance are shown in dashed and mixed lines, respectively.

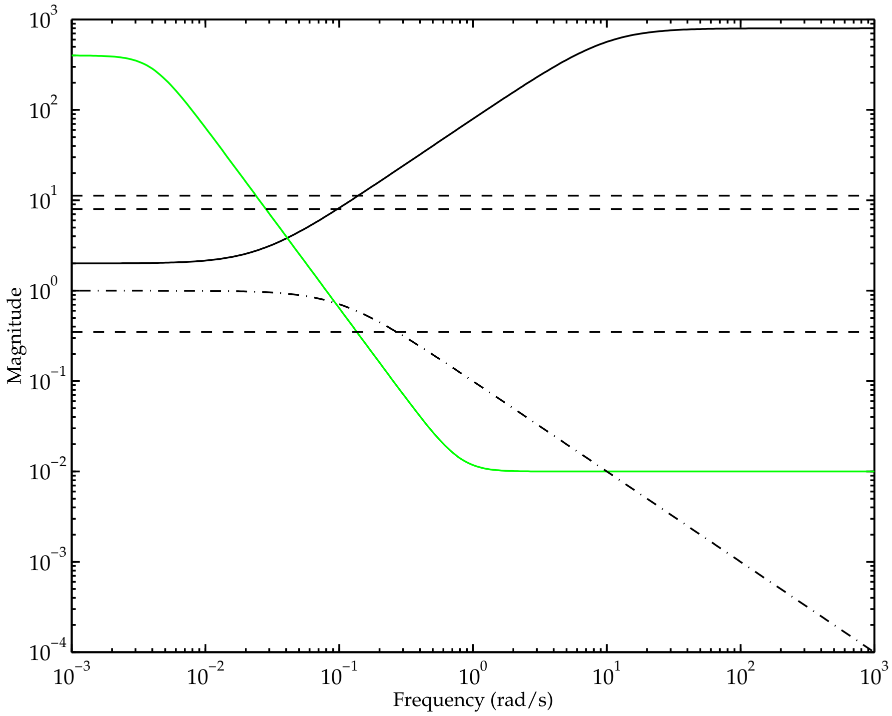

Figure 3.9: Design weights for the PWR LPV controller. The solid line is W_m; the dark shaded line, W_{T_m}; the mixed line is a weight on the disturbance input; the dashed lines correspond to W_{AO}, W_{P_I}, and W_{pos}, in decreasing order of magnitude.

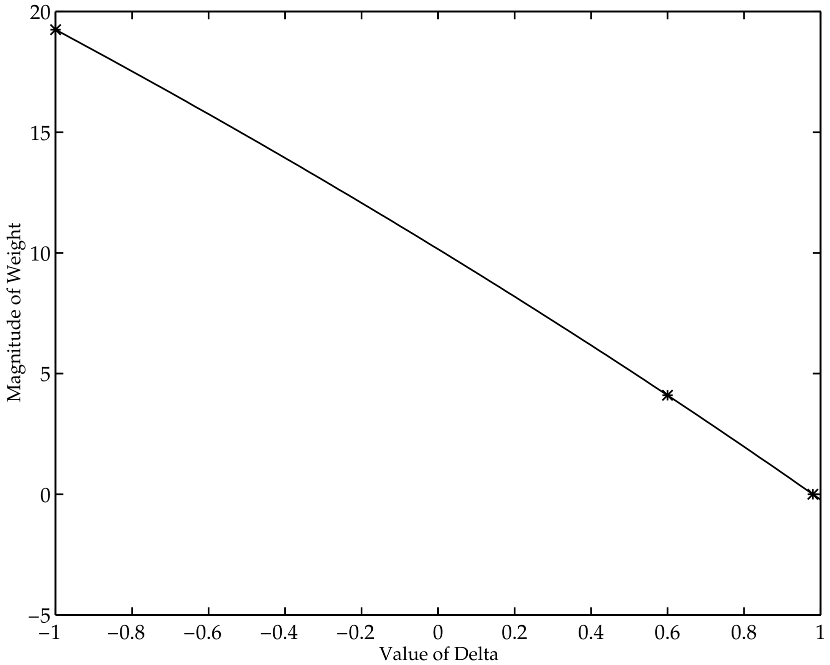

Figure 3.10: Performance weight on the axial offset as a function of δ, where $P_I = (\delta + 3)/4$. A '*' shows a value corresponding to $0.5\boldsymbol{\mathcal{P}}_n$, $0.9\boldsymbol{\mathcal{P}}_n$, or $0.99\boldsymbol{\mathcal{P}}_n$.

3.6 Evaluation

In this section we evaluate the LPV controllers by comparing them with H50 and H99.

Figure 3.12 shows the step responses of the closed-loop systems consisting of each of the controllers and a linearization of the plant at $0.99\mathcal{P}_n$. Step responses are shown because we are interested in the low frequency rejection properties of the closed-loop system. In the first column of plots, the dashed lines are the reference signals, the solid lines are the responses with the first LPV controller, the light shaded lines are with the second LPV controller, and the dark shaded lines are with H99. The second column of plots shows u_1 and u_2 for each of the controllers; u_1 is the solid line and u_2 the shaded one. Figure 3.13 is identical to Figure 3.12 except that the responses are with respect to a linearization of the plant at $0.5\mathcal{P}_n$.

At high power the plant is more difficult to control, because the control rods are almost withdrawn from the reactor. Referring to Figure 3.12, the LPV controllers are almost identical in behavior. They perform equally well, but are not as fast as H99, although they have no overshoot on the temperature. The noticeable difference is that the LPV controllers have less axial offset than H99. At this power, we consider LPV #2 the best of these controllers.

Some of this behavior is preserved in Figure 3.13, but the model is quite different here. Here H50 is slightly faster than the LPV controllers. The major difference at this power is that u_2 now has more control authority than u_1, so controllers do better to use it more, since this results in lower axial offset. H50 does use u_2 more, and the axial offset is considerably lower. At this operating point, we consider H50 the best controller.

At low power, u_2 is the dominant control, but as the power increases u_1 should be used more and more to better meet the control objectives. The LPV controllers do not change strategy between these operating points. Notice that the control plots for LPV #2 are almost identical, up to a scale change in magnitude. This is probably a result of the worst-case nature of LPV controllers. Since achieving worst-case performance does not require a change of strategy, and may in fact forbid one, the controllers do not change their use of the inputs.

Next, the behavior of the LPV controllers on a nonlinear simulator of the PWR is shown. This simulator is not the simulator the synthesis model was identified from, and has less accurate dynamics. We used this simulator because EDF would not allow us to use the original simulator. Nonetheless, the one used is reasonably accurate and provides a satisfactory way to simulate the behavior of the closed loop system. The simulator includes models for the pressurizer, steam generator, and turbine, but not the alternator. The largest underlying change between the simulator used for identification and the simulator used for evaluation is that the former simulator was used assuming the nuclear fuel was new, while the latter is configured for nuclear fuel which

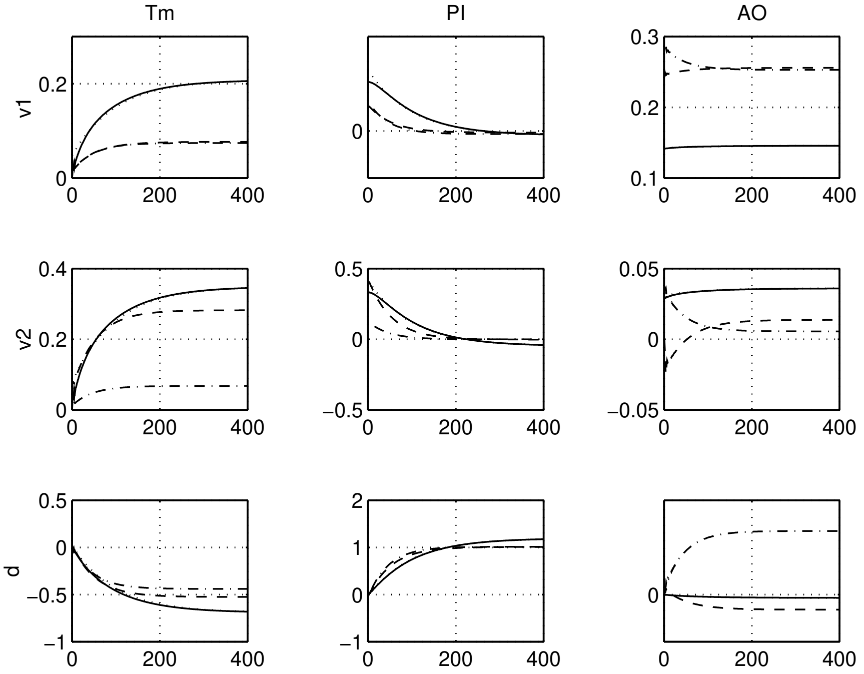

Figure 3.11: Step-responses of models identified from the new nonlinear simulator. The responses of the original model around $0.5\mathcal{P}_n$ are shown in solid lines, while the responses of the new model around $0.5\mathcal{P}_n$ is shown in dashed lines, and the responses of the new model around $0.99\mathcal{P}_n$ are shown in mixed lines.

is at half its expected lifetime. New nuclear fuel is more active than older fuel, and thus the plant dynamics are somewhat different.

To compare the new nonlinear simulator with the old, the same identification experiments that were done for the original simulator were repeated for the new simulator. Figure 3.11 shows the step responses of these models and can be compared with the identified models of Figure 3.4. In Figure 3.11, the responses of the identified model around $0.5\mathcal{P}_n$ from the original simulator are shown in solid lines, while the responses of the identified model around $0.5\mathcal{P}_n$ from the new simulator are shown in dashed lines, and the responses of the identified model around $0.99\mathcal{P}_n$ are shown in mixed lines. The responses vary — sometimes dramatically — and the performance of the LPV controllers under this simulator provides an indication of the robustness of these designs.

Also, control systems for a PWR normally have dead-bands included to prevent moving the control rods for small changes in operating conditions. These, as well as saturations on the controls, have been removed for the purposes of this study. Finally, the control system contains a static nonlinearity on u_2 which reflects an *a priori* knowledge on its loss of effectiveness as the

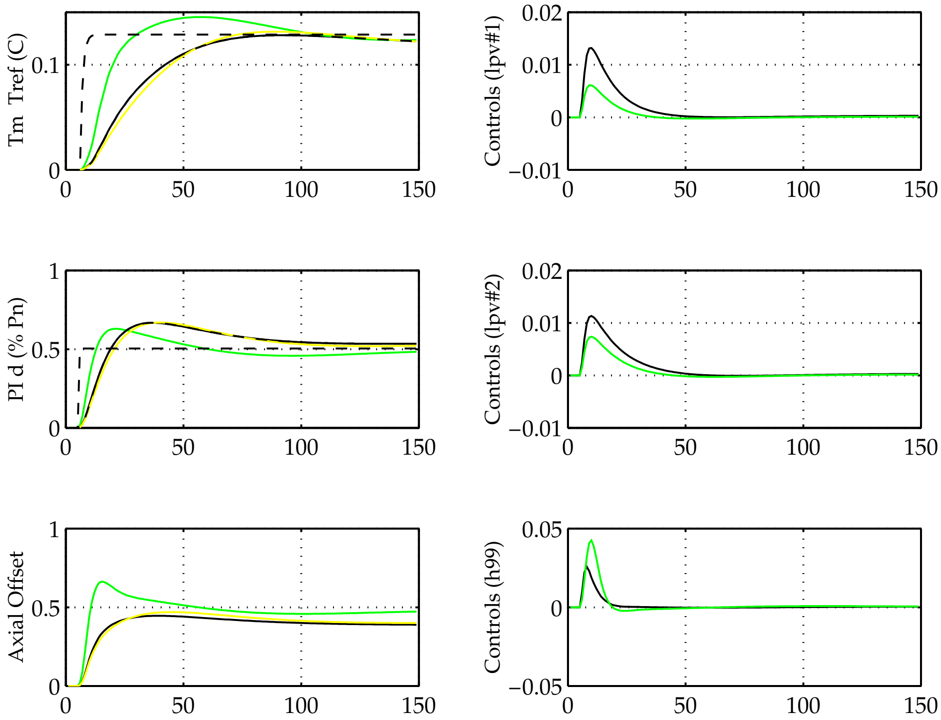

Figure 3.12: Comparison of Three Controllers at $0.99\mathcal{P}_n$.

power increases. The nonlinearity proved necessary in the identification of the models, and the control system simply inverts it out. In particular, this will explain the larger magnitude of u_2 at high power (the nonlinearity is unity at low power).

Figures 3.14 through 3.16 show the simulation results. In these figures, the response of LPV#1 is shown in shaded lines, the response of LPV#2 is shown in solid lines, and the references are shown in dashed lines. Also, u_1 is shown in shaded lines, and u_2 in solid lines. Figure 3.14 shows the response to a one percent step around $0.99\mathcal{P}_n$. LPV#2 is faster, and introduces less axial offset. This difference is even more noticeable in Figure 3.15, which is a two percent step around $0.5\mathcal{P}_n$. Comparing these results to the linear simulations, we see that there is overshoot and the response is slower. LPV#2 clearly outperforms LPV#1 in the nonlinear simulations.

Finally, the response of the LPV controllers to a large transient is shown in Figure 3.16. This is a ramp of -30%/minute from $\mathcal{P}_n$ to $0.5\mathcal{P}_n$. There is not much difference in either LPV controller on this trajectory. This is not surprising as stability for large transients is inherent in the LPV methodology, provided that the synthesis model is accurate over the operating range.

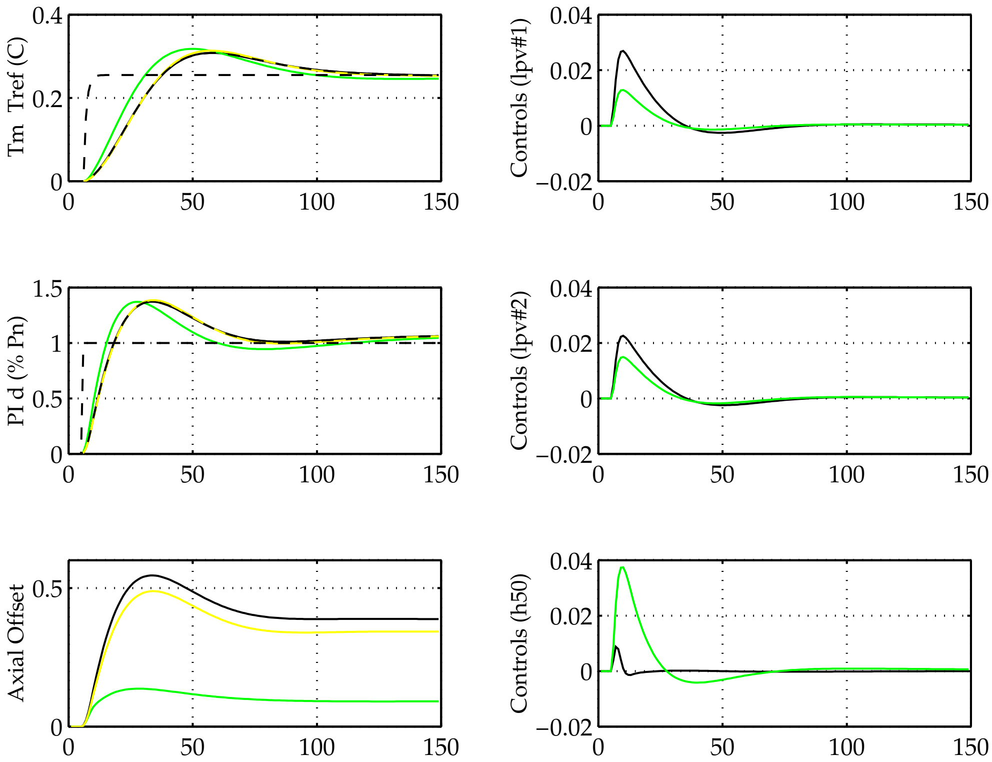

Figure 3.13: Comparison of Three Controllers at $0.5\mathcal{P}_n$.

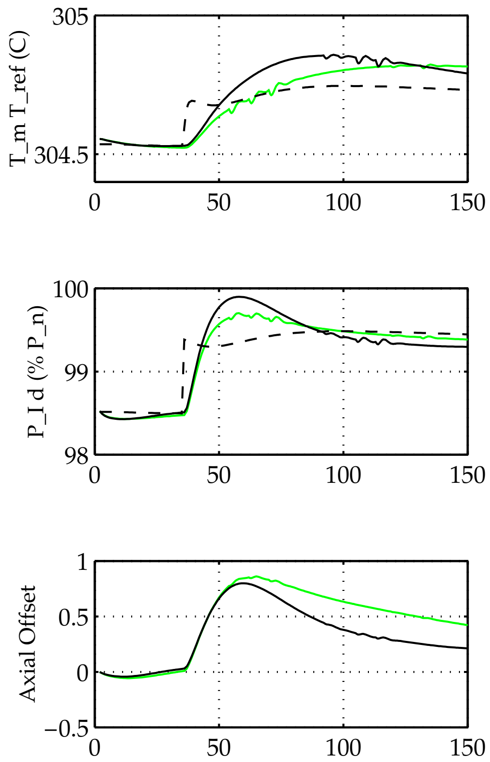

Figure 3.14: Comparison of LPV controllers on a nonlinear simulation to a step around $0.99\mathcal{P}_n$.

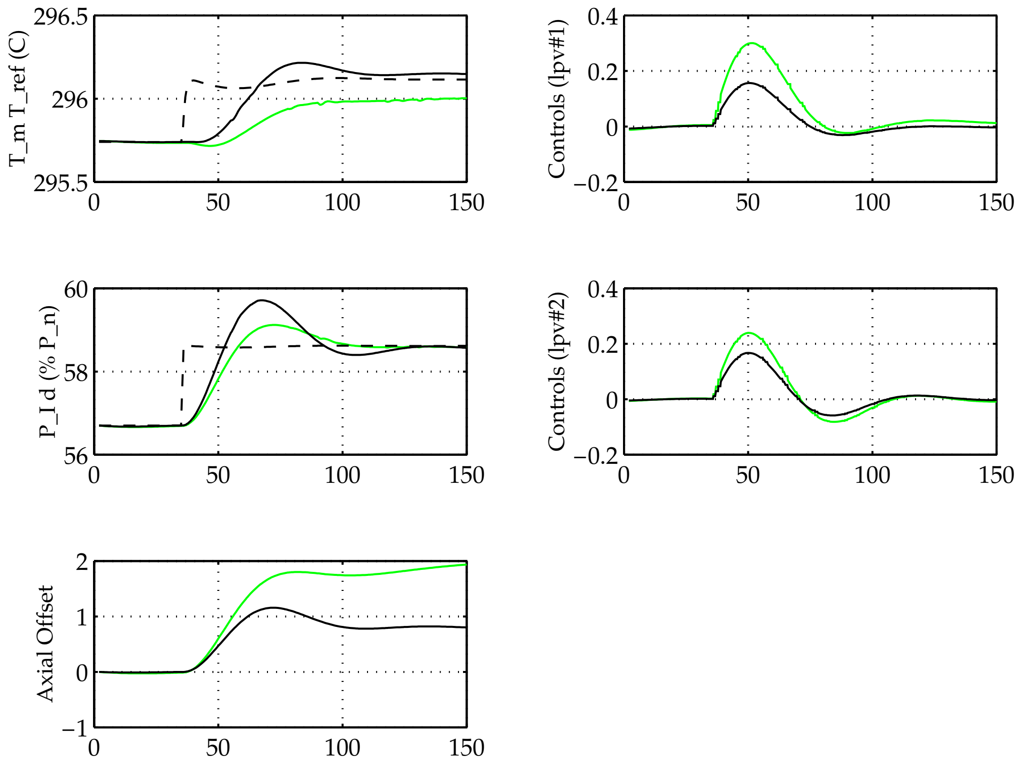

Figure 3.15: Comparison of LPV controllers on a nonlinear simulation to a step around $0.5\boldsymbol{\mathcal{P}}_n$.

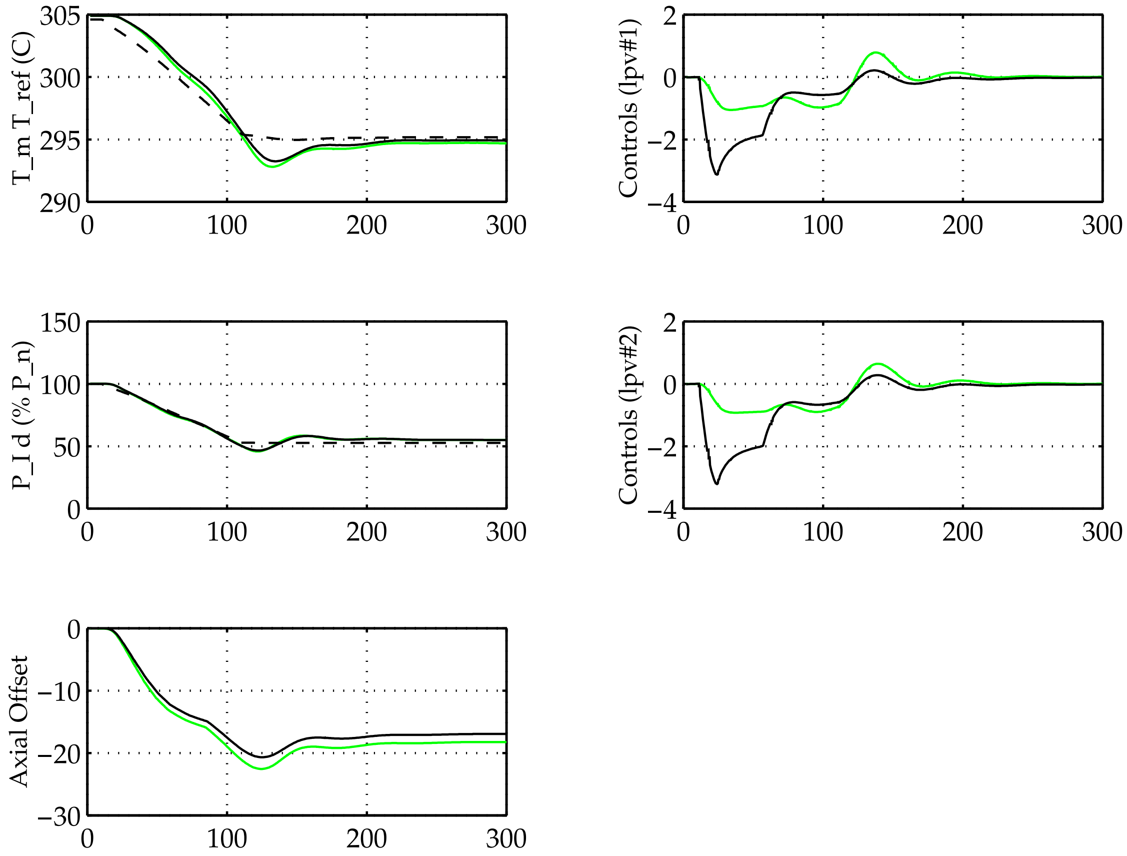

Figure 3.16: Comparison of LPV controllers on a nonlinear simulation to a large transient.

3.7 Summary

In this chapter, we have constructed an LPV controller for a nuclear power plant which performs very well, although not ideally, over the plant's operating range. This involved the construction of a parameter-dependent model from identified models, and the development of a parameter-dependent weight on its performance. The LPV controller is able to do much better over the entire operating range of the power plant than a single linear design, but does not switch strategies in its use of the control rods from low to high power, a behavior we would prefer.

Chapter 4

Application: Ducted Fan

They make mad the roaring winds,
And with tempests play.
—William Blake

This chapter introduces the second application discussed in this dissertation, a thrust-vectored ducted fan engine. The model for the ducted fan is a nonlinear first principles model, which will be presented. We detail the construction of the parameter-varying model from linearizations obtained around different operating points, and describe the synthesis of $\mathcal{H}_\infty$ and LPV controllers. The next chapter will present a detailed evaluation of these controllers in comparison with others. The general intent of this chapter is to present control designs for a system with strongly nonlinear behavior. Additionally, we demonstrate that first principles models, which in some respects are limited, can be used quite effectively in control design. The work in this chapter arose from a collaborative effort with Pascale Bendotti and Michael Kantner [27, 14].

In Section 4.1, the control problem is introduced and motivated. In Section 4.2, we describe the configuration of the ducted fan and discuss its dynamics. In Section 4.3, the derivation of the parameter-varying model is presented. Sections 4.4 and 4.5 discuss the synthesis of the $\mathcal{H}_\infty$ and LPV designs.

4.1 Introduction

Motivation

The construction of a flight control experiment at Caltech was motivated by a desire to investigate the application of linear and nonlinear control techniques to high-performance aircraft performing aggressive maneuvers. By focusing on a specific system, we hope both to generate new ideas and to investigate new techniques in control for dealing with systems having a strong nonlinear behavior. Very few design methods are available for building robust, nonlinear control laws for this class of systems. Additionally, we want to evaluate how well currently available methods and tools work on a real system with obvious nonlinearities. In this sense, the ducted fan is a testbed to answer the question

"What limits the applicability of current methods of control design to this type of system?"

There is a large body of literature on vectored propulsion systems which are gaining popularity as a method of improving the performance capabilities of modern jet aircraft. The fundamental concepts in vectored propulsion are described in the book by Gal-Or [24] (see also the survey article [23]). Most of the existing literature and experiments concentrate on control of full-scale jet engines and are primarily concerned with extending the flight envelope by extending existing (linear) control methodologies. An experiment similar to this ducted fan has been constructed by Hauser at the University of Colorado, Boulder [28].

Although we are aware of LPV control being applied to examples and simulations [4, 3, 2, 13], this is the first application of these techniques to a real physical example. Our controller performance can be compared with other linear, nonlinear, and gain scheduled controllers previously designed for the ducted fan in [27].

4.2 Description of the Fan Engine

Hardware

A picture of the experimental system, a thrust-vectored ducted fan engine, is shown in Figure 4.1. It consists of a high-efficiency electric motor with a 6-inch diameter blade, capable of generating up to 9 Newtons of thrust. In Choi *et al.* [17], a detailed description of the performance of the fan was given, including models for the thrust as a function of flap angle and fan speed, as well as some discussion of ground effects.

Overall, the experimental setup consists of the ducted fan attached to a three degree of freedom stand, as shown in Figure 4.2. The different thrust modes available are shown in Figure 4.3. The intent of the design was to have a simple ducted fan aircraft which could provide two dimensional vectored and reverse thrust. The aircraft is bolted to a rotating arm, which limits its motion to three degrees of freedom: one rotational and two translational, approximately on the surface of a sphere defined by the arm. With this geometry, the ducted fan is completely controllable with just the vectored thrust. A detailed discussion of the components is available elsewhere [17].

The aircraft is composed of a variable speed electric motor which drives a four-blade propeller. The motor and propeller assembly are bolted inside a wooden duct which has two flaps attached at the end. The pitch stability of the fan is configurable and can be changed from stable to unstable. For these experiments, the ducted fan was in a stable configuration. An optical encoder with an angular resolution of $\pi/1000$ radians is mounted on each axis.

Details of the software interface can be found in Appendix 4.B.

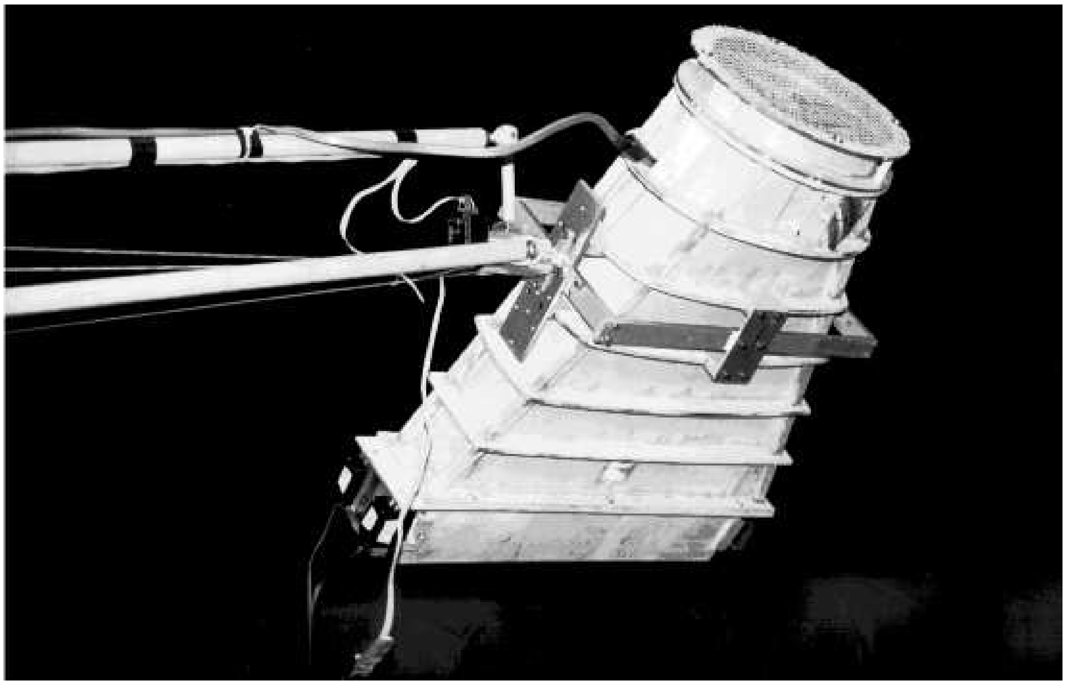

Figure 4.1: Ducted fan apparatus.

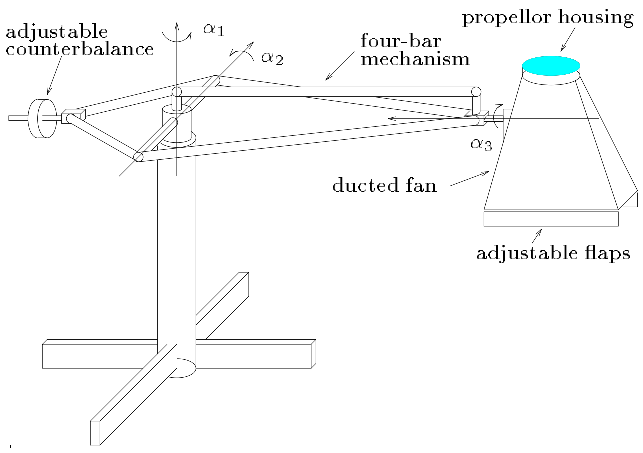

Figure 4.2: Ducted fan attached to stand.

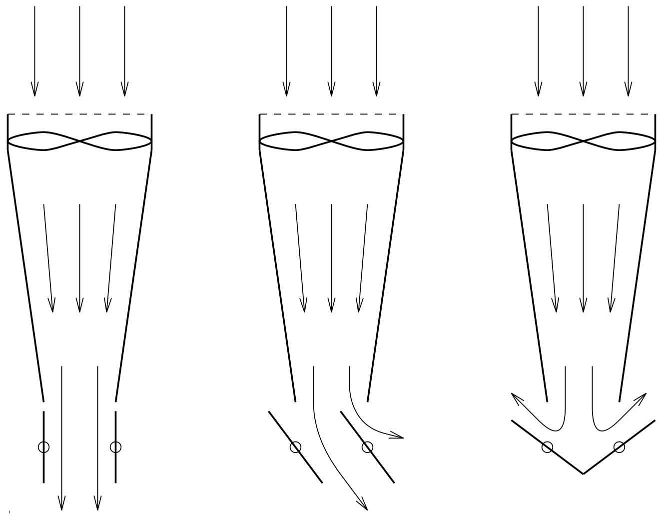

Figure 4.3: Different thrust modes for the ducted fan.

4.3 Modelling

All controllers are designed using a first principles model of the ducted fan based on standard rigid body mechanics. The state consists of the angles, α_i, and their velocities, $\dot{\alpha}_i$, $i = 1, 2, 3$. The equations of motion for the system, derived from Lagrange's equations, have the functional form

$$M(\alpha)\alpha + C(\alpha, \dot{\alpha}) + N(\alpha) = \Upsilon(\alpha, f_1, f_2),$$

where α denotes the vector of angles, $M(\alpha)$ is the generalized inertia matrix, $C(\alpha, \dot{\alpha})$ is the Coriolis matrix, $N(\alpha)$ is the matrix of gravity terms, and $\Upsilon(\alpha, f_1, f_2)$ is the matrix of applied joint torques. The ducted fan equations of motion are derived in Appendix 4.A.

The model is accurate enough for control design, although it does have limitations. Identification experiments show it is reasonably accurate near hover (*cf*. Section 6.1). Initial step responses on single axes compared favorably with experimental measurements, and a PID control test gave expected results [17]. A decoupling controller, essentially a plant inversion, worked well [27]. Nonetheless, the model omits many effects: all actuator dynamics, sensor limitations, friction, and aerodynamic effects. Static friction about the α_1 axis is significant. Aerodynamic effects have been observed in the lab during forward flight. The model also omits the gyroscopic terms that result from the angular momentum of the propeller. This term, unfortunately, is significant.

The model also assumes that the commanded forces act at a fixed point on the fan. Experiments have shown that the distance from the fan's center of mass to the point at which the force acts, r, varies as the flap angle changes, by as much as 20%. Furthermore, motor speed and flap angle, not forces, are commanded. An experimentally determined lookup table maps desired forces to motor speed and flap angle. This table is assumed to be constant, although the actual mapping varies according to the operating conditions.

Perhaps the most significant nonlinear features of the ducted fan are the effect of the pitch angle α_3 on the direction of the input forces, and the centrifugal forces. The centrifugal forces can be quite high when the fan flies rapidly, and will tend to push α_2 to a particular value.

An examination of the nonlinear model reveals that the most significant variations in parameters occur as a function of α_3, $\dot{\alpha}_1$, and α_2. The dependence on α_2 is complicated and for the trajectories we will consider it is less than that of α_3 and $\dot{\alpha}_1$, so we neglect it. The fan is strictly proper and thus the D matrix of the state space model is zero. Moreover, although the rates are not measured, an inner loop in the real-time software controller estimates them; hence in our models $C = I$. The A and B matrices are the only matrices which have parameter variations. Their structure is

$$
\left[\, A \,\middle|\, B \,\right] =
\left[\begin{array}{cccccc|cc}
1 & 0 & 0 & T & 0 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 & T & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & T & 0 & 0 \\
0 & a_{42}(\alpha_3, \dot{\alpha}_1) & a_{43}(\alpha_3) & 1 & 0 & 0 & b_{41}(\alpha_3) & b_{42}(\alpha_3) \\
0 & a_{52}(\dot{\alpha}_1) & a_{53}(\alpha_3) & 0 & 1 & 0 & b_{51}(\alpha_3) & b_{52}(\alpha_3) \\
0 & a_{62}(\alpha_3, \dot{\alpha}_1) & a_{63}(\alpha_3) & 0 & 0 & 1 & b_{61}(\alpha_3) & b_{62}(\alpha_3)
\end{array}\right] \tag{4.1}
$$

where T is the sampling rate.

Figures 4.4 and 4.5 show the parameter dependence of one term in the A matrix and all terms in the B matrix, important for the simple control model we will be employing. The remaining variations are shown in Appendix 4.C, in Figures 4.12 through 4.16. The dependence is obtained by linearizing the nonlinear model at various equilibrium operating points for different values of α_3 and $\dot{\alpha}_1$. Note the dependence of a_{42} and a_{62} on both α_3 and $\dot{\alpha}_1$, while a_{52} depends only on $\dot{\alpha}_1$; a_{43}, a_{53}, a_{63} and all terms of the B matrix depend only on α_3. In Figure 4.5, the actual values are shown as asterisks, and the least-squares fit described below is the solid line.

To derive the parameter dependence and fit it into this framework, each of the parameters mentioned above is fit with a rational function of first or second order using a least-squares technique. Immediately some approximations are made. The dependence of a_{42} and a_{62} on $\dot{\alpha}_1$ was neglected, making them depend only on α_3. The parameters a_{42}, a_{62}, b_{41}, b_{52}, and b_{61} were approximated as lines, i.e., first-order LFTs. The rest were approximated as second order LFTs. Assume in the following that $\delta_0 = z^{-1}$, $\delta_1 = \dot{\alpha}_1$, and $\delta_2 = \alpha_3$, and let

$$
\boldsymbol{\Delta} = \{\mathrm{diag}\,[\delta_0 I_{n_0}, \delta_1 I_{n_1}, \delta_2 I_{n_2}]\}\,.
$$

The resulting model with δ-dependence, $P(\delta)$, is shown in Figure 2.4, where n_0 is the size of the block corresponding to δ_0.

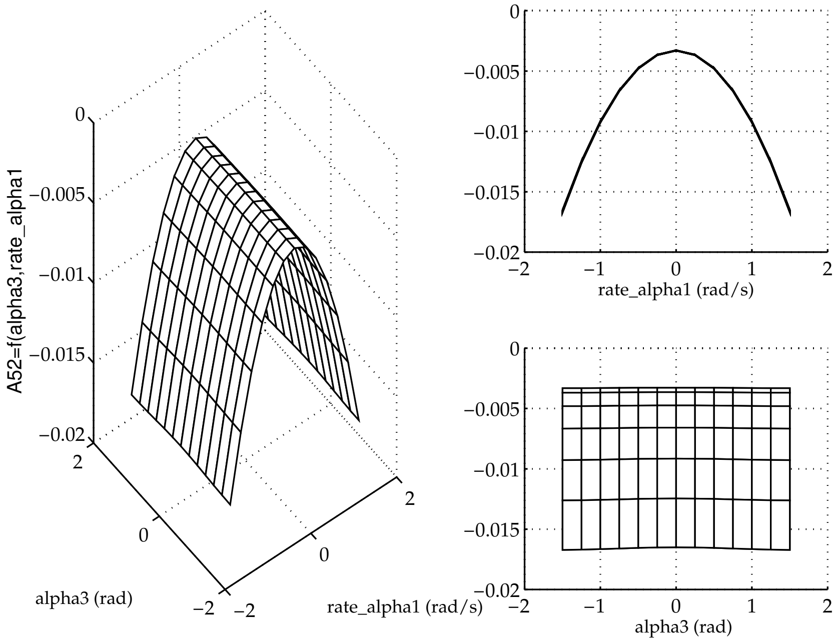

Figure 4.4: Dependence of a_{52} upon α_3 and $\dot{\alpha}_1$.

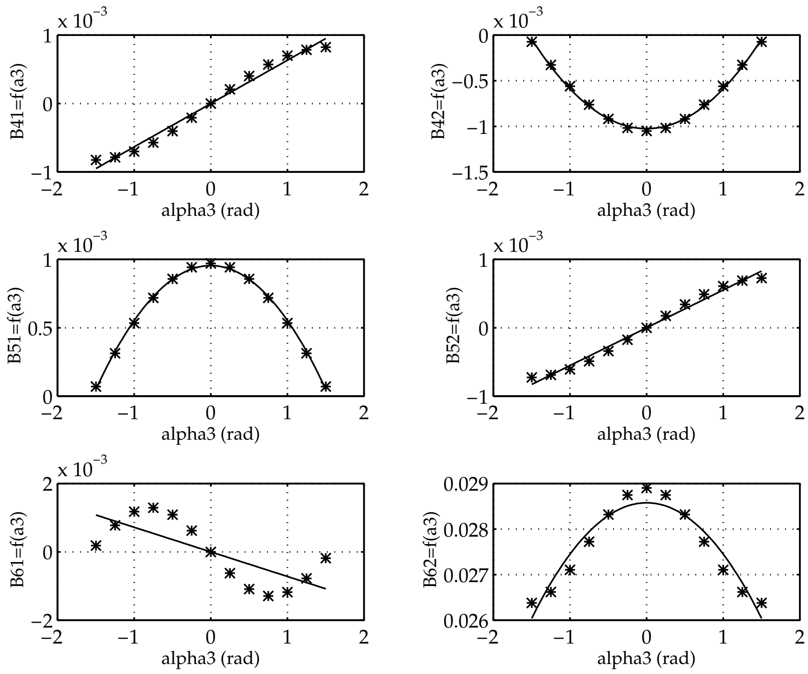

Figure 4.5: Dependence of the entries of the B matrix upon α_3.

Having done all this, an even simpler model is used for synthesis. The synthesis model considers only variations in the cross-coupling terms of the B matrix, i.e., b_{41}, b_{52}, and b_{61}, and the variation of a_{52}. The other terms are assumed to be constant with the values they take at hover. Thus $n_0 = 6$, $n_1 = 2$, and $n_2 = 2$ ($n_2 \neq 3$ because the fact that all three are lines can be exploited to reduce the size). Moreover, the range of α_3 is assumed to be from 0 radians to 1.5 radians, and the range of $\dot{\alpha}_1$ is assumed to be from 0 rad/s to 1.5 rad/s. We refer to this model as the "simplified pitch-velocity model," denoted by the pair $(P_{des}, \mathbb{\Delta}_{des})$. It is the model used in our control designs. The reasons for using this particular simplification are based primarily on intuition regarding how the ducted fan works. We will discuss these issues further in Section 6.3. There, we will see that conservatism involved in moving into the operator LPV framework affects us in an important way: we lose stabilizability in the model.

4.4 $\mathcal{H}_\infty$ Controller Synthesis

Both LPV and $\mathcal{H}_\infty$ synthesis produce controllers which reject disturbances. A tracking problem such as the ducted fan can be cast in this framework by rejecting the low frequency components of the error between the plant output and the reference. The tracking will become faster as higher frequencies are rejected. A singular value Bode plot of the linearized model is shown in dashed lines in Figure 4.6. The other lines are weights which will be described later. The dashed lines show the frequency responses of the primary gains of interest, $u_1 \to \alpha_2$ and $u_2 \to (\alpha_1, \alpha_3)$. The α_2 response is the response with an undamped mode at approximately 0.9 rad/s. This is the "pendulum mode" of the ducted fan caused by the bar connecting counterweight and fan hanging slightly below the pivot point (see Appendix 4.A). The response of α_1 has double integrator slope at low frequencies and a mode at 3.5 rad/s, the same frequency as the mode for α_3. This mode is the "rocking mode" of the fan as it rotates about the α_3 axis.

The synthesis structure used is shown in Figure 4.7, with uncertainty and performance weights included; u, y, r, and n are the controls, the measurements, the reference signals the controller must track, and sensor noises, respectively. W_u is a penalty on the control, and W_m is multiplicative actuator uncertainty. W_p is a diagonal performance weight on the signals α_1, α_2, α_3, and $\dot{\alpha}_1$. As noted in Section 2.3, $\mathcal{H}_\infty$ synthesis deals with one full block. Implicitly, there are two blocks in Figure 4.7 — one for the uncertainty at the input and one for the performance requirements. The figure is shown the way it is because we understand that for $\mathcal{H}_\infty$ synthesis all individual blocks are collected into a diagonal structure and that structure covered with one full block. Drawing the diagram the way we do reminds us of this conservatism. Note that LPV synthesis is also a "one block methodology."

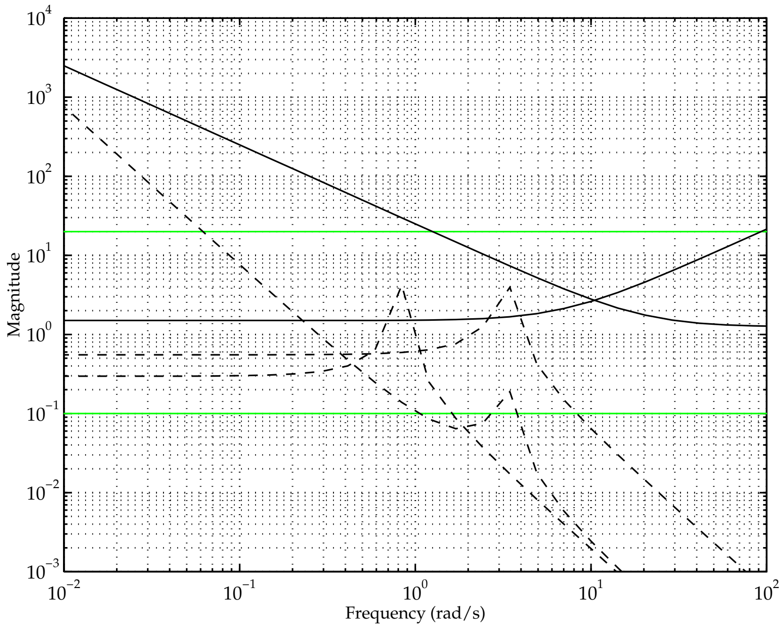

Figure 4.6: Performance and uncertainty weights for the $\mathcal{H}_\infty$ controller. The solid black lines are the performance weights on α_2 and the multiplicative uncertainty weight W_m. The shaded lines are performance weights on α_1 and α_3. The dashed lines are a Bode singular value plot of the linearized model.

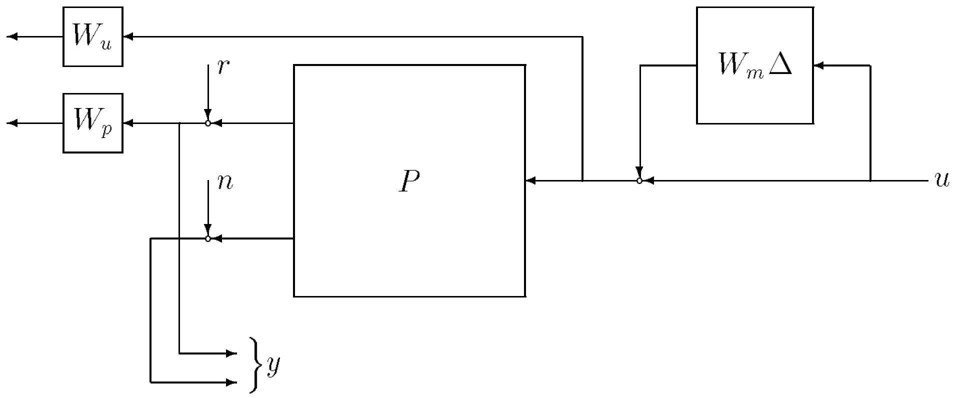

Figure 4.7: Synthesis structure used for designing the $\mathcal{H}_\infty$ and LPV controllers. See the text for an explanation of how this diagram fits into the $\mathcal{H}_\infty$ framework.

A multiplicative uncertainty structure is used at the input to the plant. This has the effect of causing the controller to roll off at high frequency. We shall see in Section 6.2 that in reality this uncertainty structure is inadequate to describe behavior the real system exhibits. W_m is the frequency-dependent weight on this uncertainty. It resembles a high pass filter and is shown in solid lines in Figure 4.6. The weight is necessary when high performance is required. That is, requiring high tracking performance produces a controller which destabilizes the system unless a frequency dependent weight is used. This adds two states to our controller.

Turning now to performance specifications, W_u is a penalty on the control action. This is placed there to keep the control signals physically realizable. Our specifications have no requirement on α_3, but some small performance weight on this output was found helpful. To provide some damping on the mode without sacrificing performance on α_1 or α_2, the weight is set to a small constant value of 0.1. A constant weight is sufficient on α_1 since it has high gain in the low frequency. A similar weight was placed on $\dot{\alpha}_1$. Notice that the gain of the α_2 channel is small. Then our controllers must have a large gain on α_2 to give good performance, and so an integral-like weight is used. This is shown in solid lines as a low pass filter in Figure 4.6, and adds one state to the controller. This performance weight seeks to make the closed loop dynamics of α_2 faster than those of α_1, in an effort to reduce the α_2 error caused when the fan is tilted at an angle.

The performance of the controller synthesized using these weights will be presented in Chapter 5.

4.5 LPV Controller Design

The LPV design process builds on the $\mathcal{H}_\infty$ design of the previous section. In particular, we use the synthesis structure shown in Figure 4.7. The basic procedure is to iterate upon the weights obtained from the $\mathcal{H}_\infty$ design to refine the LPV design. Because of the different assumptions for the underlying perturbations, our experience is that identical weights don't quite work. Nonetheless, they serve as a good starting point.

LTI weights were used; as we saw in Chapter 3, using weights depending on parts of the Δ-block can enhance performance [13], but were not needed here. The weights for the LPV design are shown in Figure 4.8. The dashed lines show the maximum singular values of the three main channels in the linearized model. The solid black line resembling a low pass filter is the performance weight on α_2, while the solid black line resembling a high pass filter is the multiplicative uncertainty weight W_m. The shaded lines are performance weights on α_1 (large) and α_3 (small), respectively.

Comparing these weights to the weights for the $\mathcal{H}_\infty$ design, they are basically identical in shape, but the values are quite different. The pole in W_m is

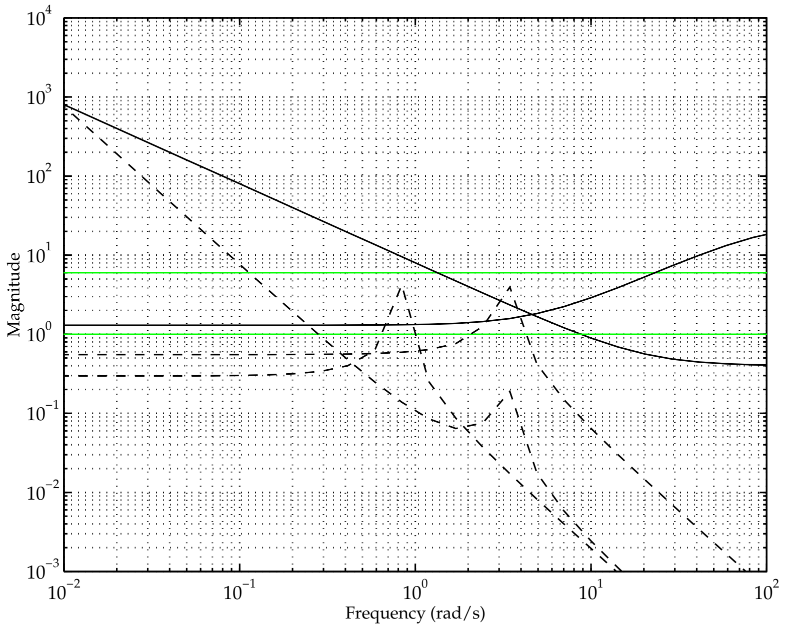

Figure 4.8: Bode plot of the ducted fan model linearized about hover (dashed lines) with weights for LPV synthesis (solid and shaded lines).

at a higher frequency, so the controller is forced to roll off more. The performance weight on α_2 is not as large, but there is a higher performance weight on α_3. The performance weight on α_1 is not as large either. An $\mathcal{H}_\infty$ controller designed with these weights would be quite sluggish, but the LPV synthesis technique produces more aggressive controllers for the same weights.

A natural question arises as to why we employ the full LPV technique to a state-feedback problem. Recall that we needed to add at least three states to the open-loop interconnection structure for the $\mathcal{H}_\infty$ control. Once W_m has states, full information for the open-loop interconnection is lost. By experimenting with different designs, we were not able to design an adequate controller using a constant weight for the multiplicative uncertainty, so we are unable to exploit the measurements of the state in the plant.

The performance of the LPV controller designed using this procedure is presented in Chapter 5.

4.6 Summary

In this chapter, we presented the ducted fan, and derived a first principles model for use in control design. From this first principles model, which is nonlinear, we constructed a parameter-varying model based on linearizations of the nonlinear model around various operating points. The parameter-varying model was then simplified to render it tractable. Based on a linearization of the model around hover, the design of an $\mathcal{H}_\infty$ controller was presented. Using this as a starting point, the design of an LPV controller proceeded by iterating upon the $\mathcal{H}_\infty$ weights as a starting point.

It is important to keep in mind an overview of what this process involves. A nonlinear model is taken, and an LPV model is constructed based on real parameter variations of several linearizations. When this model is applied to LPV routines, the underlying connection with the nonlinear model is essentially lost. What were real parameter variations now are assumed to be time-varying operators with memory. Thus, the technique is theoretically conservative.

Future work on LPV designs will involve improving the model to include more of the parameter variations than in the designs presented here.

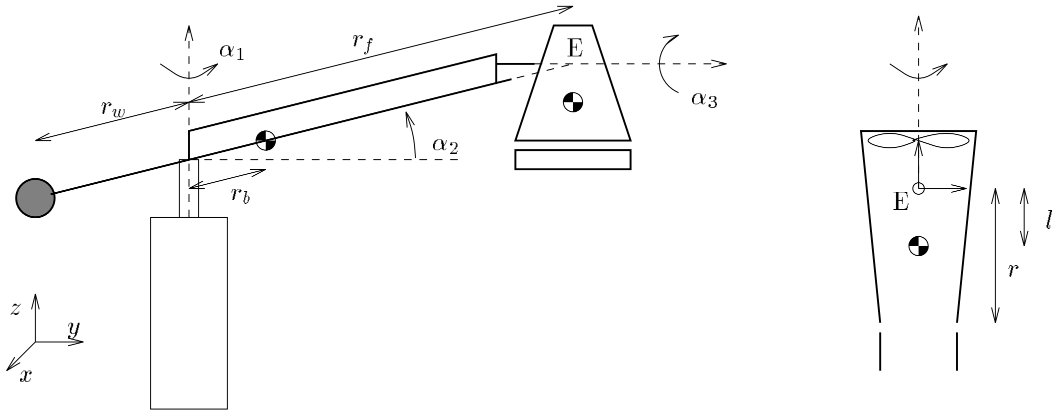

Figure 4.9: The ducted fan and stand model. The checkered circles represent the centers of mass of the bar and fan. The distance h the bar hangs below the pivot point for α_2 is not shown.

4.A Appendix: Ducted Fan Equations of Motion

In this appendix, we present a derivation of the nonlinear model of the ducted fan. While this is not exactly the model used for deriving the parameter-varying model presented in this chapter, it is very close, and captures all the details of that model. A Lagrangian analysis will be used to derive the equations of motion. We will assume a basic knowledge of kinematics. Background material for this appendix can be found in Chapter 4 of [34].

The basic geometry of the fan is shown in Figure 4.9. The standard angles α_1, α_2, and α_3 are used. Let r_w be the distance of the counterweight and m_w the mass of the counterweight. The four-bar mechanism which holds the axis of rotation for α_3 parallel to the y-axis is modelled as a slender bar with mass m_b and distance to center of mass r_b. The point of rotation for the four-bar mechanism is not at the origin of the coordinate system as is shown in the figure, but is on the z axis slightly below the origin. We denote this distance by h. Let r_f be the distance from the stand's vertical (z) axis to the point of rotation of α_3 on the fan (E), and let m_f be the mass of the fan. The distance from this point of rotation to the center of the mass of the fan is denoted by l.

There are two external forces which act on the fan, denoted by f_1 and f_2. The force f_1 acts parallel to the main axis of the fan, and f_2 perpendicular. The moment arm from the point of rotation of the fan to the point at which f_2 acts is denoted by r.

Recall that the kinetic energy T of a rigid body can be expressed as the sum of a translational and rotational component,

$$T = \frac{1}{2}m\|v\|^2 + \frac{1}{2}\omega^T \mathcal{I}\omega,$$

Symbol	Description	Value	
r_f	distance of fan center of mass from α_2 axis	143.5	cm
r_b	offset of center of mass of bar	30.0	cm
r_w	distance of counterweight from α_2 axis	-56.8	cm
l	fan center of mass offset	2.13	cm
h	bar offset from α_2 axis	4.14	cm
r	nominal distance of flaps from fan pivot	26.0	cm
m_f	mass of fan	2.22	kg
m_b	mass of bar	2.87	kg
m_w	counterweight mass	6.43	kg
I^b	bar moment of inertia, α_2 axis	0.75	kg m^2
I^f_{xx}	fan moment of inertia, α_1 axis	1.73	kg m^2
I^f_{zz}	fan moment of inertia, α_3 axis	0.05	kg m^2
g	gravitational constant	9.8	m/s^2

Table 4.1: Physical Constants for the ducted fan.

where v is the velocity at the center of mass, ω is the body angular velocity, and $\mathfrak{I}$ is the inertia tensor of the body. We will assume the inertia tensor for the bar is diagonal and equal to

$$\mathfrak{I}_b = \begin{bmatrix} I^b_{xx} & & \\ & I^b_{yy} & \\ & & 0 \end{bmatrix},$$

where the last term is 0 because we assume the bar is a slender rod. We assume that $I^b_{xx} = I^b_{yy} = I^b$. We further assume that the inertia tensor for the fan is diagonal. The fan is not symmetric with respect to two axes, so this is not true, but we assume the cross terms are negligible. Then the inertia tensor for the fan is

$$\mathfrak{I}_f = \begin{bmatrix} I^f_{xx} & & \\ & I^f_{yy} & \\ & & I^f_{zz} \end{bmatrix}.$$

The values of all physical constants for the ducted fan with a description are shown in Table 4.1.

For our analysis, we will pick the angles α_1, α_2, and α_3 as our generalized coordinates. Let α denote the vector of these angles. The Lagrangian, L, is defined as the difference between the kinetic and the potential energy of the system. Letting T and V denote the kinetic and potential energies of the system, respectively, we have

$$L(\alpha, \dot{\alpha}) = T(\alpha, \dot{\alpha}) - V(\alpha).$$

Lagrange's Theorem for the equations of motions is now given without proof (see the references cited in [34] for a proof).

Theorem 2 *The equations of motion for a mechanical system with generalized coordinates $\alpha \in \mathbb{R}^m$ and Lagrangian L are given by*

$$\frac{d}{dt}\frac{\partial L}{\partial \dot{\alpha}_i} - \frac{\partial L}{\partial \alpha_i} = \Upsilon_i,$$

where $i = 1, \ldots, m$ and Υ_i is the external force acting on the i^{th} generalized coordinate.

Thus to determine the equations of motion for the system, we only need to calculate the potential and kinetic energy of the fan in terms of α. The potential energy of the system is the sum of the potential energies of the counterweight, bar, and fan, as follows:

$$\begin{aligned} V_{\text{system}} &= V_w + V_b + V_f \\ &= g\left(m_b r_b + m_w r_w + m_f r_f\right)\sin\alpha_2 - \\ &\quad\; g\left(m_b + m_w + m_f\right) h \cos\alpha_2 - gl\cos\alpha_3. \end{aligned}$$

Likewise, the kinetic energy of the fan is

$$T_{\text{system}} = T_w + T_b + T_f.$$

In computing the kinetic energy terms, we will assume that the fan rotates through the origin, i.e., $h = 0$. There are some kinetic energy affects associated with h, but they are small and we have considered its major affect — on the potential energy. Of these terms, T_f is the most complicated, so we will present its derivation in detail. From above,

$$T_f = \frac{1}{2}m_f\|v_f\|^2 + \frac{1}{2}\omega^T \mathfrak{I}_f \omega.$$

The position of the point of rotation, p_E, is given by

$$p_E(\alpha_1, \alpha_2) = \begin{bmatrix} r_f\cos\alpha_2\cos\alpha_1 \\ r_f\cos\alpha_2\sin\alpha_1 \\ r_f\sin\alpha_2 \end{bmatrix}.$$

Hence, the position of the center of mass of the fan, p_f, is

$$p_f(\alpha_1, \alpha_2, \alpha_3) = \begin{bmatrix} r_f\cos\alpha_2\cos\alpha_1 + l\sin\alpha_3\sin\alpha_1 \\ r_f\cos\alpha_2\sin\alpha_1 - l\sin\alpha_3\cos\alpha_1 \\ r_f\sin\alpha_2 - l\cos\alpha_3 \end{bmatrix}.$$

Then v_f becomes

$$v_f(\alpha, \dot{\alpha}) = \dot{p}_f = \begin{bmatrix} -\dot{\alpha}_2 r_f\sin\alpha_2\cos\alpha_1 - \dot{\alpha}_1 r_f\cos\alpha_2\sin\alpha_1 + \\ \\ \dot{\alpha}_3 l\cos\alpha_3\sin\alpha_1 + \dot{\alpha}_1 l\sin\alpha_3\cos\alpha_1 \\ -\dot{\alpha}_2 r_f\sin\alpha_2\sin\alpha_1 + \dot{\alpha}_1 r_f\cos\alpha_2\cos\alpha_1 - \\ \\ \dot{\alpha}_3 l\cos\alpha_3\cos\alpha_1 + \dot{\alpha}_1 l\sin\alpha_3\sin\alpha_1 \\ \dot{\alpha}_2 r_f cos\alpha_2 + l\dot{\alpha}_3\sin\alpha_3 \end{bmatrix}.$$

Because of the four-bar mechanism, the fan has no angular component depending on α_2. Thus ω_f is given by

$$\omega_f = \begin{bmatrix} \dot{\alpha}_1 \\ 0 \\ \dot{\alpha}_3 \end{bmatrix}$$

and so

$$\omega_f^T \mathfrak{I}_f \omega = \dot{\alpha}_1^2 I_{xx}^f + \dot{\alpha}_3^2 I_{zz}^f.$$

This represents all the terms in the kinetic energy of the fan.

The kinetic energy of the counterweight is simply

$$T_w = \frac{1}{2} m_w r_w^2 \left(\dot{\alpha}_2^2 + \dot{\alpha}_1^2 \cos^2 \alpha_2 \right).$$

The kinetic energy of the bar is similar to this plus a rotational term. The angular velocity vector for the bar is

$$\omega_b = \begin{bmatrix} -\dot{\alpha}_2 \sin \alpha_1 \\ \dot{\alpha}_2 \cos \alpha_1 \\ \dot{\alpha}_1 \end{bmatrix}$$

and hence the kinetic energy for the bar becomes

$$T_b = \frac{1}{2} \dot{\alpha}_2^2 \left(m_b r_b^2 + I^b \right) + \frac{1}{2} \dot{\alpha}_1^2 m_b r_b^2 \cos^2 \alpha_2.$$

Finally, the external forces enter into the equations of motion as a function of α, f_1, and f_2. They are given as follows:

$$\Upsilon(\alpha, f_1, f_2) = \left[\begin{array}{c} r_f \cos \alpha_2 \begin{pmatrix} \sin \alpha_3 & -\cos \alpha_3 \\ \cos \alpha_3 & \sin \alpha_3 \end{pmatrix} \\ 0 \quad r \end{array} \right] \begin{bmatrix} f_1 \\ f_2 \end{bmatrix}.$$

We treat r as a constant, but experimental measurements show it can vary by 20% during fan operation.

Summing all the kinetic energy terms and subtracting from them the potential energy terms gives us the Lagrangian. After a detailed calculation using Theorem 2 we obtain the equations of motions for the fan, which we write in the form

$$M(\alpha)\alpha + C(\alpha, \dot{\alpha}) + N(\alpha) = \Upsilon(\alpha, f_1, f_2).$$

$M(\alpha)$ is called the generalized inertia matrix, $C(\alpha, \dot{\alpha})$ is the matrix of Coriolis terms, and $N(\alpha)$ is the matrix of gravity terms. Letting

$$J = m_w r_w^2 + m_b r_b^2 + m_f r_f^2,$$

these matrices are:

$$M(\alpha) = \begin{bmatrix} \begin{pmatrix} J\cos^2\alpha_2 + I_{xx}^f + \\ m_f l^2 \sin^2\alpha_3 \end{pmatrix} & -m_f r_f l \sin\alpha_2 \sin\alpha_3 & -m_f r_f l \cos\alpha_2 \cos\alpha_3 \\ -m_f r_f l \sin\alpha_2 \sin\alpha_3 & J + I^b & m_f r_f l \cos\alpha_2 \sin\alpha_3 \\ -m_f r_f l \cos\alpha_3 \cos\alpha_2 & m_f r_f l \cos\alpha_2 \sin\alpha_3 & I_{zz}^f + m_f l^2 \end{bmatrix},$$

$$C(\alpha, \dot\alpha) = \begin{bmatrix} -2J\dot\alpha_1\dot\alpha_2 \cos\alpha_2 \sin\alpha_2 + 2m_f l^2 \dot\alpha_1 \dot\alpha_3 \cos\alpha_3 \sin\alpha_3 + \\ m_f r_f l \left(\dot\alpha_3^2 - \dot\alpha_2^2\right) \cos\alpha_2 \sin\alpha_3 \\ \\ J\dot\alpha_1^2 \cos\alpha_2 \sin\alpha_2 + m_f r_f l \dot\alpha_3^2 \cos\alpha_2 \cos\alpha_3 - \\ 2m_f r_f l \dot\alpha_1 \dot\alpha_3 \sin\alpha_2 \cos\alpha_3 \\ \\ -m_f l^2 \dot\alpha_1^2 \sin\alpha_3 \cos\alpha_3 - m_f r_f l \dot\alpha_2^2 \sin\alpha_2 \sin\alpha_3 + \\ 2m_f r_f l \dot\alpha_1 \dot\alpha_2 \sin\alpha_2 \cos\alpha_3 \end{bmatrix},$$

and

$$N(\alpha) = \begin{bmatrix} 0 \\ \\ g\cos\alpha_2 \left(m_b r_b + m_w r_w + m_f r_f\right) + \\ h \sin\alpha_2 \left(m_b + m_w + m_f\right) \\ \\ g m_f l \sin\alpha_3 \end{bmatrix}.$$

Note that the inertia matrix is not constant. The inertia of the system changes as a function of altitude and pitch angle, since the counterweight and fan move with respect to the vertical axis as a function of these angles.

This model does not include the effect of a non-zero h term on the the kinetic energy. Nor does it include the gyroscopic terms that result from the angular momentum of the whirling fan blade. The fan blade spins at approximately 10,000 rpm, so even though the blade itself is light, the Coriolis forces it generates can be significant.

4.B Appendix: Software Interface

The experiment is interfaced to an 80486 computer running an MS-DOS-based real-time kernel called Sparrow [35]. Custom hardware is used to read in joint angles via the encoders and generate PWM signals for radio controlled (R/C) servos which control the flap angles. Note that the servos are not controlled via radio, these are simply the type of servos common in the R/C industry. Currently, the joint angles are read in at 200Hz and the PWM signals are output at 50Hz, the standard update rate for R/C servos. A voltage-to-current amplifier powers the fan engine, and consequently controls the fan force.

Controllers are designed and simulated using MATLAB on Sun workstations. Sparrow loads linear, LPV, and some types of gain-scheduled controllers directly from MATLAB data files. Once a controller is designed, it can be tested immediately. Nonlinear controllers, implemented as MATLAB S-functions, require a small amount of revision before they are linked to Sparrow.

The LPV package of Sparrow supports arbitrary scaling and offsets of the δ parameters, and arbitrary Δ size. Thus, LPV controllers can be designed for a variety of plant parameter ranges and model parameterizations without requiring any software modifications. The basic structure for implementing LPV controllers is shown in Figure 4.10.

Since the LPV model is obtained by linearizing the nonlinear model about several equilibrium points, the state is measured, and based on this measurement the nominal forces required to maintain equilibrium at this operating point are fed forward continuously into the plant. The nominal offset forces were computed *a priori* and stored in a lookup table. The nominal forces are shown in Figure 4.11; Sparrow actually implements a flattened version of this which avoids saturating the motor at large values of α_3.

Since an LPV controller is an LFT on a Δ of measured parameters, implementing it requires a real-time matrix inversion. To reduce computations, though, the LFT can be eliminated using past values of Δ. This approximation allows faster sampling rates. As long as the δ-values do not change much between samples, the approximation is very accurate. All experiments in this paper were run using this approximation.

The software for implementing LPV controllers holds any δ at 0 if it becomes negative, and at 1.5 if it exceeds this.

4.C Appendix: Other Parameter Variations

In this appendix, the remaining parameter variations of the model of Equation 4.1 are shown in Figures 4.12 through 4.16.

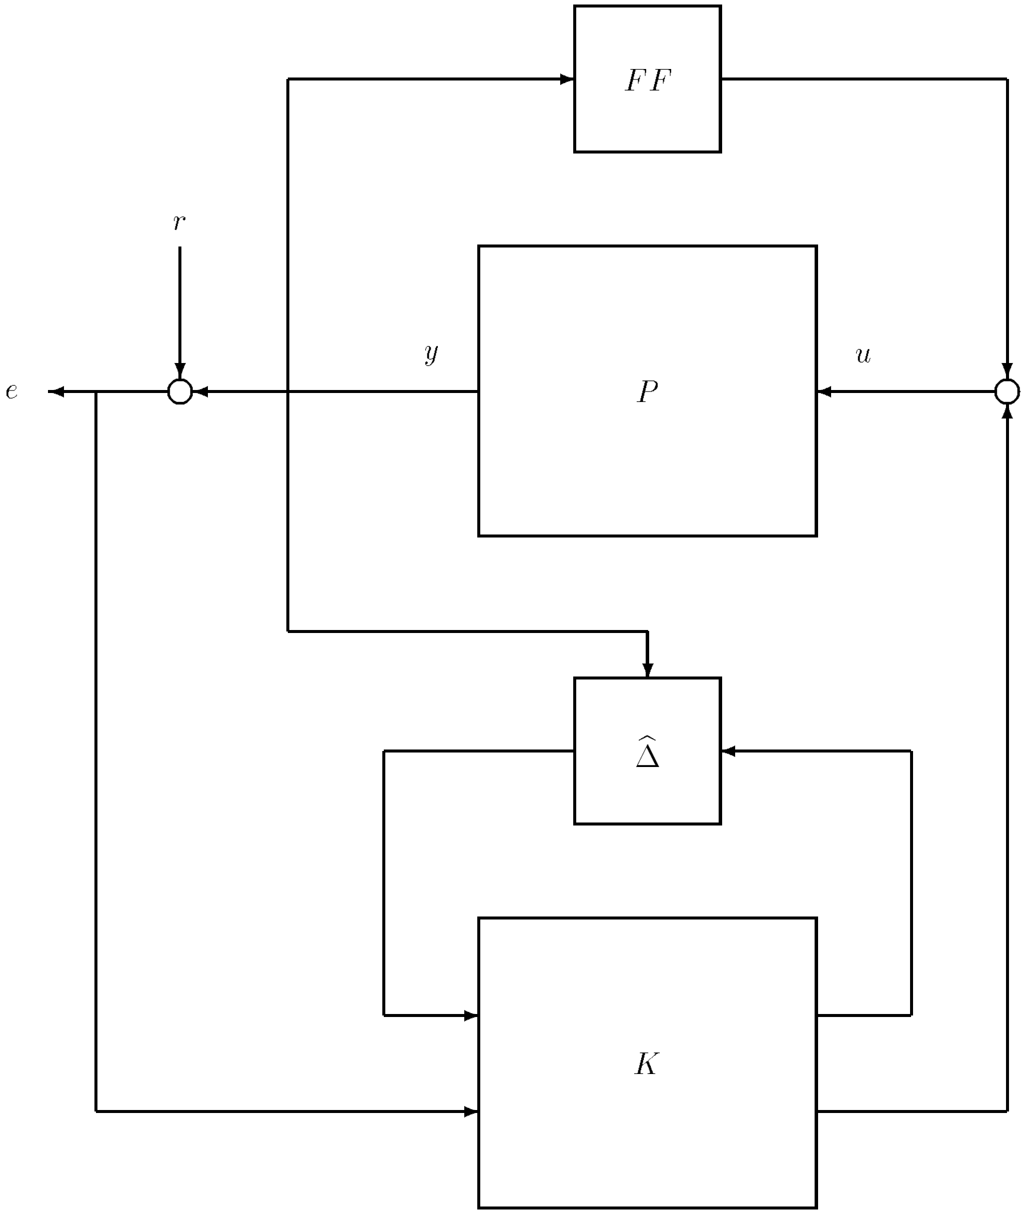

Figure 4.10: This block diagram shows how LPV controllers are implemented under Sparrow.

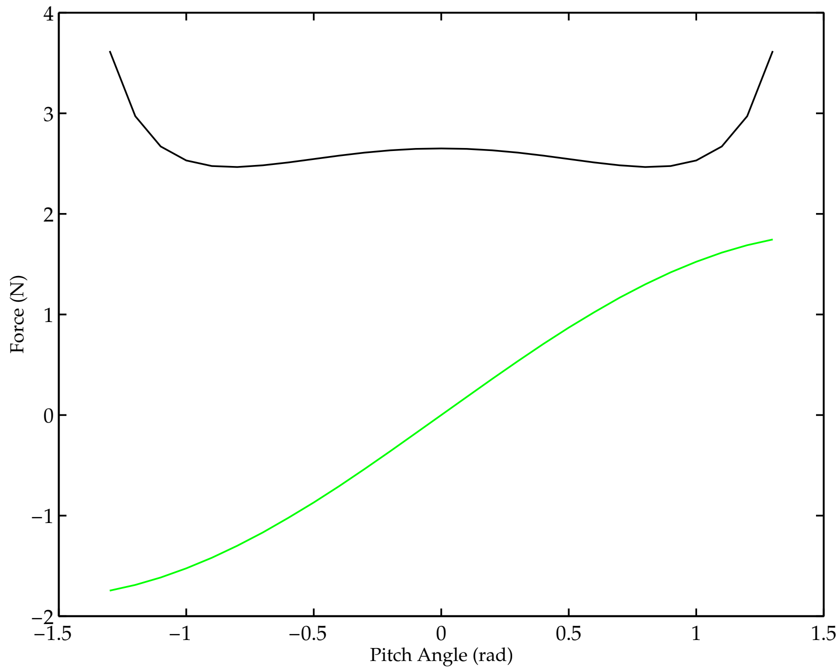

Figure 4.11: Nominal forces required to maintain equilibrium as a function of α_3. The solid line is the u_1 feedforward force and the shaded line the u_2 feedforward force.

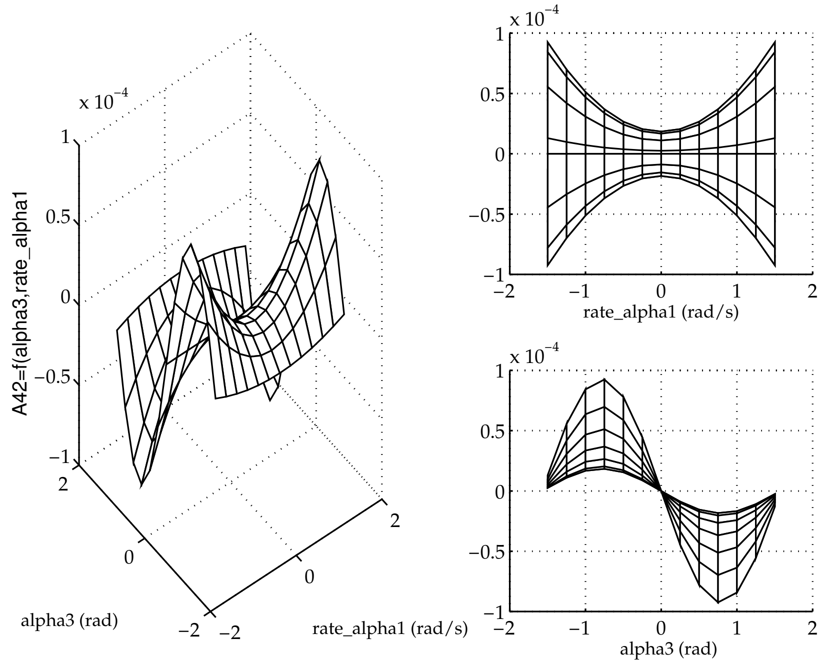

Figure 4.12: Dependence of a_{42} upon α_3 and $\dot{\alpha}_1$.

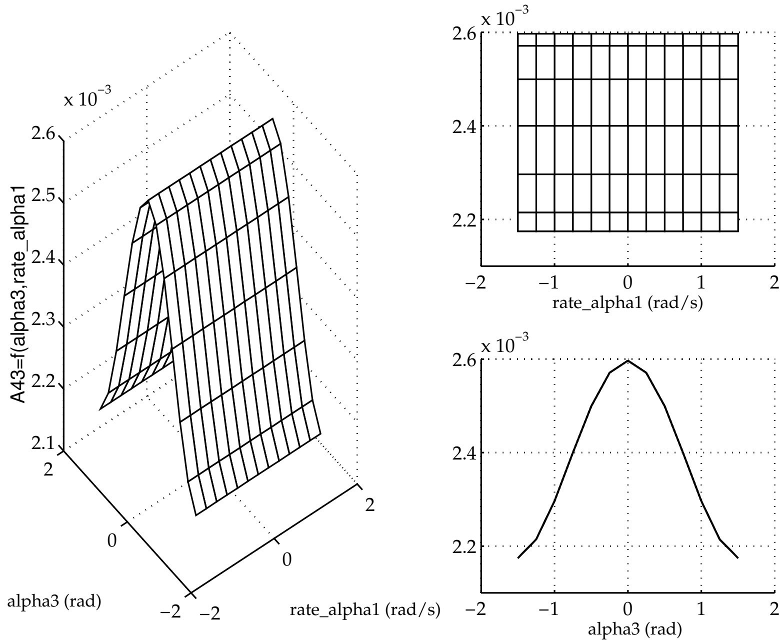

Figure 4.13: Dependence of a_{43} upon α_3 and $\dot{\alpha}_1$.

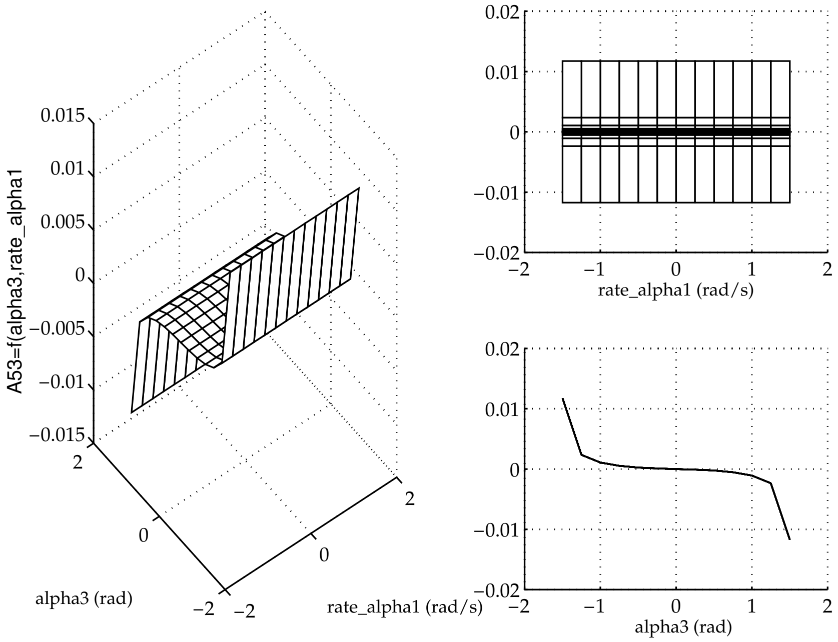

Figure 4.14: Dependence of a_{53} upon α_3 and $\dot{\alpha}_1$.

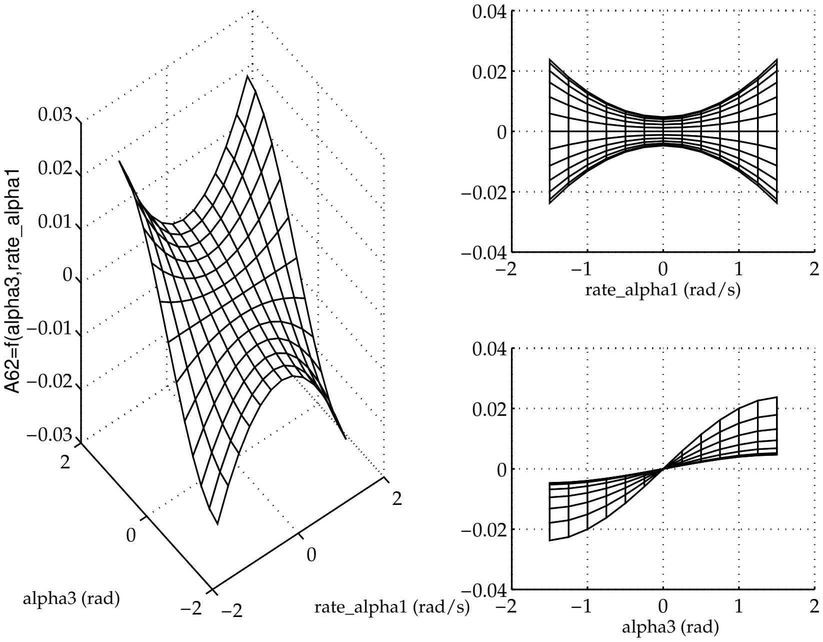

Figure 4.15: Dependence of a_{62} upon α_3 and $\dot{\alpha}_1$.

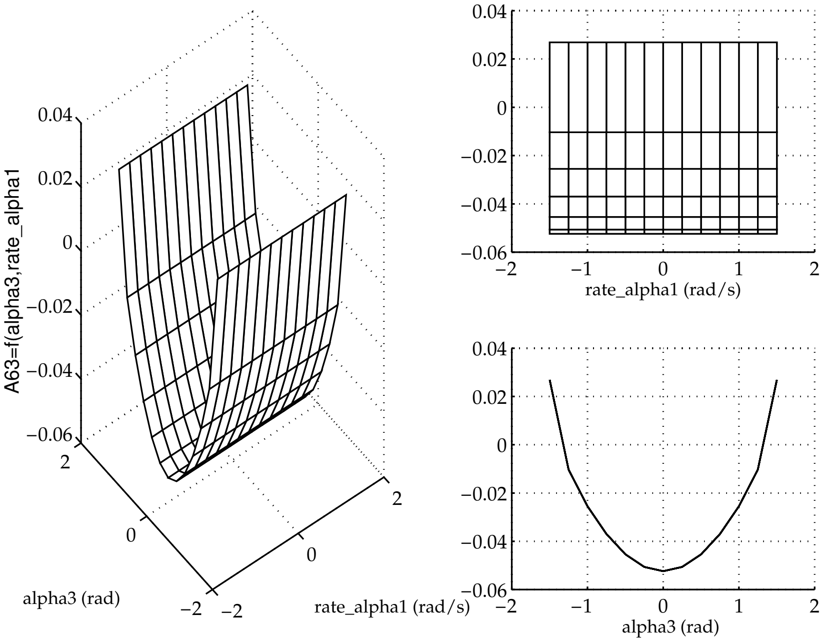

Figure 4.16: Dependence of a_{63} upon α_3 and $\dot{\alpha}_1$.

Chapter 5

Evaluation of Ducted Fan Controllers

Experiment escorts us last–
His pungent company
Will not allow an Axiom
An Opportunity.
—Emily Dickinson

In this chapter, the performance results on the controllers designed in the previous chapter are presented and they are compared with other controllers which have been designed for the ducted fan. The criteria by which a controller is judged to perform well or not is an interesting area of investigation, and one method for evaluating the controllers is presented, comprising many different performance measures. The selection and development of these measures arose from joint work with Richard Murray, Pascale Bendotti, Michael Kantner, Carolyn Beck, and Geir Dullerud, but are presented here for the first time.

Section 5.1 describes the trajectories on which the controllers were tested. Section 5.2 describes the performance measures by which the controllers are evaluated. Section 5.3 presents the results of the controller tests, along with representative plots for various controllers (complete results may be found in Appendix 5.A). Section 5.4 explains the experimental procedure used.

5.1 Description of Trajectories

The controllers were tested on four trajectories. Two of the trajectories are simple and command changes on only one axis. The third trajectory is demanding and commands rapid changes to the α_1 and α_2 axes simultaneously. The fourth trajectory is similar to the third, but slower.

The first trajectory is a one radian change on the α_1 axis over 5 seconds. The second is a 0.1 radian step change on α_2. While these trajectories are not challenging, they should demonstrate the controllers' abilities to track each axis independently.

The third trajectory is more complex and commands the fan to fly rapidly in the positive α_1 direction. During forward flight, the fan achieves $\dot{\alpha}_1$ of 0.628

rad/s, over three times greater than in the first trajectory. While in forward flight, a sinusoidal variation is commanded in α_2, with a magnitude of 0.2 radians, and a frequency of 1.26 rad/s. Thus, for every complete revolution of the fan in the lateral direction, it will go up and down twice. This trajectory is referred to as the first "rolling" trajectory.

The fourth trajectory is also a rolling trajectory. In this trajectory the fan achieves $\dot{\alpha}_1$ of 0.314 rad/s, with a sinusoidal variation of α_2 of magnitude 0.2 radians and a frequency of 0.943 rad/s. Thus the fan goes up and down about 1.5 times for every revolution in the lateral direction. This trajectory will be referred to as the second rolling trajectory.

5.2 Performance Measures

Giving a fair comparison of controllers on a real system, which have possibly been designed with radically different techniques, is a non-trivial task. A set of performance measures which attempts to give as fair an overview of the controllers as possible is presented here. The controllers are quantitatively compared based on several figures of merit, some of which have been used for other controller comparisons on the ducted fan [27]. These results are summarized in Tables 5.1 through 5.4.

The 10-90% rise time is a standard figure of merit for step responses. For the ramp in α_1 it provides a measure of how closely the ramp follows the signal. A 90% delay factor is computed by measuring the difference, in seconds, between when the commanded trajectory reaches 90% of its final value and when the system reaches this same value. Steady state error is computed by averaging the absolute value of the error over the last four seconds of the trajectory. The percent (%) overshoot is self-explanatory. In the tables, e_i will denote the error in the ith channel, i.e., the difference between α_i commanded and α_i measured. Several figures of merit involving the norms on the signal in either the α_i channel or on the error are then presented. These are signal norms and their meaning is self-explanatory. The settling time for a signal is defined as the amount of time it takes for the signal to be within ϵ of the reference signal. Thus it only applies to trajectories which end at an equilibrium position. The standard value of ϵ for a step response is 5%, and this value will be used for all calculations. The settling time is denoted by T_{s_i} for the measurement on the ith channel.

Note that $\|e_i\|_2$ is the RMS error. This is an important figure of merit because it is what the $\mathcal{H}_\infty$ and LPV controllers are attempting to minimize. Also, controllers typically will not perform well if $\|\alpha_3\|_\infty$ exceeds $\pi/2$, since the action of the forces switches at this angle, i.e., for $\alpha_3 = \pi/2$, u_1 drives the fan laterally, and u_2 vertically.

Additionally, weighted ℓ_∞ norms are computed for some trajectories. The weighting method uses a pre-specified envelope that the signal must lie within

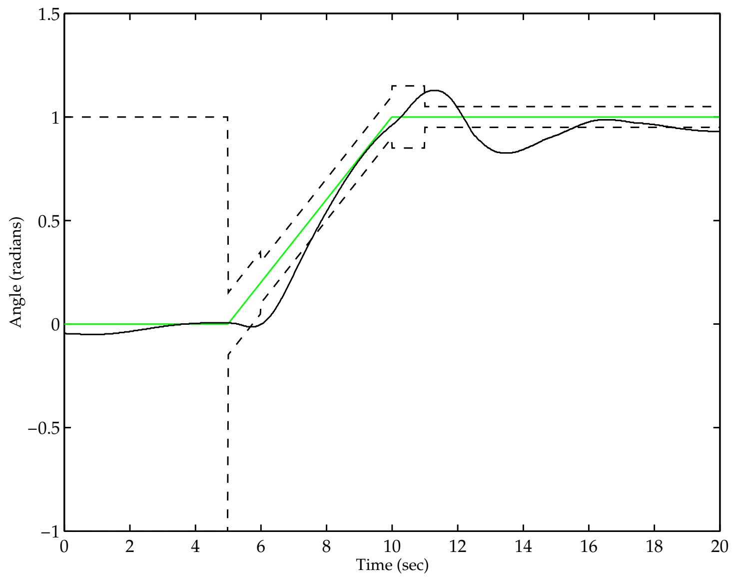

Figure 5.1: An example of the ramp envelope specification, with the response from the LPV controller. The shaded line is the desired response and the dashed lines are the allowable error bounds. The solid line is the actual response.

over the duration of the trajectory. For the ramp trajectory, the envelope consists of three parts: an overshoot limit at the beginning and end of the ramp, an error limit during the ramp, and a maximum steady state error at the end of the ramp. An example of this type of envelope, with an error signal from the LPV controller, is shown in Figure 5.1. See [16] for further discussion of this type of performance specification.

Another envelope is a step envelope, which consists of a specification on the allowed overshoot of a signal and the settling time of the signal. The final envelope used for weighting the norm is an exponential envelope. This envelope is one at time $t = 0$, and exponentially decreases at a specified exponential decay until the signal should have settled, at which time a 5% error is allowed until the end of the signal.

The weighted ℓ_∞ norms are denoted by ζ in the tables; ζ_{re} denotes the weighted norm with the ramp envelope, ζ_{se} with the step envelope, and ζ_{ee} with the exponential envelope. If ζ is less than one, then the signal stayed within the envelope throughout the test. A ζ greater than one indicates the signal exceeded the envelope, and provides a measure of how much it exceeded it by.

Finally, a performance measure is included which gives a measure for how much actuator bandwidth is being used by the controller. Ideally, a controller should perform with as low a bandwidth on the actuators as possible. The measure is computed by high-pass filtering the command channels and measuring the energy of the resulting signal. A lower number is better. These measures are denoted by $\|W * u_i\|_2$ for $i = 1, 2$.

In the table, a row ranking is computed for each measure, and the rank of the controller in comparison to the others is shown in the upper right hand corner. One way of determining how good each controller is would then be to add up the ranks. The controller with the lowest sum is the best. This method of evaluation has problems: it assigns equal weight to all measures and doesn't take into account how much better a controller performed than another on a given measure. Additionally, some performance measures are complementary, i.e., one can only be raised at the expense of another. A simple example of this would be having a large $\|\dot{\alpha}_1\|$ and small $\|\alpha_3\|$. If $\|\alpha_3\| = 0$ then $\dot{\alpha}_1 = 0$, so some tradeoff is required. Nonetheless, this is reasonable, since often performance specifications are complementary, and must be traded off. The rank total for each controller is also presented. One important note about the rankings is that each entry in the table is presented in finite precision, but the rankings were computed from the actual data, and thus specifications which appear tied really may not be.

The controllers presented in the tables are the ones discussed previously, namely the $\mathcal{H}_\infty$, LPV, and locked LPV controllers, in addition to some new ones. The design of these new controllers can be found in [27], and nothing beyond a few superficial remarks will be made here. The LQR design is standard; an integrator is added to the controller on the α_2 channel. The gain scheduled controller is a combination of three LQR designs around different equilibrium points of the fan's forward flight and hover. Bumpless transfer is ensured since the controllers all share one integrator, and otherwise have no states.

The ducted fan is I/O linearizable with respect to any pair of outputs (α_i, α_j). The I/O decoupled controller is a nonlinear controller consisting of an I/O linearization stage followed by a "loop-at-a-time" heuristic design. We note in passing that standard I/O linearization designs use pole-placement, which does not work well for the ducted fan.

5.3 Results and Evaluations

Table 5.1 shows the figures of merit for the ramp in α_1, Table 5.2 shows the results for the step in α_2, Table 5.3 shows the results for the first rolling trajectory, and Table 5.4 shows the results for the second rolling trajectory. The LPV controller locked at hover is denoted "LPV$_{hl}$" in the tables and when locked in forward flight it is denoted "LPV$_{fl}$."

A plot for each trajectory showing behavior of the $\mathcal{H}_\infty$, LPV, and locked LPV controllers is found in Appendix 5.A. Figures for the other controllers may be found in [27]. These plots show the α_1, α_2, and α_3 channels, with desired and actual values, in addition to the commanded forces. Recall that

Criteria	Controllers						
	LQR	$\mathcal{H}_\infty$	I/O D.	Gain S	LPV	LPV$_{fl}$	LPV$_{hl}$
α_1 rise time	3.34 ^{3}	3.42 ^{4}	3.68 ^{5}	3.24 ^{2}	3.04 ^{1}	4.62 ^{6}	4.88 ^{7}
α_1 90% delay	0.62 ^{3}	1.16 ^{4}	1.62 ^{5}	0.48 ^{2}	0.10 ^{1}	1.80 ^{6}	2.12 ^{7}
α_1 % ov'shoot	7.24 ^{2}	3.15 ^{1}	18.87 ^{6}	8.81 ^{3}	12.78 ^{4}	15.61 ^{5}	25.35 ^{7}
steady state e_1	0.03 ^{3}	0.02 ^{1}	0.07 ^{5}	0.02 ^{1}	0.04 ^{4}	0.11 ^{6}	0.23 ^{7}
$\|e_1\|_\infty$	0.26 ^{3}	0.34 ^{5}	0.41 ^{7}	0.25 ^{2}	0.22 ^{1}	0.34 ^{4}	0.40 ^{6}
$\|e_1\|_2$	2.53 ^{2}	3.96 ^{4}	6.09 ^{6}	2.35 ^{1}	2.79 ^{3}	5.18 ^{5}	6.68 ^{7}
$\|e_1\|_1$	54 ^{2}	74 ^{4}	137 ^{6}	49 ^{1}	67 ^{3}	126 ^{5}	169 ^{7}
$\|\alpha_2\|_\infty$	0.05 ^{3}	0.05 ^{3}	0.03 ^{1}	0.07 ^{6}	0.04 ^{2}	0.06 ^{5}	0.09 ^{7}
$\|e_2\|_\infty$	0.05 ^{3}	0.05 ^{3}	0.03 ^{1}	0.07 ^{6}	0.04 ^{2}	0.06 ^{5}	0.09 ^{7}
$\|e_2\|_2$	0.51 ^{4}	0.44 ^{2}	0.37 ^{1}	0.68 ^{5}	0.45 ^{3}	0.72 ^{6}	1.21 ^{7}
$\|e_2\|_1$	10.4 ^{4}	8.1 ^{1}	8.8 ^{2}	14.9 ^{6}	9.3 ^{3}	13.8 ^{5}	26.3 ^{7}
$\|\alpha_3\|_\infty$	1.28 ^{7}	1.12 ^{4}	0.80 ^{3}	1.20 ^{5}	1.20 ^{5}	0.76 ^{1}	0.78 ^{2}
$\|\dot{\alpha}_3\|_\infty$	4.14 ^{7}	2.29 ^{4}	2.05 ^{2}	3.35 ^{6}	2.66 ^{5}	1.63 ^{1}	2.07 ^{3}
ζ_{re}	1.72 ^{2}	2.95 ^{3}	4.14 ^{6}	1.69 ^{1}	3.48 ^{5}	3.36 ^{4}	5.07 ^{7}
T_{s_1}	6.70 ^{2}	5.92 ^{1}	13.42 ^{4}	9.00 ^{3}	14.98 ^{6}	14.80 ^{5}	14.98 ^{7}
ζ_{ee}	1.07 ^{4}	0.94 ^{3}	0.63 ^{1}	1.32 ^{6}	0.75 ^{2}	1.26 ^{5}	1.88 ^{7}
$\|W * u_1\|_2$	1.98 ^{2}	0.87 ^{1}	4.88 ^{7}	3.76 ^{6}	2.35 ^{5}	1.98 ^{3}	2.17 ^{4}
$\|W * u_2\|_2$	9.98 ^{5}	0.89 ^{1}	13.83 ^{7}	10.20 ^{6}	3.40 ^{3}	4.59 ^{4}	3.23 ^{2}
Rank total	59 ^{3}	49 ^{1}	75 ^{5}	68 ^{4}	58 ^{2}	87 ^{6}	108 ^{7}

Table 5.1: Performance data for the α_1 ramp. See the text for an explanation of the controllers and measures.

Criteria	Controllers						
	LQR	$\mathcal{H}_\infty$	I/O D.	Gain S	LPV	LPV$_{fl}$	LPV$_{hl}$
$\|\alpha_1\|_\infty$	0.01 ^{2}	0.04 ^{5}	0.00 ^{1}	0.02 ^{3}	0.15 ^{7}	0.10 ^{6}	0.03 ^{4}
$\|e_1\|_\infty$	0.01 ^{2}	0.04 ^{5}	0.00 ^{1}	0.02 ^{3}	0.15 ^{7}	0.10 ^{6}	0.03 ^{4}
$\|e_1\|_2$	0.18 ^{2}	0.82 ^{5}	0.07 ^{1}	0.23 ^{3}	2.03 ^{6}	2.30 ^{7}	0.30 ^{4}
$\|e_1\|_1$	4.8 ^{2}	22.3 ^{5}	1.6 ^{1}	5.9 ^{3}	49.4 ^{6}	65.4 ^{7}	7.3 ^{4}
α_2 rise time	1.12 ^{7}	0.96 ^{5}	0.64 ^{1}	1.10 ^{6}	0.76 ^{3}	0.74 ^{2}	0.78 ^{4}
α_2 90% delay	1.90 ^{7}	1.72 ^{5}	1.02 ^{1}	1.76 ^{6}	1.34 ^{2}	1.36 ^{3}	1.36 ^{4}
α_2 % ov'shoot	10.0 ^{1}	22.5 ^{3}	25.6 ^{4}	19.4 ^{2}	47.7 ^{6}	50.8 ^{7}	44.5 ^{5}
steady st. e_2	0.00 ^{5}	0.00 ^{6}	0.01 ^{7}	0.00 ^{4}	0.00 ^{3}	0.00 ^{1}	0.00 ^{2}
$\|\alpha_2\|_\infty$	0.11 ^{1}	0.12 ^{3}	0.13 ^{4}	0.12 ^{2}	0.15 ^{6}	0.15 ^{7}	0.14 ^{5}
$\|e_2\|_2$	0.75 ^{6}	0.74 ^{5}	0.58 ^{1}	0.71 ^{2}	0.72 ^{4}	0.75 ^{7}	0.71 ^{3}
$\|e_2\|_1$	10.0 ^{5}	10.3 ^{6}	9.8 ^{2}	8.6 ^{1}	10.0 ^{4}	10.7 ^{7}	9.8 ^{3}
$\|\alpha_3\|_\infty$	0.09 ^{2}	0.21 ^{4}	0.05 ^{1}	0.09 ^{3}	0.63 ^{7}	0.21 ^{5}	0.28 ^{6}
$\|\dot{\alpha}_3\|_\infty$	0.35 ^{2}	0.51 ^{6}	0.28 ^{1}	0.41 ^{4}	1.03 ^{7}	0.36 ^{3}	0.45 ^{5}
ζ_{se}	0.40 ^{1}	0.45 ^{4}	0.40 ^{1}	0.40 ^{1}	0.95 ^{6}	1.02 ^{7}	0.89 ^{5}
ζ_{ee}	0.92 ^{5}	0.90 ^{4}	0.74 ^{1}	0.87 ^{2}	0.95 ^{6}	1.02 ^{7}	0.89 ^{3}
T_{s_2}	1.32 ^{6}	1.22 ^{4}	0.66 ^{1}	1.22 ^{5}	0.96 ^{2}	2.32 ^{7}	1.00 ^{3}
$\|W * u_1\|_2$	2.27 ^{6}	1.21 ^{1}	10.97 ^{7}	1.96 ^{3}	2.20 ^{5}	2.16 ^{4}	1.90 ^{2}
$\|W * u_2\|_2$	4.90 ^{6}	0.35 ^{1}	9.21 ^{7}	4.71 ^{5}	1.95 ^{2}	1.96 ^{3}	2.23 ^{4}
Rank total	68 ^{3}	72 ^{5}	42 ^{1}	58 ^{2}	87 ^{6}	96 ^{7}	70 ^{4}

Table 5.2: Performance data for the α_2 step. See the text for an explanation of the controllers and measures.

Criteria	Controllers						
	LQR	$\mathcal{H}_\infty$	I/O D.	Gain S	LPV	LPV$_{fl}$	LPV$_{hl}$
e_1 at $t = t_f$	0.39 ^{5}	0.60 ^{6}	0.37 ^{4}	0.18 ^{1}	0.34 ^{3}	3.31 ^{7}	0.30 ^{2}
$\|\alpha_1\|_\infty$	13.74 ^{2}	13.53 ^{1}	13.76 ^{3}	13.95 ^{5}	13.80 ^{4}	17.43 ^{7}	14.43 ^{6}
$\|e_1\|_\infty$	0.41 ^{3}	0.73 ^{5}	0.74 ^{6}	0.26 ^{1}	0.36 ^{2}	3.31 ^{7}	0.68 ^{4}
$\|e_1\|_2$	11.89 ^{3}	19.43 ^{6}	13.50 ^{5}	6.33 ^{1}	6.49 ^{2}	54.58 ^{7}	12.17 ^{4}
$\|e_1\|_1$	404 ^{4}	659 ^{6}	441 ^{5}	213 ^{2}	188 ^{1}	1472 ^{7}	382 ^{3}
e_2 at $t = t_f$	0.06 ^{5}	0.04 ^{3}	0.16 ^{7}	0.05 ^{4}	0.01 ^{1}	0.03 ^{2}	0.09 ^{6}
$\|\alpha_2\|_\infty$	0.15 ^{1}	0.30 ^{7}	0.30 ^{6}	0.23 ^{3}	0.26 ^{4}	0.18 ^{2}	0.27 ^{5}
$\|e_2\|_\infty$	0.25 ^{6}	0.14 ^{2}	0.16 ^{3}	0.20 ^{4}	0.13 ^{1}	0.28 ^{7}	0.23 ^{5}
$\|e_2\|_2$	4.35 ^{6}	1.60 ^{1}	2.63 ^{4}	2.15 ^{3}	1.64 ^{2}	5.00 ^{7}	3.21 ^{5}
$\|e_2\|_1$	123 ^{6}	43 ^{2}	72 ^{4}	48 ^{3}	39 ^{1}	141 ^{7}	85 ^{5}
$\|\alpha_3\|_\infty$	1.49 ^{6}	1.42 ^{5}	1.04 ^{1}	1.08 ^{2}	1.84 ^{7}	1.18 ^{3}	1.27 ^{4}
$\|\dot{\alpha}_3\|_\infty$	2.03 ^{5}	3.00 ^{7}	0.98 ^{2}	1.49 ^{4}	2.45 ^{6}	0.63 ^{1}	1.26 ^{3}
$\|W * u_1\|_2$	2.97 ^{2}	3.78 ^{4}	6.59 ^{7}	4.09 ^{5}	5.54 ^{6}	3.38 ^{3}	2.52 ^{1}
$\|W * u_2\|_2$	5.40 ^{5}	0.87 ^{1}	8.26 ^{7}	5.45 ^{6}	3.38 ^{3}	3.43 ^{4}	1.24 ^{2}
Rank Total	59 ^{4}	56 ^{3}	69 ^{6}	44 ^{2}	43 ^{1}	71 ^{7}	61 ^{5}

Table 5.3: Performance data for the first rolling trajectory. See the text for an explanation of the controllers and measures.

Criteria	Controllers			
	$\mathcal{H}_\infty$	LPV	LPV$_{fl}$	LPV$_{hl}$
e_1 at $t = t_f$	0.76 ^{3}	0.15 ^{1}	1.33 ^{4}	0.11 ^{2}
$\|\alpha_1\|_\infty$	6.30 ^{1}	6.91 ^{2}	8.40 ^{4}	6.95 ^{3}
$\|e_1\|_\infty$	0.76 ^{3}	0.33 ^{1}	1.33 ^{4}	0.56 ^{2}
$\|e_1\|_2$	19.61 ^{4}	5.78 ^{1}	14.70 ^{3}	11.39 ^{2}
$\|e_1\|_1$	658.46 ^{4}	175.84 ^{1}	369.28 ^{3}	340.70 ^{2}
e_2 at $t = t_f$	0.01 ^{1}	0.12 ^{3}	0.25 ^{4}	0.05 ^{2}
$\|\alpha_2\|_\infty$	0.29 ^{2}	0.32 ^{3}	0.28 ^{1}	0.35 ^{4}
$\|e_2\|_\infty$	0.24 ^{1}	0.27 ^{3}	0.27 ^{2}	0.29 ^{4}
$\|e_2\|_2$	2.13 ^{1}	2.86 ^{2}	4.11 ^{4}	3.24 ^{3}
$\|e_2\|_1$	47.32 ^{1}	67.48 ^{2}	104.86 ^{4}	85.34 ^{3}
$\|\alpha_3\|_\infty$	0.96 ^{1}	1.58 ^{3}	1.01 ^{2}	1.59 ^{4}
$\|\dot{\alpha}_3\|_\infty$	1.05 ^{2}	2.77 ^{4}	0.62 ^{1}	2.22 ^{3}
$\|W * u_1\|_2$	1.57 ^{1}	3.34 ^{4}	3.14 ^{3}	2.96 ^{2}
$\|W * u_2\|_2$	0.18 ^{1}	2.05 ^{3}	1.12 ^{2}	2.26 ^{4}
Rank total	26 ^{1}	33 ^{2}	41 ^{4}	37 ^{3}

Table 5.4: Performance data for the second rolling trajectory. See the text for an explanation of the controllers and measures.

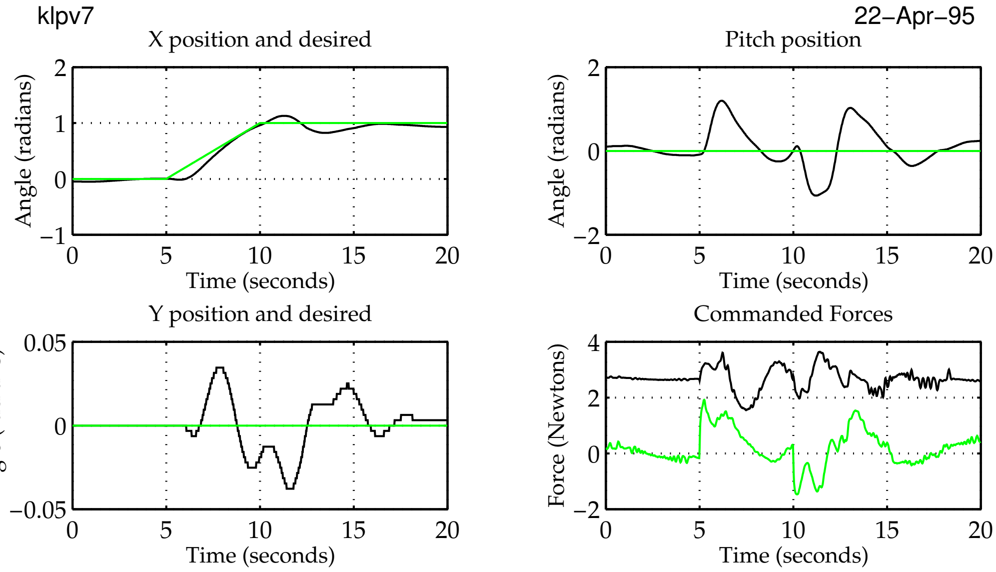

Figure 5.2: Closed loop response of the LPV controller for an α_1 ramp.

the "X" direction is α_1, the "Y" direction is α_2, and the pitch is α_3.

For the α_1 trajectory data shown in Table 5.1, several results are apparent. The LPV is the fastest controller, but has problems settling. This is shown in Figure 5.2. Due to its problems settling, it does not perform well on the ramp envelope. The $\mathcal{H}_\infty$ controller is slower than any of the gain-scheduled, LPV, or LQR controllers, but has the lowest overshoot, steady state error, and settling times. The I/O decoupling controller does the best at holding the fan at a constant α_2 during the trajectory; this is not surprising in view of its design. In terms of actuator bandwidth, both the $\mathcal{H}_\infty$ and LPV controllers are markedly lower than the I/O decoupling, LQR, and gain-scheduled controllers. The I/O decoupling controller is particularly bad, as is shown in Figure 5.3. Based on these results, it seems fair to conclude that the $\mathcal{H}_\infty$ controller is the best of the group if the low errors are desired, while the LPV is best if speedy response at the expense of some error is desired. All controllers are able to follow the ramp in α_1.

The performance results for the step in α_2 are shown in Table 5.2. The I/O decoupling controller is an extremely good controller on this trajectory. This is somewhat surprising, since, given the nature of the trajectory, no decoupling is required for it. The I/O decoupling controller has the worst steady-state error of the group; this shows the lack of an integrator on α_2. The LQR and gain-scheduled controllers are quite similar, which is not surprising either, since no scheduling occurs on the trajectory. The $\mathcal{H}_\infty$ controller is rather slow. The LPV controller does not have very good performance compared to the other controllers. It has a fast rise-time, but a considerable amount of overshoot. Note, however, that it settles very quickly and is within the step envelope

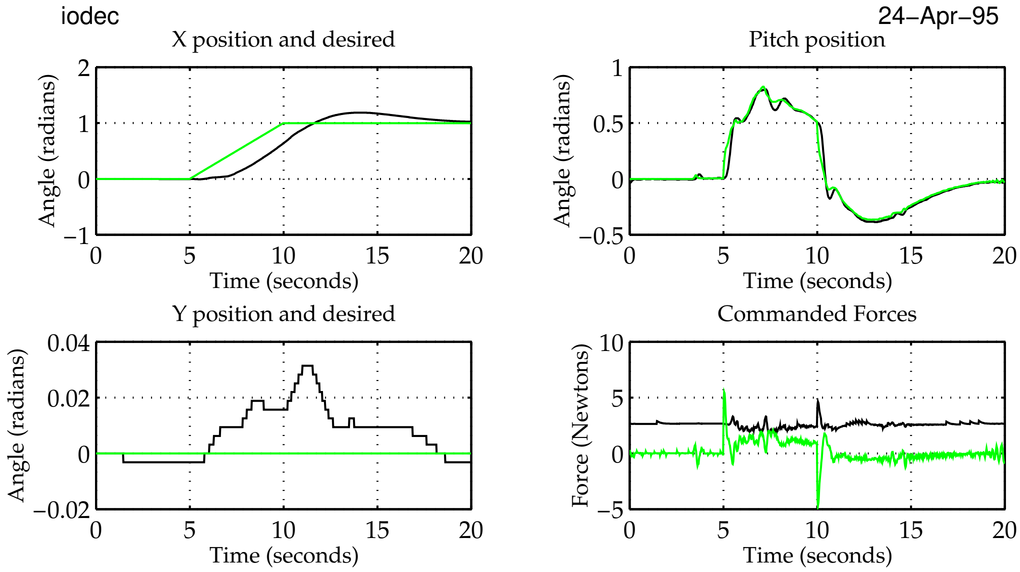

Figure 5.3: Closed loop response of the I/O decoupled controller for an α_1 ramp.

specification.

The most surprising thing about the LPV controller is a rather bizarre offset in the α_1 channel for this trajectory, shown in Figure 5.4. This behavior is not fully understood at this time, but is not an artifact of unmodelled dynamics since the full nonlinear simulation also shows this behavior.

The results for the first rolling trajectory are tabulated in Table 5.3. The gain-scheduled and LPV controllers are close in performance, as is illustrated in Figures 5.5 and 5.6. The gain-scheduled controller is slightly better on α_1 while the LPV controller is slightly better on α_2. The LPV controller is better overall on tracking α_1 and α_2, but is penalized by the ranking scheme for tilting the fan more to do it. As evidenced by the 1-norms, though, the LPV controller tracks the signal significantly better. It is also interesting to note that this is the first example of a trajectory where the linear and I/O decoupled controllers are unable to keep up.

When the trajectory is slowed down, the picture changes somewhat, as evidenced in Table 5.4. Unfortunately, the full spectrum of ducted fan controllers is unavailable for this trajectory. The LPV controller outperforms the $\mathcal{H}_\infty$ controller significantly on the α_1 measures, but is slightly edged in the α_2 measures. The LPV has more overshoot in α_1 than the $\mathcal{H}_\infty$. The LPV controller does a significantly better job of tracking the references overall, though.

It is interesting to note how poorly the locked controllers do. They are almost uniformly bad. This illustrates that the LPV controller is exploiting the parameter variations and changes in feedforward forces.

Finally, since ranking is obviously not the best way of evaluating a con-

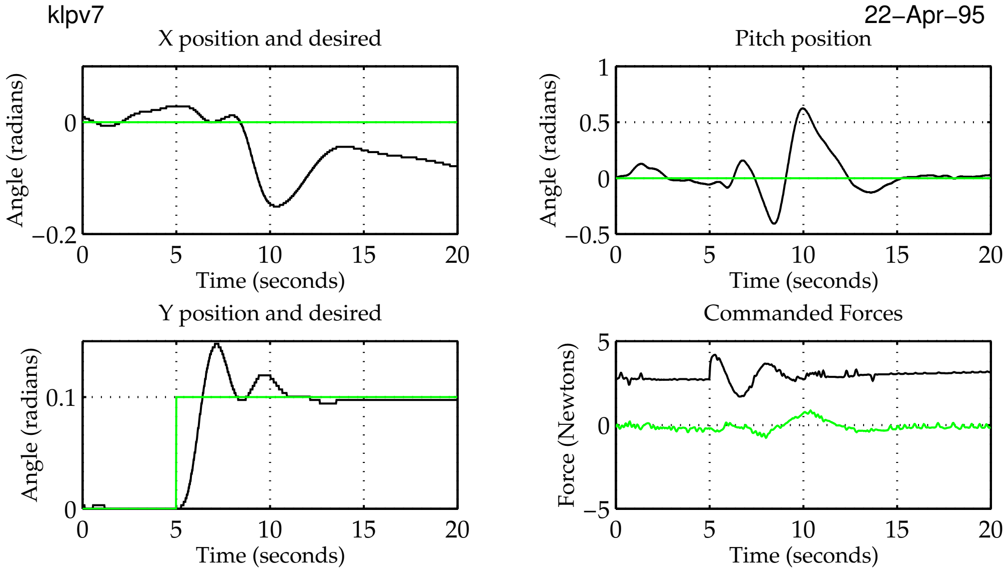

Figure 5.4: Closed loop response of the LPV controller for an α_2 step.

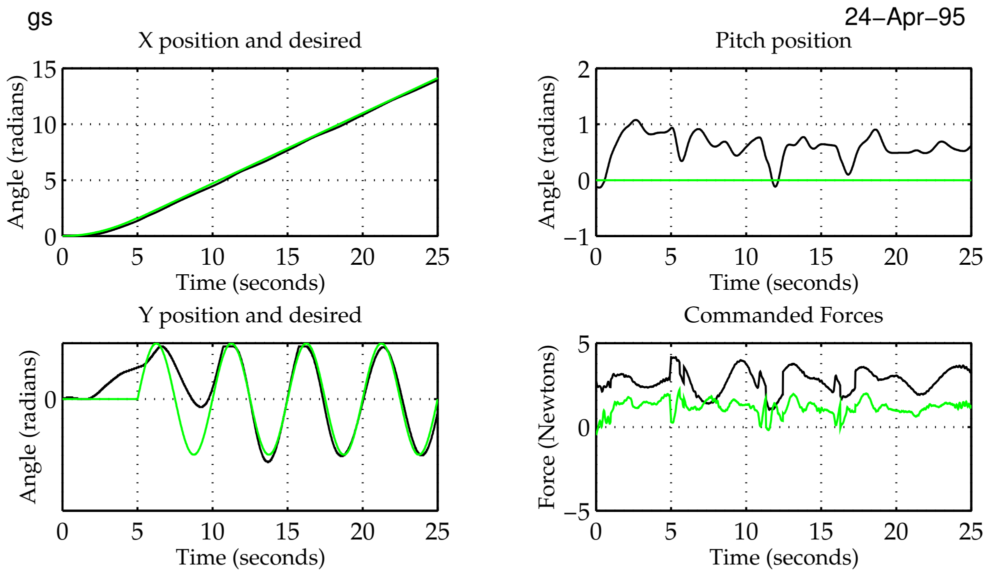

Figure 5.5: Closed loop response of the gain-scheduled controller for the first rolling trajectory.

Criteria	Controllers				
	LQR	$\mathcal{H}_\infty$	I/O D.	Gain S	LPV
e_1 at $t = t_f$	0.65^{4}	1.00^{5}	0.62^{3}	0.31^{1}	0.56^{2}
$\|\alpha_1\|_\infty$	0.99^{2}	0.97^{1}	0.99^{3}	1.00^{5}	0.99^{4}
$\|e_1\|_\infty$	0.55^{3}	0.98^{4}	1.00^{5}	0.35^{1}	0.49^{2}
$\|e_1\|_2$	0.61^{3}	1.00^{5}	0.69^{4}	0.33^{1}	0.33^{2}
$\|e_1\|_1$	0.61^{3}	1.00^{5}	0.67^{4}	0.32^{2}	0.29^{1}
e_2 at $t = t_f$	0.38^{4}	0.25^{2}	1.00^{5}	0.31^{3}	0.06^{1}
$\|\alpha_2\|_\infty$	0.49^{1}	1.00^{5}	0.98^{4}	0.74^{2}	0.86^{3}
$\|e_2\|_\infty$	1.00^{5}	0.56^{2}	0.65^{3}	0.80^{4}	0.52^{1}
$\|e_2\|_2$	1.00^{5}	0.37^{1}	0.61^{4}	0.49^{3}	0.38^{2}
$\|e_2\|_1$	1.00^{5}	0.35^{2}	0.59^{4}	0.39^{3}	0.32^{1}
$\|\alpha_3\|_\infty$	0.81^{4}	0.77^{3}	0.57^{1}	0.59^{2}	1.00^{5}
$\|\dot{\alpha}_3\|_\infty$	0.68^{3}	1.00^{5}	0.33^{1}	0.50^{2}	0.82^{4}
$\|W_{act} * u_1\|_2$	0.45^{1}	0.57^{2}	1.00^{5}	0.62^{3}	0.84^{4}
$\|W_{act} * u_2\|_2$	0.65^{3}	0.11^{1}	1.00^{5}	0.66^{4}	0.41^{2}
Rank Total	59^{4}	56^{3}	69^{6}	44^{2}	43^{1}
Primary	5.43^{5}	5.26^{4}	4.83^{3}	3.00^{2}	2.89^{1}
Primary II	5.80^{4}	5.51^{3}	5.83^{5}	3.31^{2}	2.95^{1}
with $\|\alpha_i\|$, $i = 1, 2$	7.28^{3}	7.48^{4}	7.80^{5}	5.05^{2}	4.80^{1}
with $\|W * u\|$	8.39^{4}	8.16^{3}	9.80^{5}	6.33^{2}	6.05^{1}
with $\|\alpha_3, \dot{\alpha}_3\|_\infty$	9.87^{3}	9.93^{4}	10.69^{5}	7.41^{1}	7.87^{2}

Table 5.5: Normalized performance data for the first rolling trajectory, with primary and secondary performance measures.

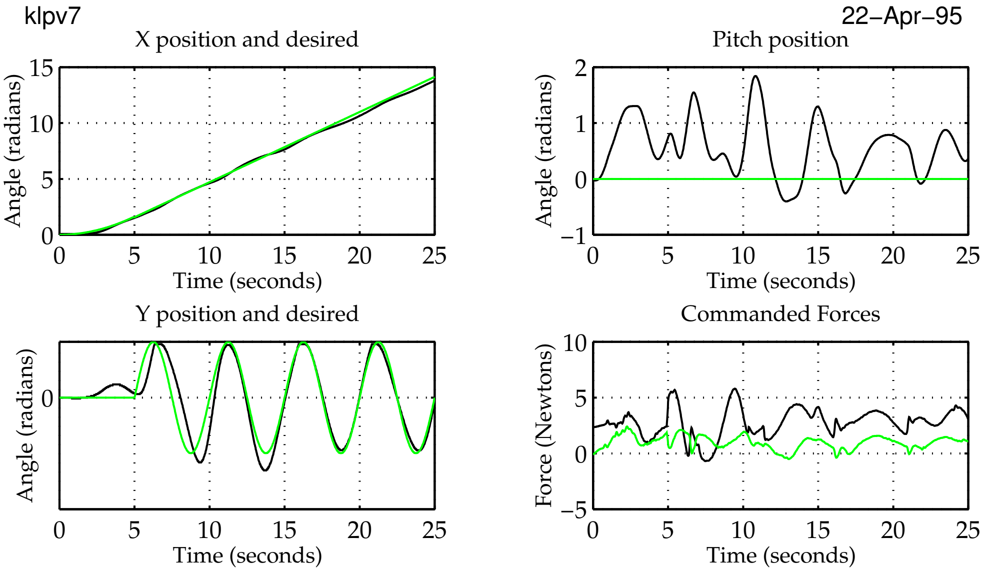

Figure 5.6: Closed loop response of the LPV controller for the first rolling trajectory.

troller, we present an improved scheme for the rolling trajectory, shown in Table 5.5 for the non-locked controllers. First, each row of the performance measures is normalized by dividing by the worst performance measures. To evaluate the controllers we then add up the respective numbers and present a sum. The better controller is the one with the lowest sum. This type of evaluation will take into consideration how close various controllers are in a fairer way than simple ranking.

The performance measures are also grouped into primary and secondary measures in Table 5.5. A primary performance measure is one in which the open loop fan will not perform well, and a secondary measure is one in which it will. For example, $\|\alpha_3\|_\infty$ is clearly a secondary measure, since the fan hovering will not change pitch, and thus this will be zero. The RMS error in either channel is an example of a primary measure. The primary measures for the rolling trajectory are the norms of e_i, and the final error on e_1. Because the rolling trajectory ends at the same height that it started at, the final error in α_2 is not primary, but as this is more an artifact of our trajectory, we consider it with the primary performance measures in the row labelled "Primary II." We then start evaluating the controllers as secondary measures are added: first including the ∞-norms on α_1 and α_2, then the actuator bandwidth measures, and finally the ∞-norms on α_3 and $\dot{\alpha}_3$.

The most noticeable thing about these measures is the grouping: the LPV and gain scheduled controllers are always quite close, usually within 10% of one another. All other controllers have much larger sums, and the I/O decoupling controller is usually the worst. On the primary performance specifications the

LPV controller is the clear winner, but is edged out by the gain scheduled controller when the norms on α_3 and $\dot{\alpha}_3$ are included.

5.4 Experimental Procedure

The previous results present single runs of controllers on a given trajectory. All data was taken during the same experimental session, but important questions are how repeatable are the results and how do the performance measures vary from experiment to experiment.

Such an analysis is very time consuming, and we shall not do a complete one, but rather give an indication of the results by analyzing one particular controller, the $\mathcal{H}_\infty$ controller, over ten different experiments. We shall use very basic statistical analysis; see [9] for more detail.

Let n be the number of experiments we have, and let $\widehat{x}$ denote the mean value of the set, i.e.,

$$\widehat{x} = \frac{1}{N}\sum_{i=1}^{n} x_i.$$

The standard deviation of the sample is

$$S_x = \sqrt{\frac{(\sum_{i=1}^{n} x_i^2) - n\widehat{x}^2}{n-1}}.$$

Assume the underlying distribution of the samples is Gaussian. The 95% confidence interval is defined as the interval within which the true mean value ξ of the population will lie with probability 0.95. The confidence interval is given by

$$\widehat{x} - t_{0.025}\frac{S_x}{\sqrt{n}} < \xi < \widehat{x} + t_{0.025}\frac{S_x}{\sqrt{n}}$$

where $t_{0.025}$ is computed from the t-distribution and for ten samples is 2.228.

The confidence intervals for the performance measures on the rolling trajectory with ten runs of the $\mathcal{H}_\infty$ controller are shown in Table 5.6. The first column shows the performance measures for the first experimental run of the controller on the trajectory, which are also the results reported in Table 5.3. The second column shows the lower bound of the confidence interval for the performance measures and is denoted by "$\mathcal{H}_\infty$ lb;" the third column shows the mean value of performance measures for the ten experiments and is denoted by "$\mathcal{H}_\infty$ mn;" the fourth column shows the upper bound of the confidence interval and is denoted by "$\mathcal{H}_\infty$ ub." The fifth column shows the performance data for the LPV controller, which is repeated for ease of reference. Finally, the last column describes what other non-locked controllers from Table 5.3 have performance measures within the confidence interval. We see that while some of the controllers have performance measures which overlap, the data we

Criteria	Controllers					Within
	$\mathcal{H}_\infty$	$\mathcal{H}_\infty$ lb	$\mathcal{H}_\infty$ mn	$\mathcal{H}_\infty$ ub	LPV	
e_1 at $t = t_f$	0.60	0.45	0.56	0.68	0.34	None
$\|\alpha_1\|_\infty$	13.53	13.45	13.57	13.68	13.80	None
$\|e_1\|_\infty$	0.73	0.73	0.79	0.84	0.36	I/O
$\|e_1\|_2$	19.42	17.70	18.87	20.04	6.49	None
$\|e_1\|_1$	658.47	585.24	628.75	672.25	188.47	None
e_2 at $t = t_f$	0.04	0.02	0.04	0.06	0.01	LQR, GS
$\|\alpha_2\|_\infty$	0.30	0.30	0.31	0.33	0.26	I/O
$\|e_2\|_\infty$	0.14	0.13	0.15	0.17	0.13	I/O, LPV
$\|e_2\|_2$	1.60	1.60	1.95	2.31	1.64	GS, LPV
$\|e_2\|_1$	43.45	42.79	53.54	64.29	39.26	GS
$\|\alpha_3\|_\infty$	1.42	1.44	1.52	1.60	1.84	LQR
$\|\dot{\alpha}_3\|_\infty$	3.00	2.91	3.25	3.60	2.45	None
$\|W_{act} * u_1\|_2$	3.78	2.77	3.24	3.72	5.54	LQR
$\|W_{act} * u_2\|_2$	0.87	0.89	1.01	1.14	3.38	None

Table 5.6: 95% confidence intervals for the performance measures of the rolling trajectory for the $\mathcal{H}_\infty$ controller.

have chosen to present is a fair evaluation of the $\mathcal{H}_\infty$ controller. Clearly, a complete analysis would have confidence intervals for all the other controllers as well.

5.5 Summary

In this chapter we have evaluated the performance of a set of controllers for the ducted fan. For a given trajectory the best controller is difficult to choose *a priori* but, based on its performance on complex trajectories and in the α_1 ramp, the LPV controller is the best controller overall. We demonstrated that a first principles model was adequate for designing controllers which achieve good performance and that standard linear techniques seem inadequate for good tracking performance. Future work should include the investigation of more aggressive trajectories and the improvement of the nonlinear model to include gyroscopic and friction effects.

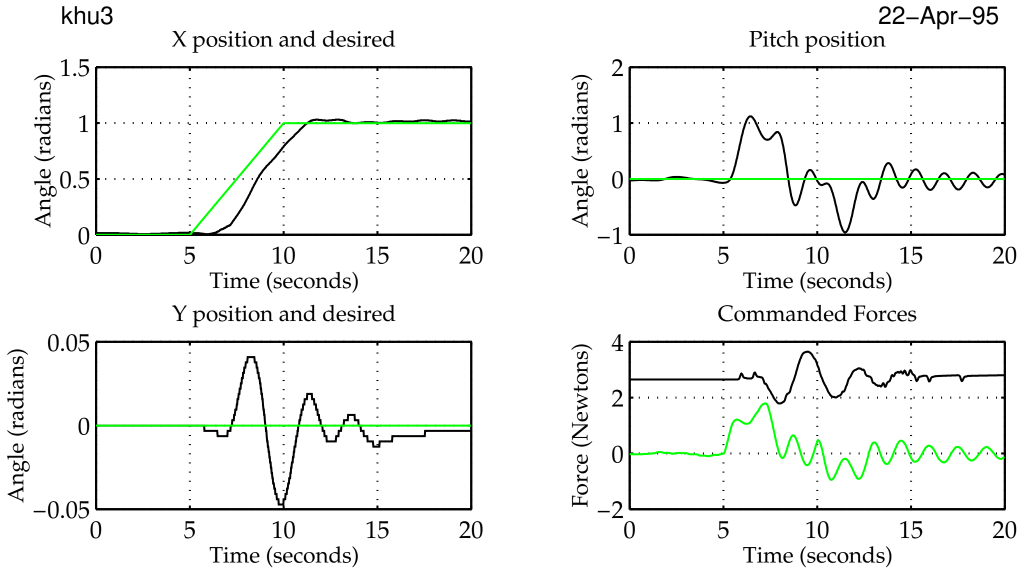

Figure 5.7: Closed loop response of the $\mathcal{H}_\infty$ controller for an α_1 ramp.

5.A Appendix: Plots of Controller Runs

In this appendix, plots of the controller experiments for the $\mathcal{H}_\infty$, LPV, and locked controllers are shown in Figures 5.7 through 5.22. They are grouped by trajectory, for the ramp in α_1, step in α_2, and two rolling trajectories. In the plots, reference signals are shown in shaded lines for the α_1, α_2, and α_3 directions ("X," "Y," and "Pitch" in the plots) while the controllers' responses are shown in solid lines. The commanded forces u_1 and u_2 are shown as well, in solid and shaded lines respectively.

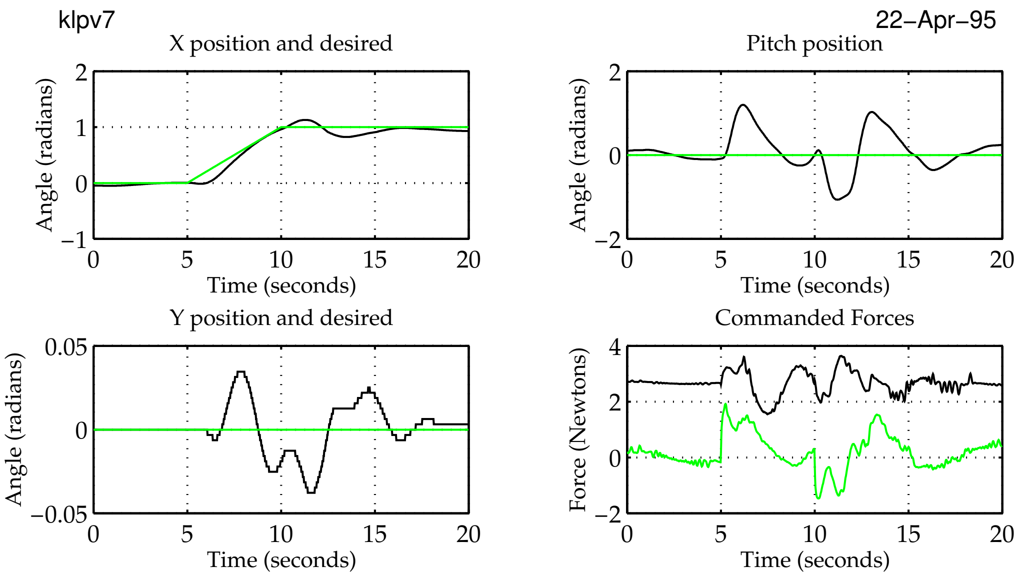

Figure 5.8: Closed loop response of the LPV controller for an α_1 ramp.

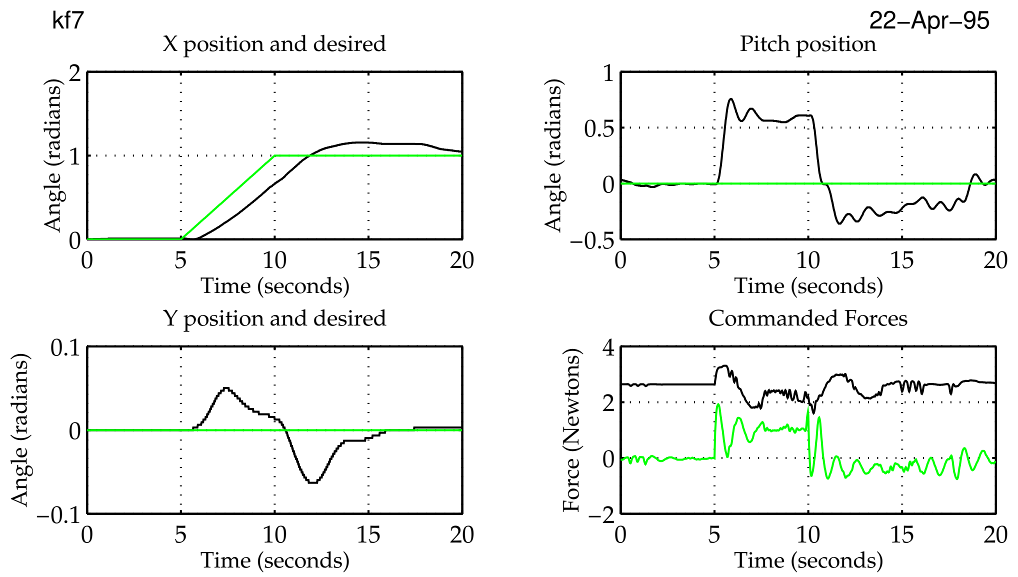

Figure 5.9: Closed loop response of the LPV controller locked in forward flight for an α_1 ramp.

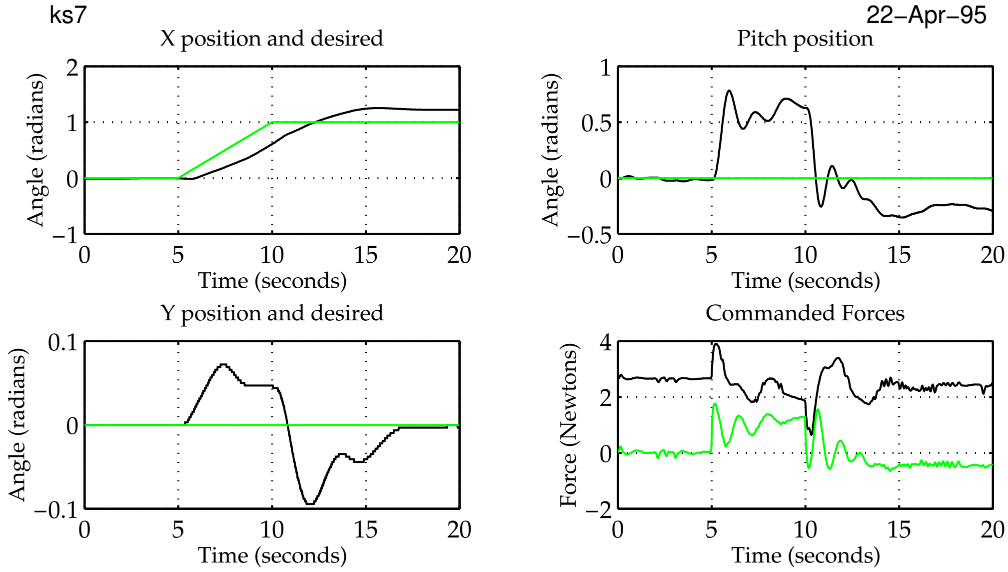

Figure 5.10: Closed loop response of the LPV controller locked at hover for an α_1 ramp.

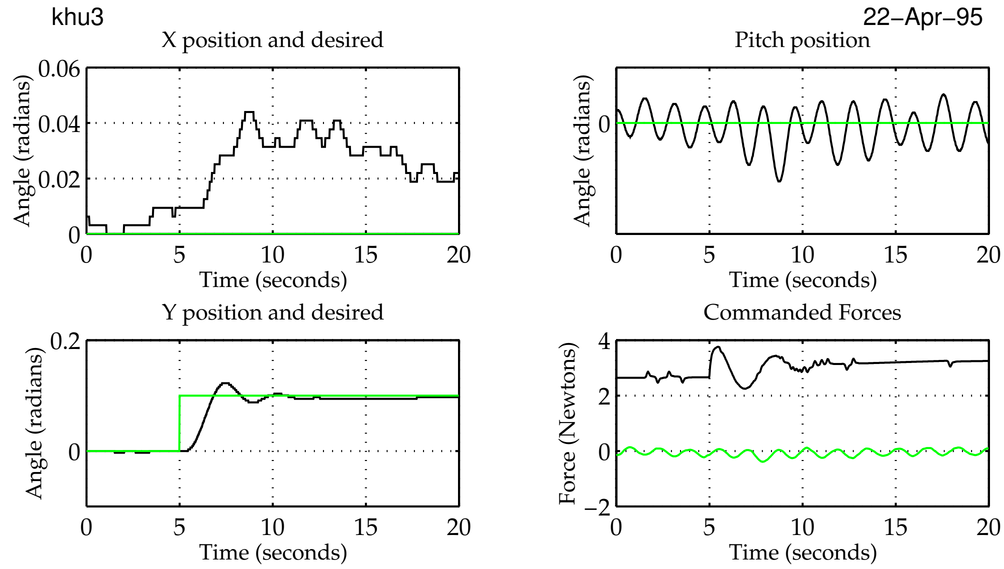

Figure 5.11: Closed loop response of the $\mathcal{H}_\infty$ controller for an α_2 step.

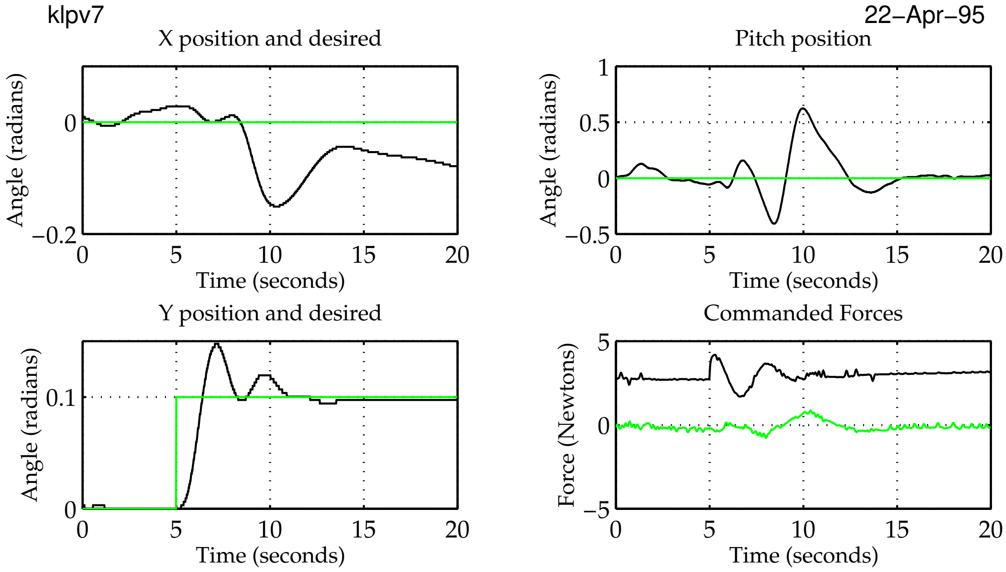

Figure 5.12: Closed loop response of the LPV controller for an α_2 step.

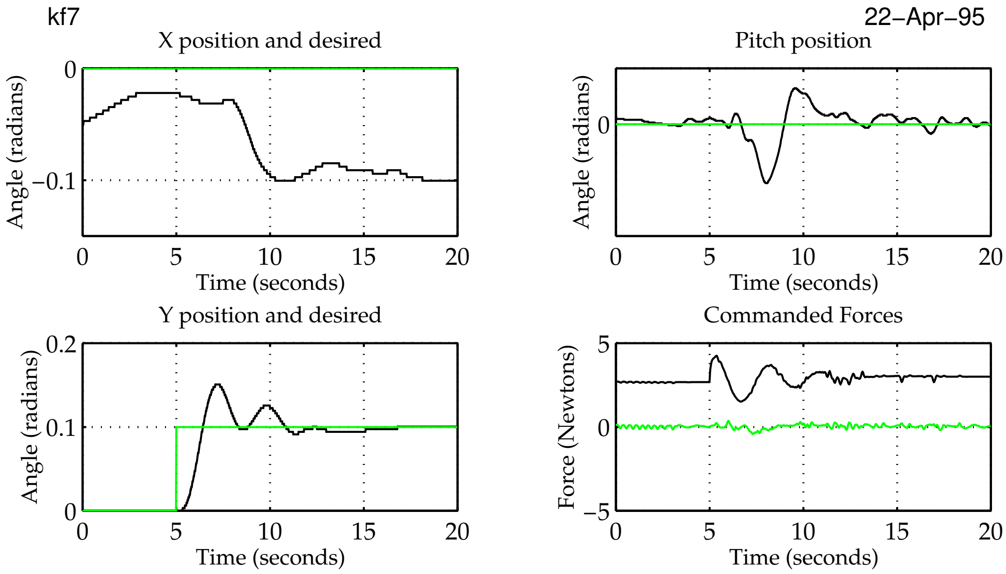

Figure 5.13: Closed loop response of the LPV controller locked in forward flight for an α_2 step.

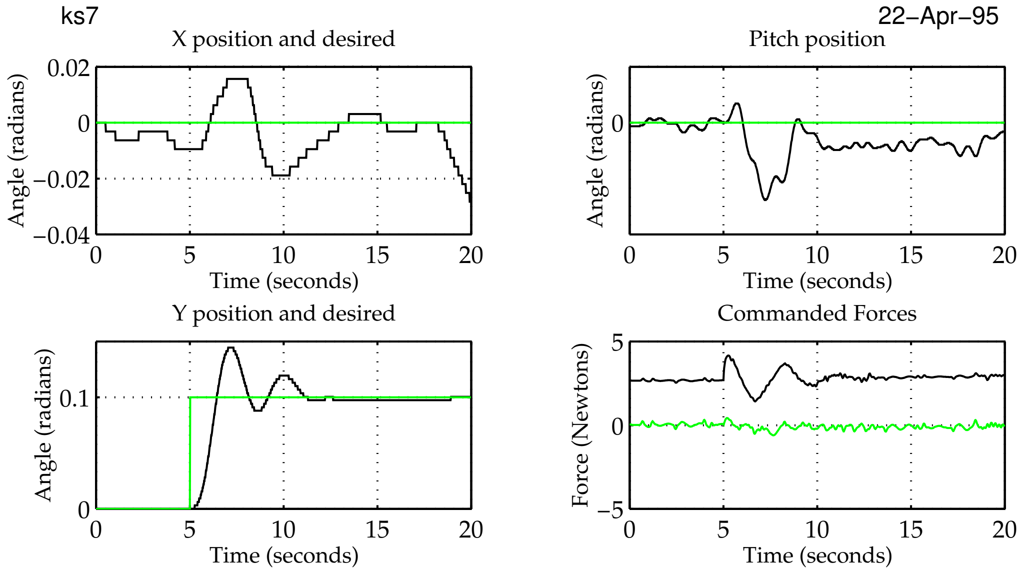

Figure 5.14: Closed loop response of the LPV controller locked at hover for an α_2 step.

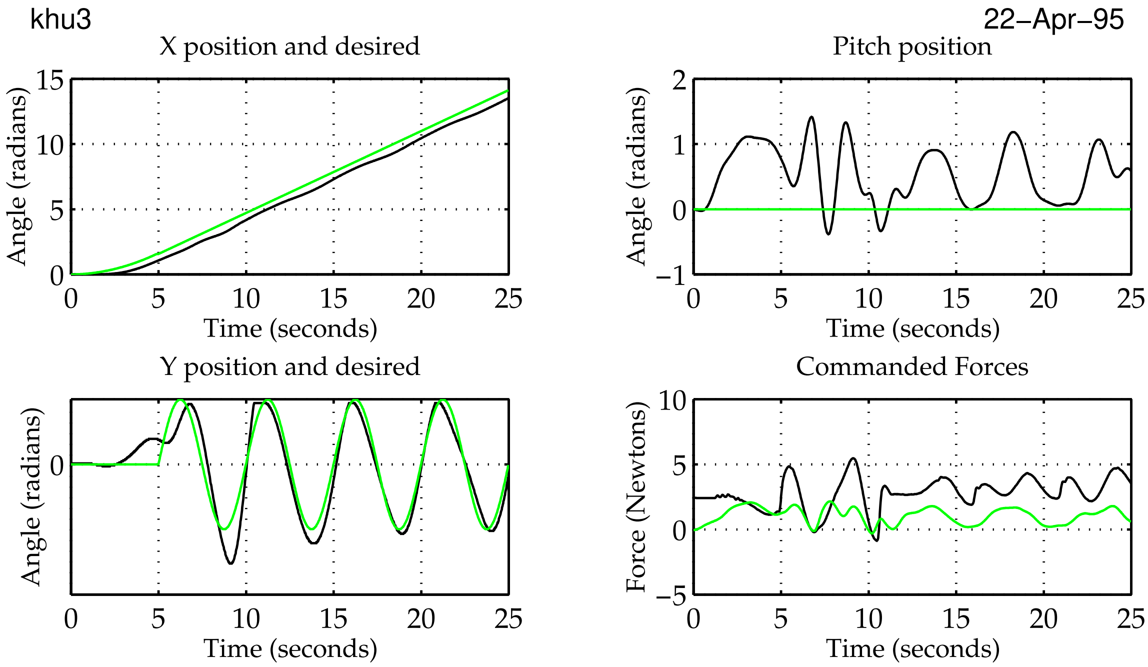

Figure 5.15: Closed loop response of the $\mathcal{H}_\infty$ controller for the first rolling trajectory.

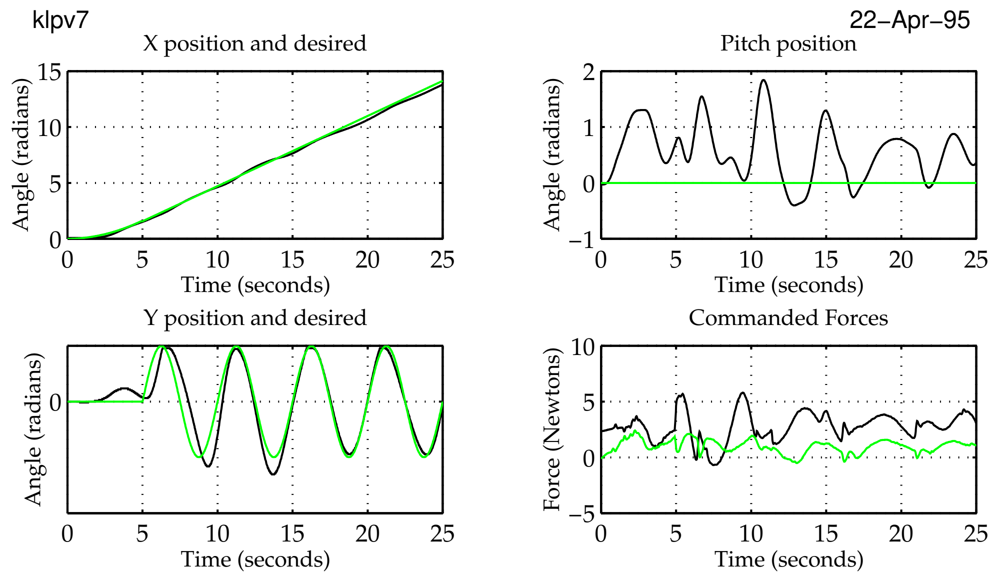

Figure 5.16: Closed loop response of the LPV controller for the first rolling trajectory.

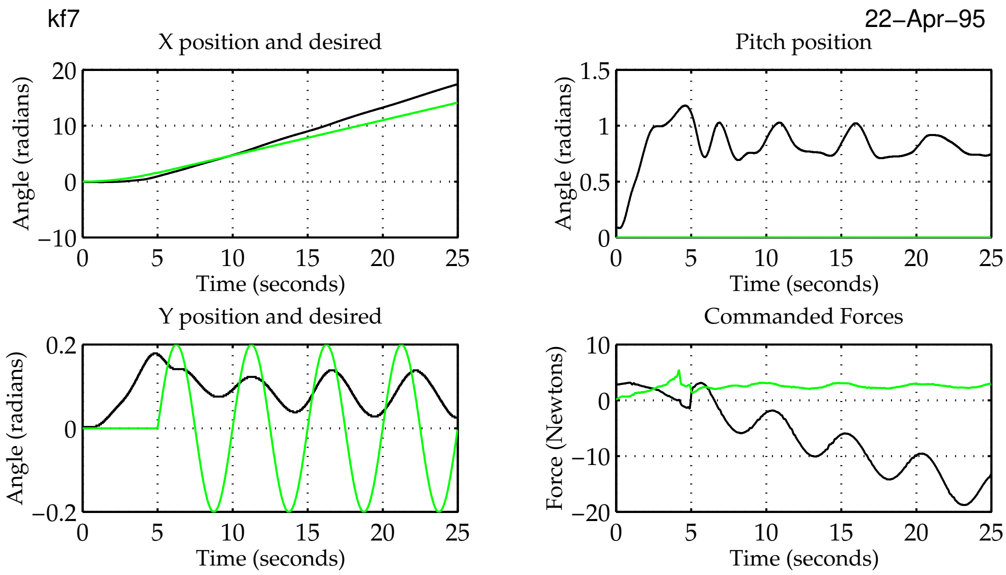

Figure 5.17: Closed loop response of the LPV controller locked in forward flight for the first rolling trajectory.

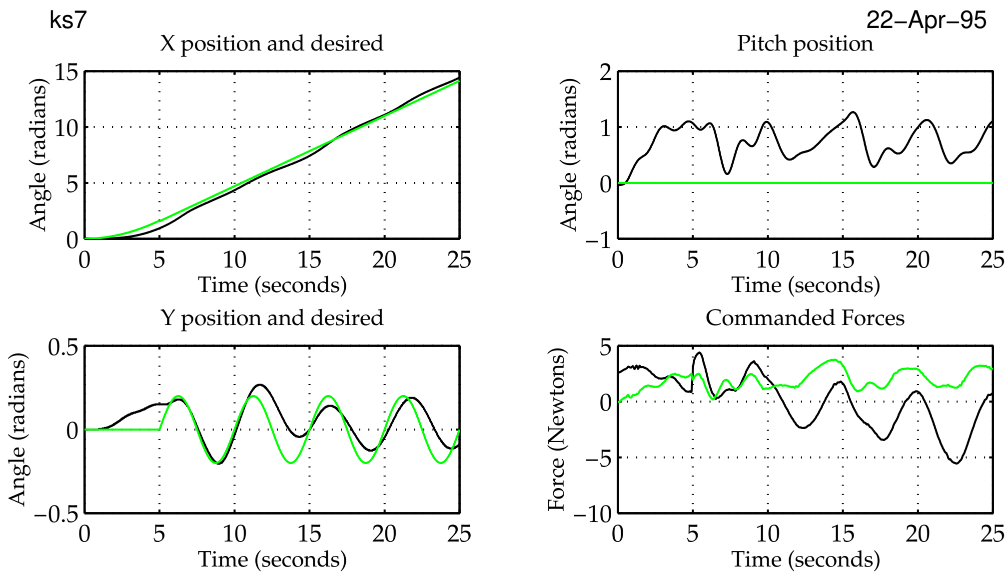

Figure 5.18: Closed loop response of the LPV controller locked at hover for the first rolling trajectory.

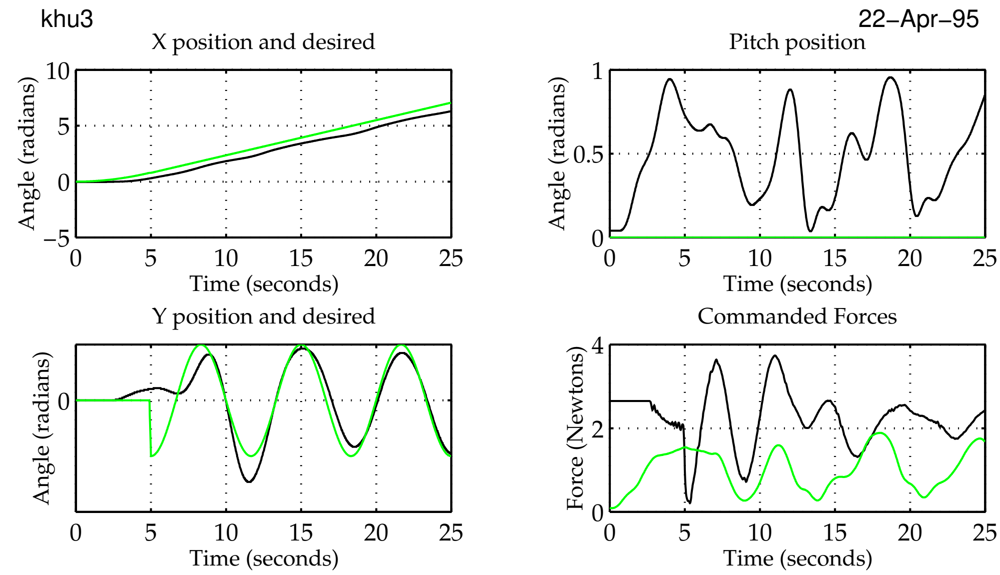

Figure 5.19: Closed loop response of the $\mathcal{H}_\infty$ controller for the second rolling trajectory.

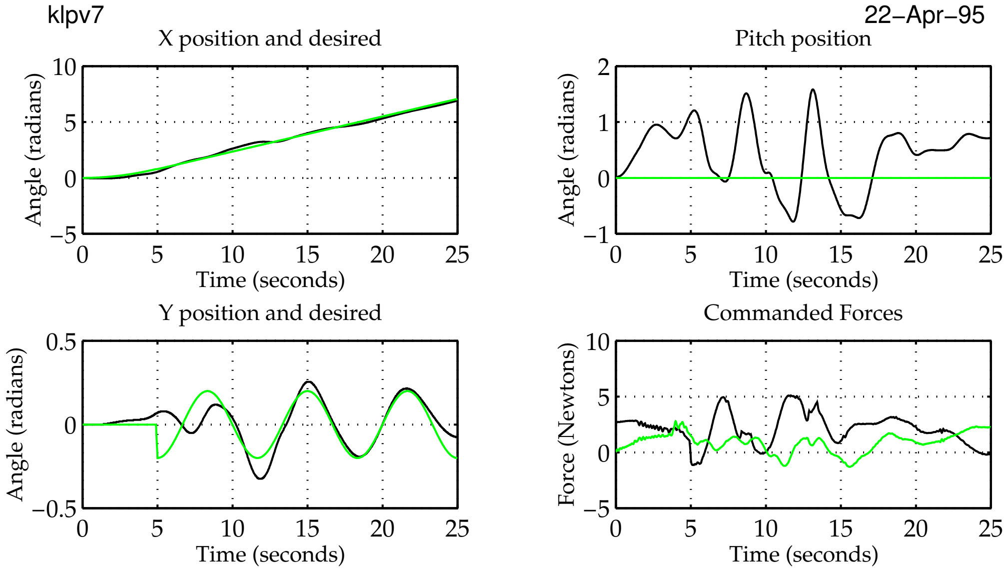

Figure 5.20: Closed loop response of the LPV controller for the second rolling trajectory.

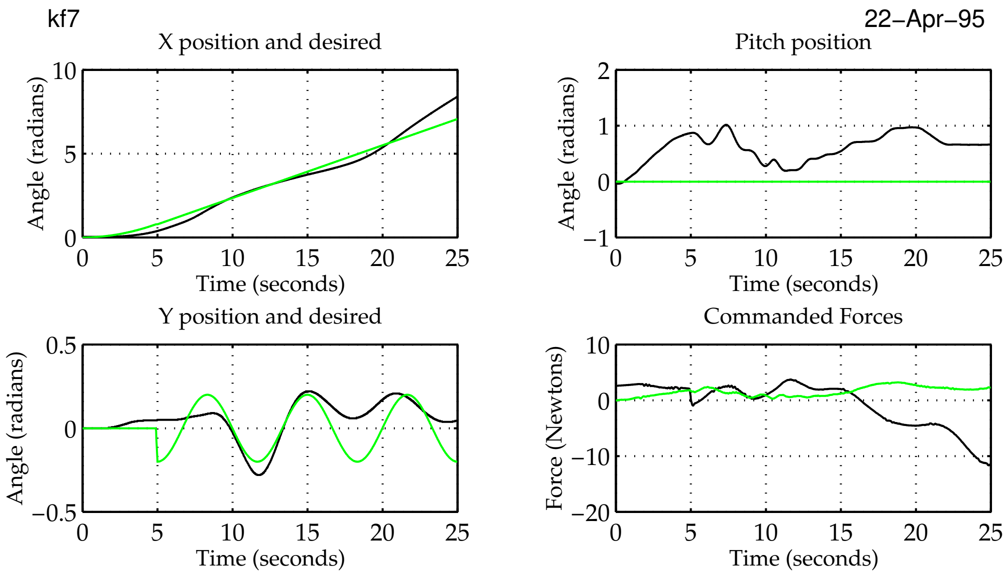

Figure 5.21: Closed loop response of the LPV controller locked in forward flight for the second rolling trajectory.

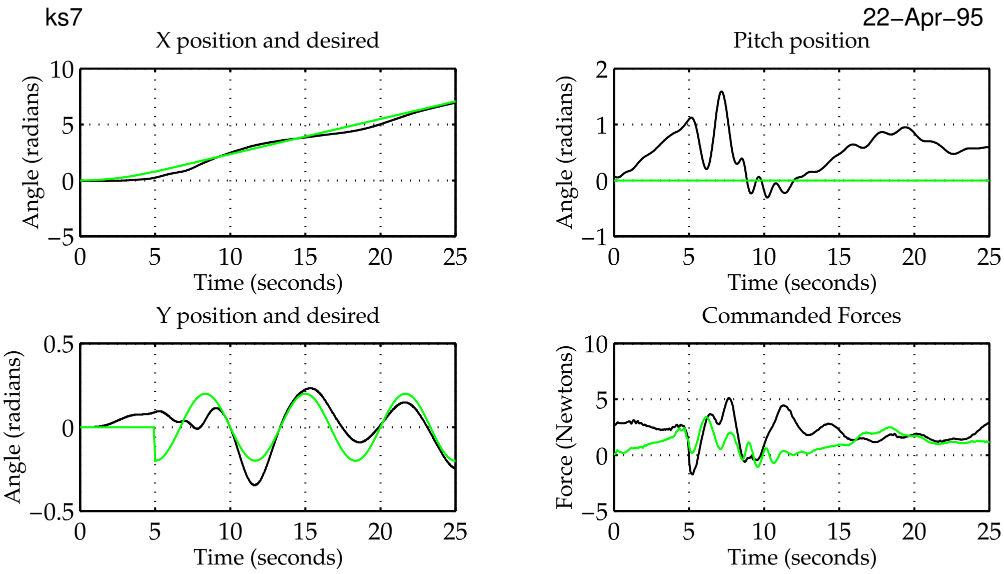

Figure 5.22: Closed loop response of the LPV controller locked at hover for the second rolling trajectory.

Chapter 6

Special Topics

Knowledge the shade of a shade,
Yet must thou sail after knowledge,
Knowing less than drugged beasts. phtheggometha thasson
$\phi\theta\epsilon\gamma\gamma\omega\mu\epsilon\theta\alpha \quad \theta\widehat{\alpha}\sigma\sigma o\nu$

—Ezra Pound

This chapter explores several interesting issues associated with the ducted fan. We present them here since presenting them earlier would have taken us too far afield from our discussion of controller design. In Section 6.1, identification experiments are performed on the fan to develop a linear model for the fan near hover which is based on the actual system rather than a first principles model. In Section 6.2, model validation techniques are employed on both the identified and linearized model and prove an excellent analysis tool for giving insight into the models. Finally, in Section 6.3, we present the first application of newly developed model reduction techniques [5, 6] to the ducted fan.

We will assume the reader has some familiarity with the basic concepts in these sections, as a detailed review of all the background material involved would be too burdensome.

6.1 Identifying a Model for the Ducted Fan

In this section we attempt to infer a linear model of the ducted fan from measured data. Based on this identification, we hope to construct a more accurate model of the ducted fan than previously available. The identification might, in turn, lead to an improved first principles model if we can deduce, for example, the friction and drag coefficients. The material in this section was a joint effort with Pascale Bendotti.

Review of the Identification Method

The identification method we will employ is a parameter estimation method. A complete discussion of the method can be found in [29], but see the parallel discussion in Section 3.3. Let $\varepsilon(t)$ denote the prediction-error, i.e.,

$$\varepsilon(t) = y(t) - G(\theta, z)u(t) \tag{6.1}$$

where θ is a vector of parameters to be identified, $y(t)$ is a measured series of data from the system when an input $u(t)$ is applied, and

$$G(\theta, z) = C(\theta)(zI - A(\theta))^{-1}B(\theta) + D(\theta) \tag{6.2}$$

is the transfer function obtained from state-space matrices $A(\theta)$, $B(\theta)$, $C(\theta)$, and $D(\theta)$, which have entries depending on the parameter vector.

The optimal parameter estimate is then given by

$$\widehat{\theta} = \arg \min_{\theta \in \mathcal{D}} \det \left[\frac{1}{N} \sum_{t=1}^{N} \varepsilon(t)\varepsilon^T(t) \right] \tag{6.3}$$

where $\mathcal{D}$ is the set of allowed parameter vectors, N is the length of the data record, and "arg min" means the minimizing argument of the function. Note, there is no guarantee that $\widehat{\theta}$ is unique. If it isn't we obtain a set of minimizing arguments.

Most of the effort in using this type of identification method is in defining the structure of the state-space matrices and in providing good initial values for the parameter vector. This is critical because the optimization is non-convex and may contain many local minima.

The structure we will employ, obtained after much trial and error, is one obtained using the first principles model linearized about hover as a starting point. Unfortunately, we have been unable to satisfactorily obtain an identified model using exactly the same structure (based on the ability of the resulting model to predict the data set) and our final structure is

$$[A|B] = \left[\begin{array}{cccccc|cc} 1 & 0 & 0 & T & 0 & 0 & 0 & 0 \\ 0 & \phi_1 & 0 & 0 & \phi_2 & 0 & \phi_3 & 0 \\ 0 & 0 & \phi_4 & 0 & 0 & \phi_5 & 0 & \phi_6 \\ 0 & 0 & \phi_7 & 1 & 0 & 0 & 0 & \phi_8 \\ 0 & \phi_9 & 0 & 0 & 0 & 0 & \phi_{10} & 0 \\ 0 & 0 & \phi_{11} & 0 & 0 & 0 & 0 & \phi_{12} \end{array} \right] \tag{6.4}$$

with $C = I_4$ and $D = 0$. Note in particular that this model only provides α_1, α_2, α_3, and $\dot{\alpha}_1$ as outputs. We were unable to obtain good convergence properties when $\dot{\alpha}_2$ and $\dot{\alpha}_3$ were included while keeping C an identity matrix. It may be possible to obtain results for these two signals if the constraint that $C = I$ is relaxed.

The data set used for the identification is shown in Figure 6.1. In this figure the measured signals are shown in solid lines while filtered signals are shown in shaded lines. Filtering is a practical necessity to aid the convergence

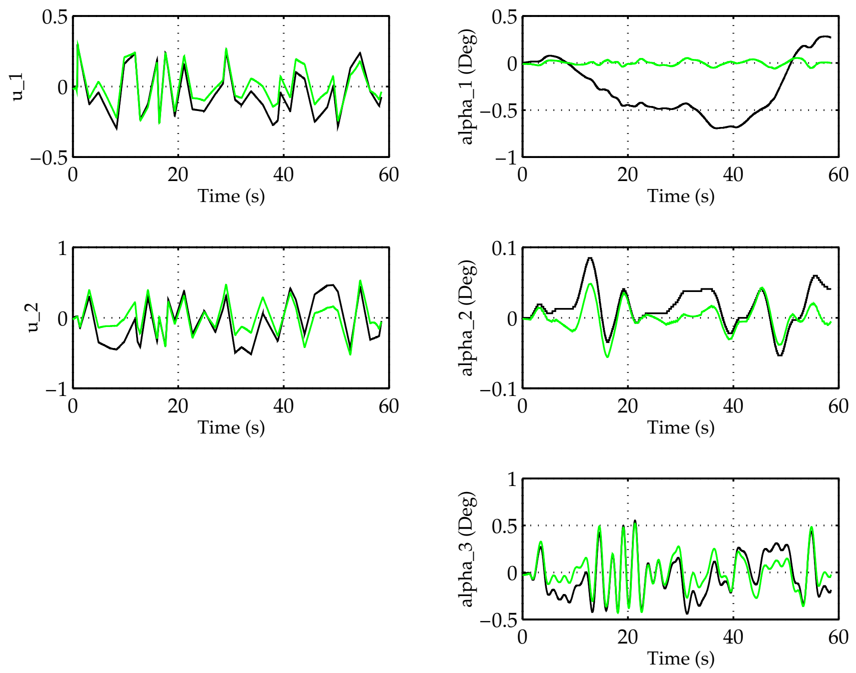

Figure 6.1: Filtered (shaded lines) and unfiltered (solid lines) signals for model identification.

of the estimation routine. The figure shows the inputs, u_1 and u_2, and the main outputs, α_1, α_2, and α_3.

The difficulty in obtaining a good model able to predict the data when using the structure of the linearized model is primarily associated with bad prediction of the α_1 channel and — to a lesser extent — the α_2 channel. In the linearized model both the channels have undamped modes, whereas significant damping is present on the real system. Moreover, α_1 in the undamped model is a double integrator ($1/s^2$). When damping is added this transfer function becomes $1/(s^2 + ks)$ which at low frequencies has a slope closer to $1/s$. Also, we are best able to identify the least damped mode in the system, α_3. We thus conclude that the linearized structure is inappropriate as a structure for identification of the damped system.

Using the structure of Equation 6.4, the identification procedure produces a model which is able to predict this specific data set with great accuracy, as shown in Figure 6.2. In this figure, the solid lines are the outputs shown in Figure 6.1, while the outputs produced by simulating the identified model with the same inputs shown previously are shown in shaded lines. The match is almost perfect on α_3, good on α_2, but not very good on α_1. As a comparison, a simulation with the linearized model is shown in Figure 6.3, where it is seen that the prediction is quite terrible. This is not unexpected, since the linearized model is undamped, but is still quite striking.

As an additional comparison between the identified and linearized models,

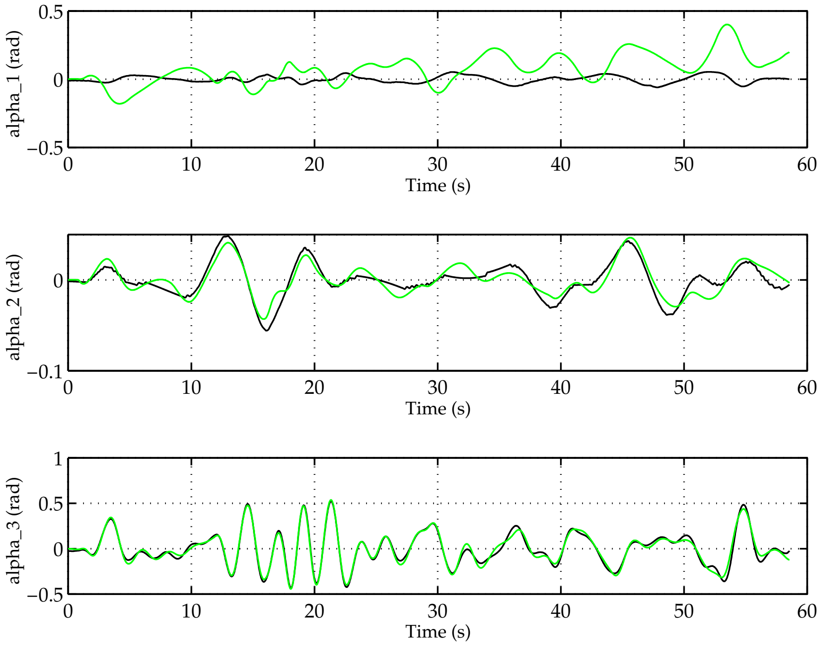

Figure 6.2: Prediction results by the identified model using the original data set. The solid lines are the predictions and the shaded lines the actual data.

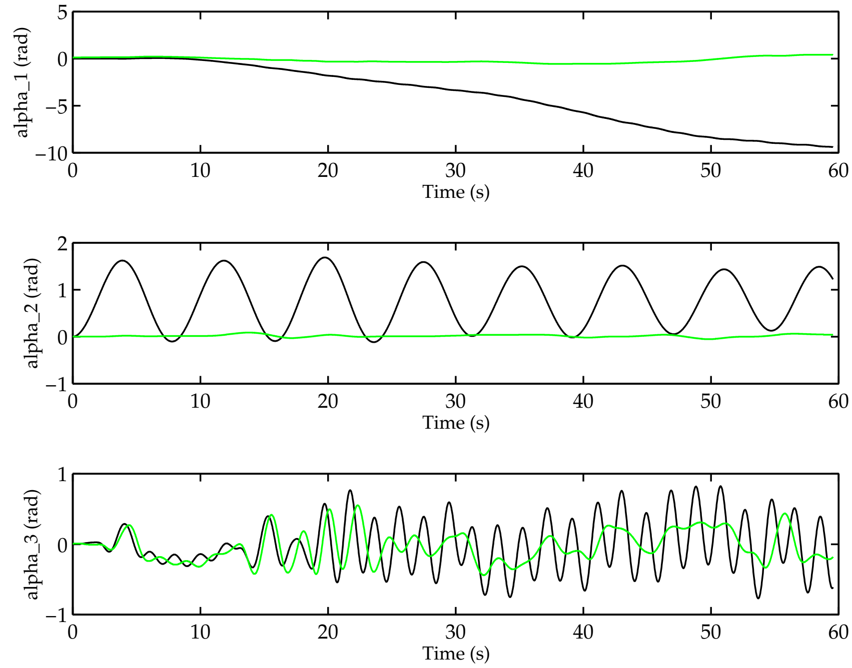

Figure 6.3: Prediction results by the linearized model using the original data. The solid lines are the predictions and the shaded lines the actual data.

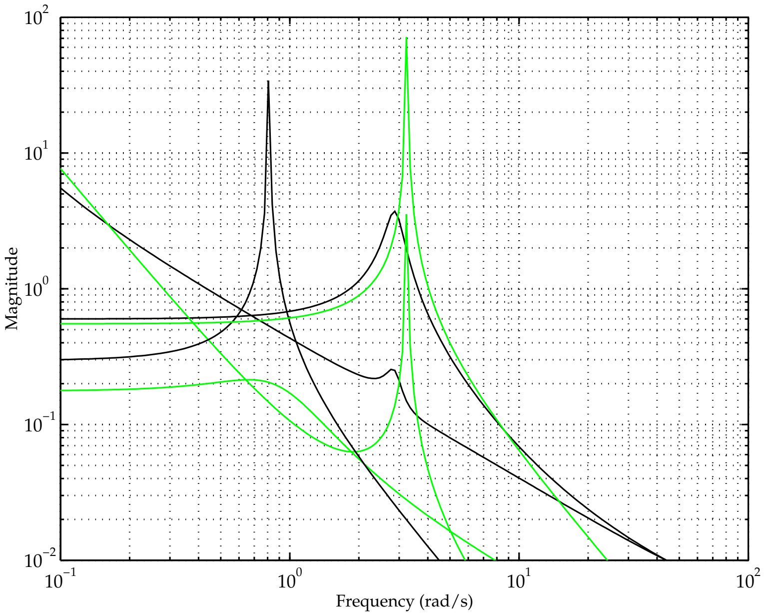

Figure 6.4: The Bode singular value plot of the identified and linearized models. The solid lines are the identified model and the shaded lines are the linearized model.

Figure 6.4 shows the Bode singular value plot for the α_1, α_2, and α_3 channels. The solid line is that of the identified model and the shaded line that of the linearized. The identified model is damped, particularly the α_2 channel, which is expected due to friction and other effects. Additionally, the frequencies are shifted slightly between the linearized and identified models. Since the values going into the linearized model were obtained solely by measurement of physical parameters, this also is unsurprising.

To confirm that the identified model is reasonable when given another data set, a different data set was simulated. The results are presented in Figure 6.5. This figure shows the same general behavior as on the previous data set, and although not quite as good, the controller is still able to predict the behavior of α_2 and α_3 fairly well. Additionally, an $\mathcal{H}_\infty$ controller was designed for the ducted fan using this model and successfully implemented.

Summary

In this section we have identified a superior linear model of the ducted fan. We are able to identify a very good linear model of α_3, and a reasonably good model of α_2. The model for the α_1 channel needs improvement. In the next section, we will attempt to analyze this model further in the context of determining whether a given uncertainty structure is adequate to predict

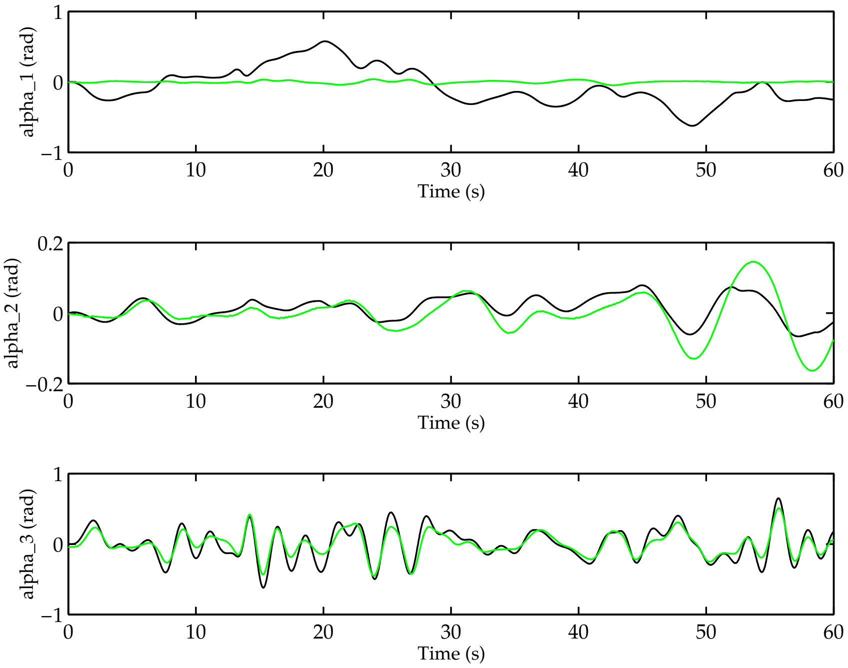

Figure 6.5: Prediction results for a different data set by the identified model. The solid lines are the actual data and the shaded lines the prediction.

observed behavior.

6.2 An Application of Model Validation

In this section we discuss an application of model validation to the ducted fan. We will see that model validation provides valuable information into the refinement of both the linearized and identified models discussed previously. Based on this information, we propose a method to develop future uncertainty models for the ducted fan and critique where both the linearized and identified models should be improved. This work arose out of a collaborative effort with John Morris and Matt Newlin.

Introduction to Model Validation

Model validation is concerned with determining whether a mathematical model is consistent with (or covers) a collection of experimental data. The constant matrix version of this problem has been formulated as a generalization of the structured singular value μ_g in [37]. The main result is that when the norms of the noise, disturbances, and uncertainty are of size less than or equal to $1/\gamma$, the data and uncertain model are consistent if $\mu_g \geq \gamma$. As is the case for μ, μ_g cannot in general be exactly computed and thus upper and lower bounds are employed. An upper bound for μ_g using LMIs was developed in [37], and a general algorithm has been developed to compute this upper bound [33]. The upper bound is tight in the sense that equality between the upper bound and μ_g is achieved for certain classes of uncertainty [37]. In the following, we will denote this upper bound of μ_g by γ_{ub}. Additionally, a modified power algorithm was developed for a lower bound of μ_g which we will denote by γ_{lb}.

To make use of the constant matrix results with experimental data and a robust model involves two steps. Figure 6.6 shows the general interconnection structure of a robust model for a model validation problem. Here d is the exogenous disturbance, u is the control input, and n is the noise input. Δ is an element of a prescribed uncertainty set $\mathbf{\Delta}$ which is assumed to consist of LTI perturbations. W_n is a weight on the noise signal n.

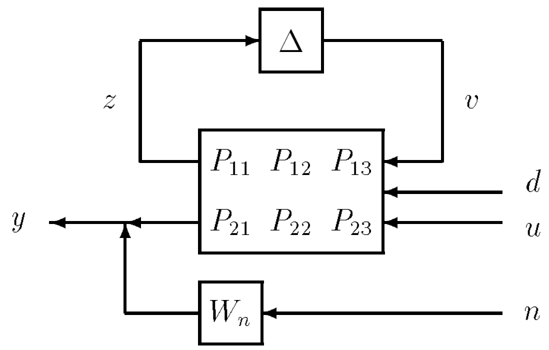

Figure 6.6: Block diagram for the model validation problem.

The best software currently available for solving this problem performs a frequency domain analysis; thus, any time-domain data must be transformed into the frequency domain. Then a sweep over frequency is performed. At each frequency point the robust model and the frequency domain data are reduced to a constant matrix problem. It is not necessary to utilize all available data in the frequency sweep as a subset of the data corresponding to the important frequency points can be used instead, similar to μ-analysis and μ-synthesis, [20].

Given a plant P and frequency-domain input/output data, u and y, the constant matrix bounds are computed over Ω, a subset of frequency points in the FFT of u and y. Define the following bounds across all frequencies:

$$\mu_g^\star(P) = \min_{\omega\in\Omega} \mu_g(P(\omega)) \tag{6.5}$$

$$\gamma_{lb}^\star(P) = \min_{\omega\in\Omega} \gamma_{lb}(P(\omega)) \tag{6.6}$$

$$\gamma_{ub}^\star(P) = \min_{\omega\in\Omega} \gamma_{ub}(P(\omega)). \tag{6.7}$$

Then the following is true:

$$\gamma_{lb}^{\star}(P) \leq \mu_g^{\star}(P) \leq \gamma_{ub}^{\star}(P). \tag{6.8}$$

The important result is that for all $\gamma < \gamma_{lb}^{\star}$ the robust model and the data are γ-consistent on the frequency set Ω. The following defines what is meant by the statement "the robust model and the data are consistent," and particularly when they are consistent for a specific value of γ.

> **Definition 1** *A robust model and data are* γ-consistent *if* $\exists \Delta \in \mathbf{\Delta}$, $\|\Delta\| \leq 1/\gamma$, $\|d\| \leq 1/\gamma$ *and* $\|n\| \leq 1/\gamma$ *such that* $y = W_n n + (\Delta \star P)\,[d' u']'$, *where* P, W_n, d, n, u, *and* y *are as in Figure 6.6, with* u *and* y *experimental data,* d *exogenous disturbances, and* n *measurement noise.* $\mathbf{\Delta}$ *describes the set of structured uncertainty.*

As an aside, note that in the synthesis of a robust controller the μ-analysis results for the same block structure as used in model validation should be smaller than $\mu_{lb}^{\star}$. When this occurs, the controller will be robust to the disturbances, noises and uncertainties necessary for the model to be consistent with the measured data. Simply put, if this is the case, then the model is consistent with the data, and the controller is consistent with the model.

Models and Generation of Signals

The current section concerns itself exclusively with linear models for the ducted fan around hover. The "linearized model" is the nonlinear model of the ducted fan presented in Chapter 4 and linearized about hover. Recall that this is the model from which the $\mathcal{H}_\infty$ controller was designed. In particular, the uncertainty weight we will use for model validation is the weight W_m depicted in Figure 4.6. We will also use the identified model for the ducted fan discussed in the previous section. Note that the uncertainty weight was developed in conjunction with the linearized model and may not be as well-suited for controller design with the identified model.

To generate a data set to employ for model validation a zero mean random signal was constructed around the equilibrium hover point caused by a force pair of $(2.65, 0)$N. The input data set is shown in Figure 6.7. The output from the ducted fan for the α_1, α_2, and α_3 channels is shown in Figure 6.8.

Since we must perform the model validation tests in the frequency domain, we transform our finite length discrete time data to the frequency domain with the discrete Fourier transform (DFT). Recall that an implicit assumption in the use of the DFT is that the time domain data repeats periodically forever. Thus, high quality frequency domain data requires the use of time domain data that looks as though it could repeat periodically. Consequently, the test input signals end with a quiescent period, so that the system might return to the initial rest state. This strategy works well for the states associated

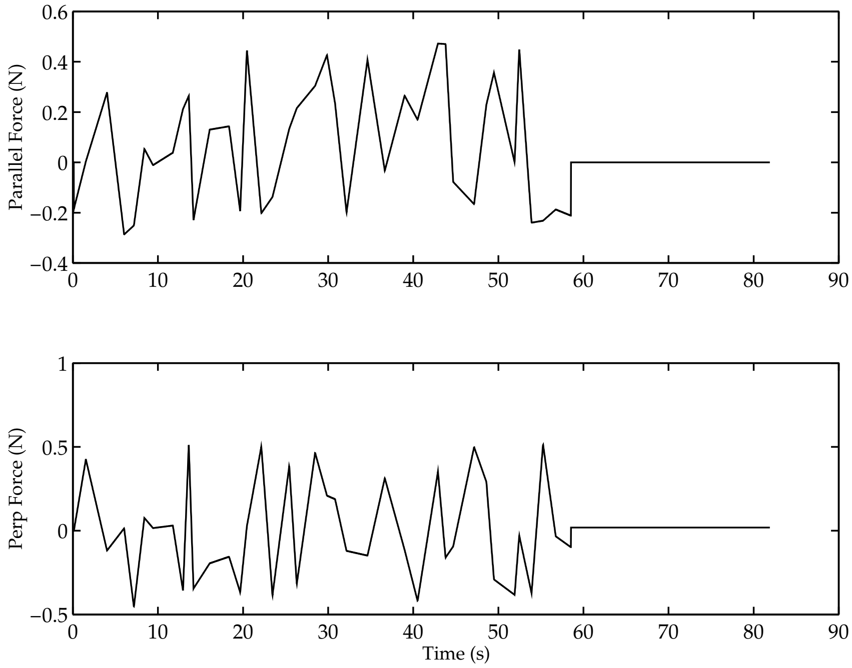

Figure 6.7: A set of random inputs to the ducted fan. The top graph is u_1 without the feedforward force of 2.65N and the graph is u_2.

with α_2, α_3, and $\dot{\alpha}_1$, but not with α_1, which exhibits a slow drift, caused by nonlinearities in the model.

As we aren't concerned with the system behavior at very low frequencies, this α_1 drift is of little concern except that it might corrupt the data at other, more interesting, frequencies. This corruption is due to the frequency content in the step transition from the end of the data record to the beginning of the record as it repeats periodically.

To minimize this effect the α_1 data set is filtered with a fourth-order acausal high-pass filter with a cut-off frequency of 0.25 rad/s. The filter was chosen to cause minimal phase distortion over the frequency range of interest while making the processed time history appear suitable for the DFT. Note that the filtering corrupts the data below 0.25 rad/s.

Transfer Function Data

Model validation acts on frequency domain data. Representative plots of the transfer function data are now presented. The particular example we show is for the identified model. Shown are the transfer functions from $(u_1, u_2) \longrightarrow (\alpha_1, \alpha_2, \alpha_3)$. The remaining transfer function data is shown in the Appendix.

In Figures 6.9 through 6.14, the solid line shows the FFT of the measured output data divided by the FFT of the input data. An 'x' in the figure shows

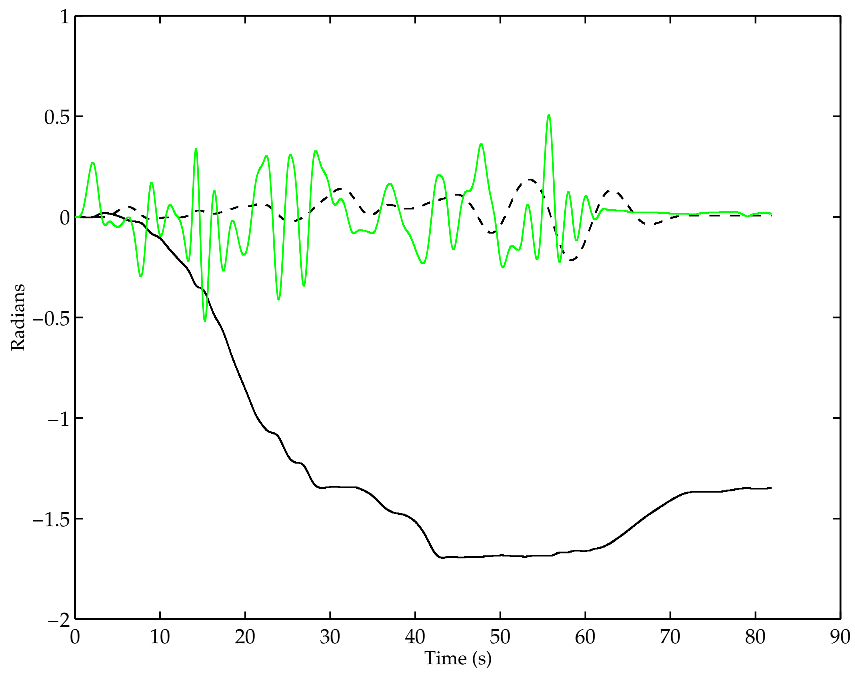

Figure 6.8: The outputs from the ducted fan. The solid line is α_1, the dashed is α_2 and the shaded is α_3.

a point at which model validation will be done. The shaded line shows the FFT of the output data generated by simulating the identified model with the input data used to generate the real output data.

Note that these plots are not sufficient to predict what model validation will calculate, since phase information is important to the model validation process.

Model Validation Results

In this section, the results of model validation are presented. A very important point to note when interpreting the results below is that we are *not* using the results to verify the robustness guarantees of a closed-loop system. This is the standard way of considering model validation, so readers familiar with the subject should be wary of falling into this habit of thinking in this work. Instead, we are using model validation to determine if a model and uncertainty description can capture the dynamic structure of a true system.

Because $\mathcal{H}_\infty$ synthesis, when posed as an output tracking problem, makes no distinction between a command input and a noise input, we have little empirical feel for the noise weight required by the validation procedure. Preliminary model validation data was used to iterate on the magnitude of W_n; the final choice for W_n was $W_n^{mv} = 0.05I$. W_m was unchanged. These preliminary results seemed to indicate that the problem is not particularly sensitive

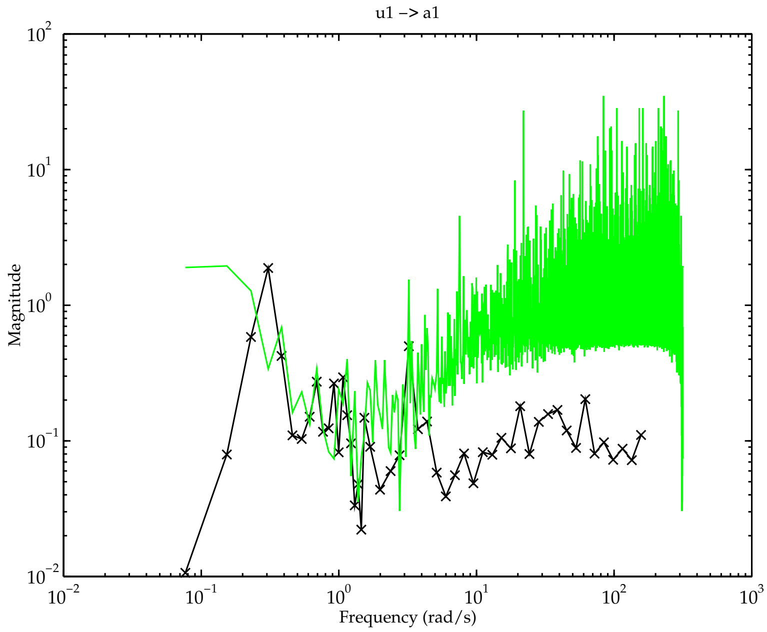

Figure 6.9: Transfer function from u_1 to α_1 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

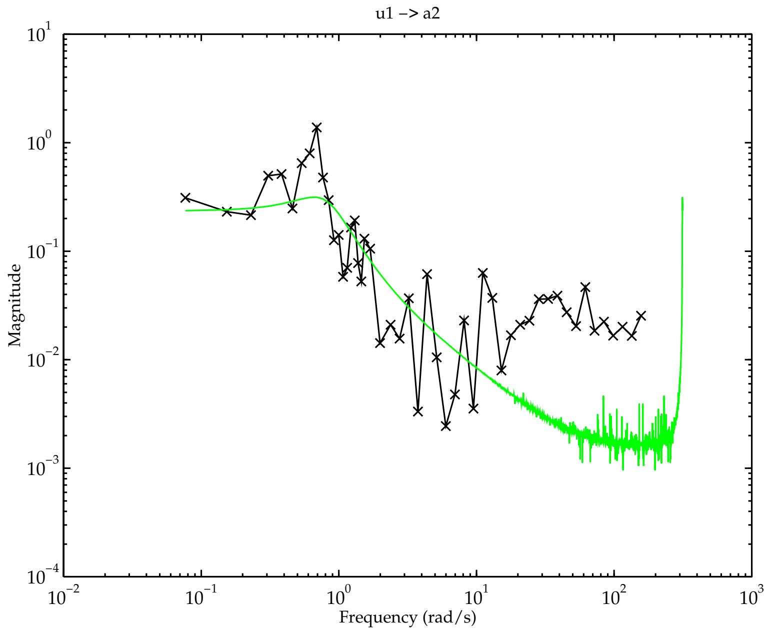

Figure 6.10: Transfer function from u_1 to α_2 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

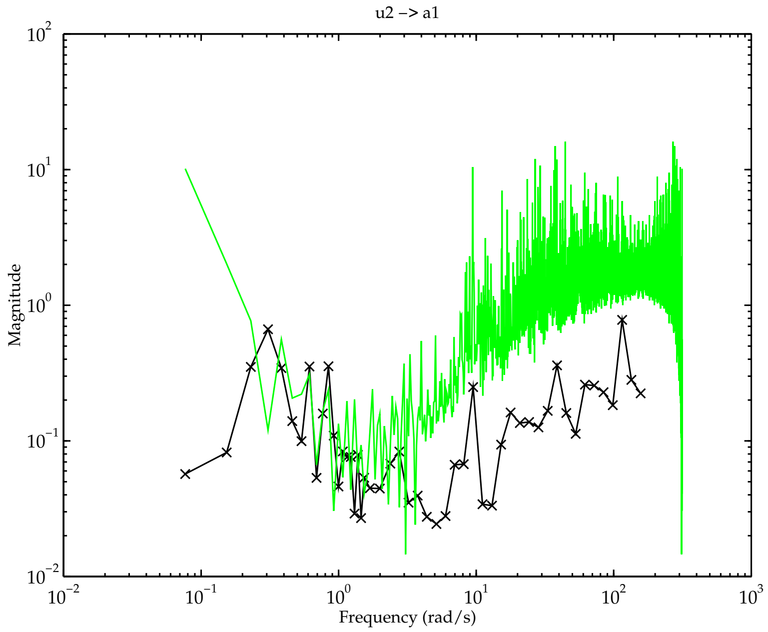

Figure 6.11: Transfer function from u_2 to α_1 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

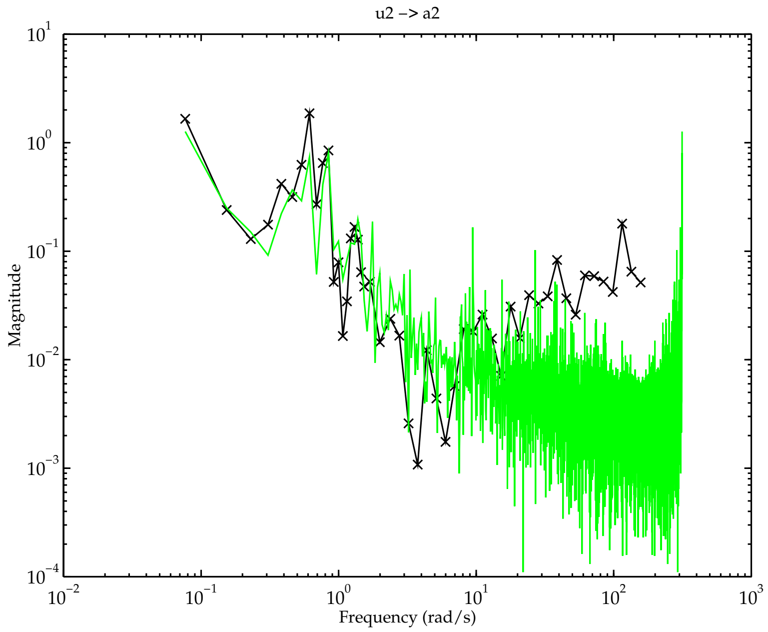

Figure 6.12: Transfer function from u_2 to α_2 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

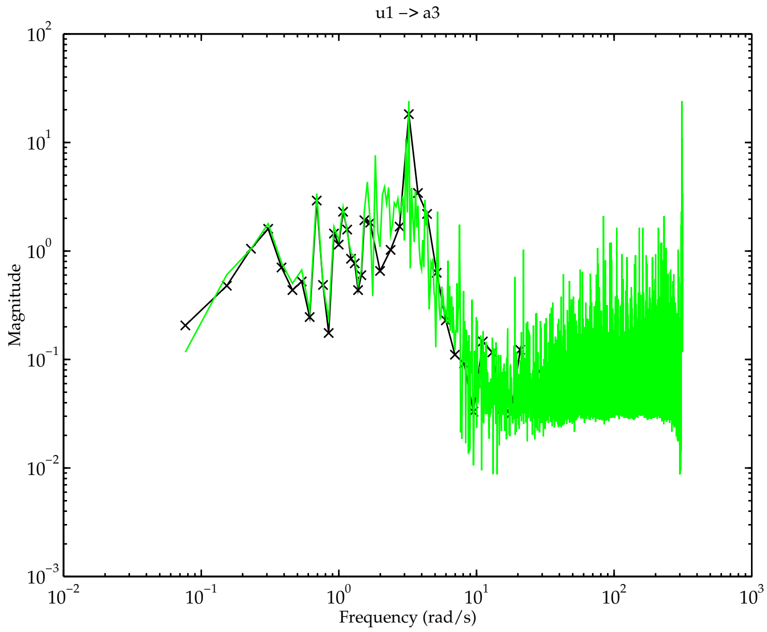

Figure 6.13: Transfer function from u_1 to α_3 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

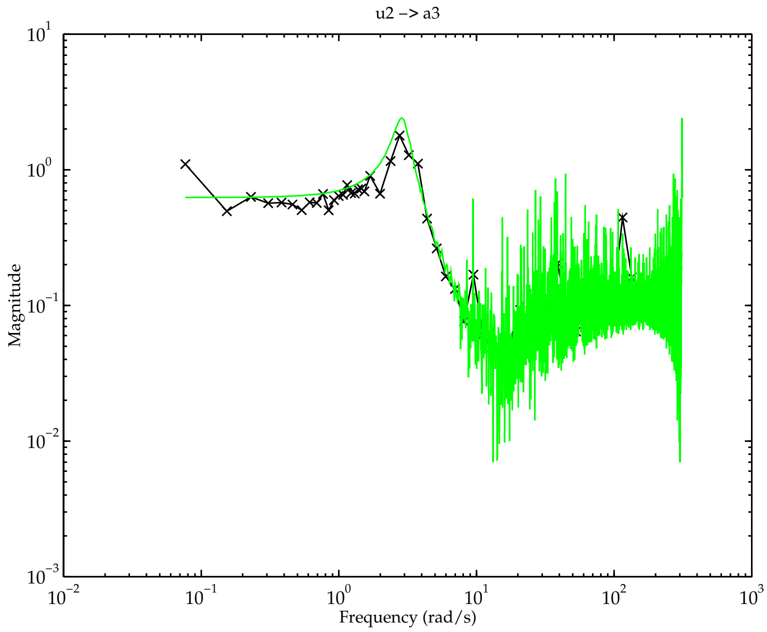

Figure 6.14: Transfer function from u_2 to α_3 from measured data (solid with an 'x' at the actual data point) and simulated with identified model (shaded).

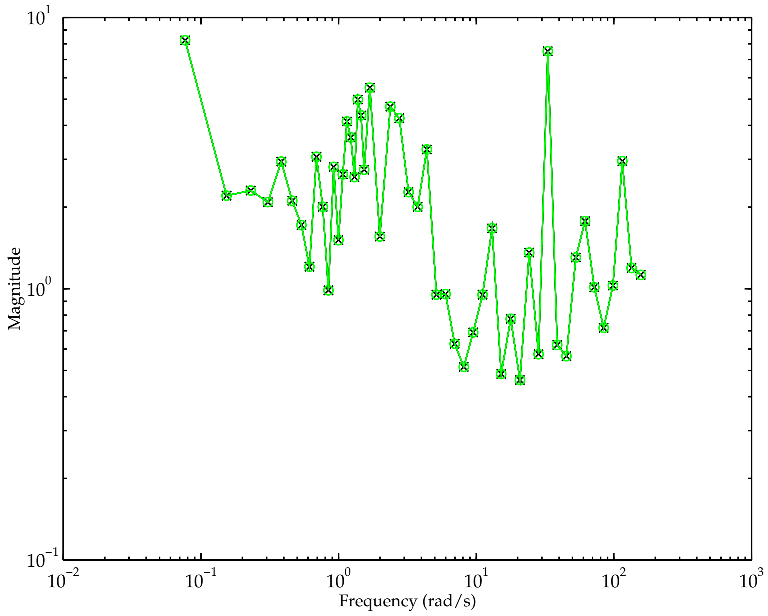

Figure 6.15: Model validation bounds for the identified model on α_1.

to the choice of W_n, so it was set to a low value.

Figure 6.15 shows that the MISO transfer function, $(u_1, u_2) \longrightarrow \alpha_1$, is consistent for $\gamma \approx 0.5$. Over most of the frequency range, the bounds are much better, but they drop around 10 rad/s. This is interesting, as it is after the mode shown in Figure 4.6. Figure 6.16 shows that the MIMO transfer function $(u_1, u_2) \longrightarrow (\alpha_2, \alpha_3)$ is consistent for $\gamma \approx 1.5$.

The previous two results might lead us to believe that the complete MIMO model is γ-consistent for $\gamma \approx 0.5$. Figure 6.17 shows the bounds for the MIMO transfer function $(u_1, u_2) \longrightarrow (\alpha_1, \alpha_2, \alpha_3)$ and demonstrates clearly that this is not the case. In fact, the model and data only become γ-consistent at high frequency, where the uncertainty weight becomes very large. This implies that either there are unmodelled dynamics in the ducted fan that have a significant effect on the α_1 channel, or that the specific uncertainty required to obtain model and data consistency is different for the α_1 channel than the (α_2, α_3) channels. We will address this question further below.

Appendix 6.B presents additional model validation runs on sets of different channels. These results are tabulated in Table 6.1. The "transfer function" column shows the output channels which were selected for a particular validation computation; the inputs were always (u_1, u_2). The second and third columns show the lower and upper bounds, $\gamma^\star_{lb}$ and $\gamma^\star_{ub}$, for the identified and linearized models, respectively. For the single channel case, the value of μ_g can be computed and that result is given.

Finally, the fourth column shows model validation results where the data

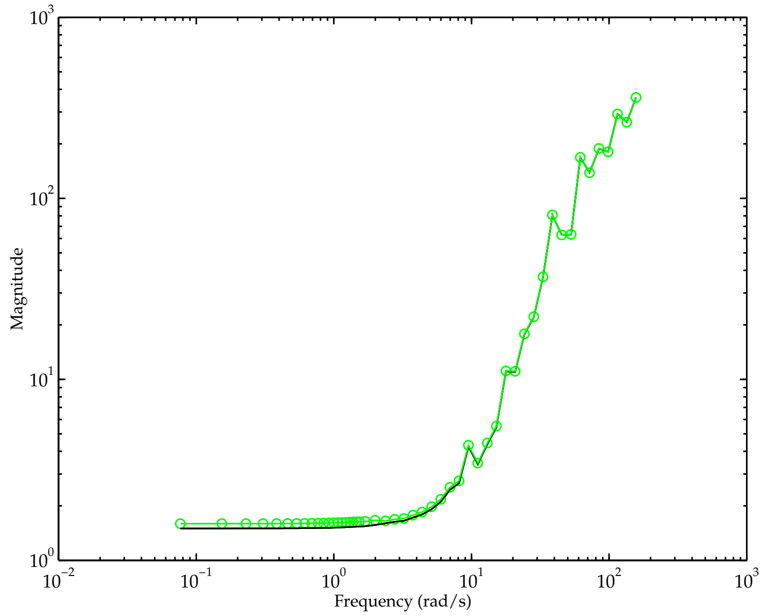

Figure 6.16: Model validation bounds for the identified model on α_2 and α_3.

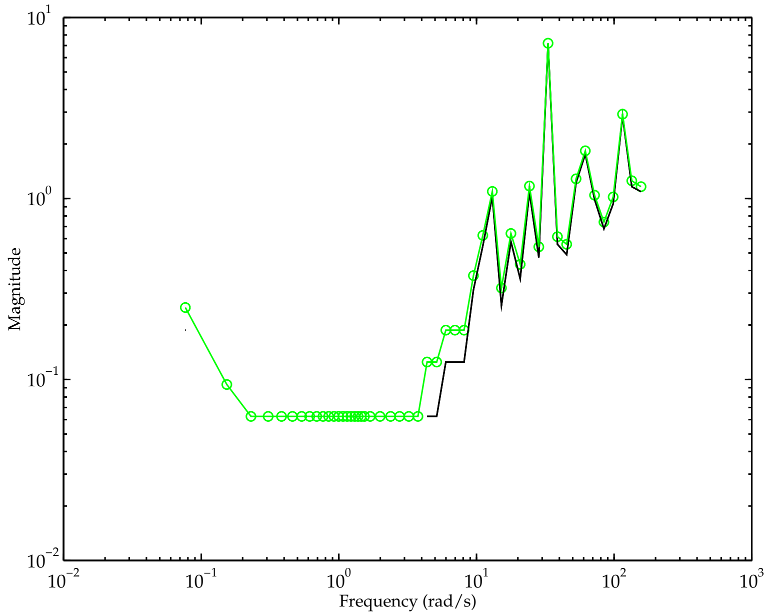

Figure 6.17: Model validation bounds for the identified model on α_1, α_2 and α_3.

Transfer Function	Identified Model		Linearized Model		Linearized Model, simulated data	
	$\gamma^\star_{lb}$	$\gamma^\star_{ub}$	$\gamma^\star_{lb}$	$\gamma^\star_{ub}$	$\gamma^\star_{lb}$	$\gamma^\star_{ub}$
α_1	0.4613		0.3269		0.2248	
(α_2, α_3)	1.5011	1.5949	1.5009	1.5947	1.3592	1.3837
$(\alpha_1, \alpha_2, \alpha_3)$	0	0.0625	0	0.0625	0	0.0547
$\dot{\alpha}_1$	1.5386		1.5374		1.4347	
α_3	1.5361		1.5361		1.3765	
α_2	1.5271		1.5269		1.4025	
(α_1, α_3)	0	0.0625	0	0.0625	0	0.0625

Table 6.1: A table showing the results of various model validation results. Where the upper and lower bounds are theoretically equal, the value of μ_g is given.

y was not obtained from an actual experiment but from applying the input u shown in Figure 6.7 to a nonlinear simulation. Since the nonlinear simulation is undamped, the resulting output was then windowed in the time domain before being transformed into the frequency domain. Thus, these results do not compare with the other columns. Examining them shows that the results are very similar to the ones for the identified and linearized models on real data, which allows us to conclude that the discrepancy between the model validation results for (α_2, α_3) and $(\alpha_1, \alpha_2, \alpha_3)$ are not due to unmodelled dynamics (as there are none in the nonlinear simulation), but that it must be due to the uncertainty structure not being representative.

Conclusions

From the results shown in the previous section, it is clear that the present uncertainty description, that of a lumped multiplicative uncertainty at the input, is insufficient to account for the data in all three channels simultaneously. The problem channel appears to be α_1, since a model validation on α_2 and α_3 shows those channels work well together, but neither works well with α_1. Note, though, that a model validation on α_1 alone shows that the model is reasonable. This is a striking example, in a practical application, of why thinking about MIMO systems in a "loop-at-a-time" framework is incorrect. To improve the model, a different uncertainty description is needed to account for this difference.

A possible explanation for the difference between α_1 and the other data is that α_1 is essentially an integral times a gain of α_3. If this gain were uncertain, or varied, then the resulting signal could be out of phase with α_3. From the dynamics of the stand and from data taken from the force-torque sensors for computing the force-lookup table of the fan, we know that the parameter r, which is the distance from the center of rotation of the fan to the point at which the forces act, can vary by as much as 20% during normal operation, and that this indeed might cause the gain variation. In any event, further investigation of these two facets of the ducted fan model is needed.

In comparing the results for the identified and linearized models from Appendix 6.B, we see that on every test, the identified model is better than the linearized model. This coincides well with the prediction results given in Section 6.1. This is particularly interesting since the uncertainty structure was developed, by trial and error, to account for inaccuracies in the linearized model.

For linear systems, the frequency domain provides a great simplification of dynamical system representation in that convolution in the time domain becomes multiplication in the frequency domain. Thus data which is coupled across time is decoupled in frequency, and the associated robustness analysis and model validation problems are greatly simplified by this decoupling. While the frequency domain is natural for continuous time infinite horizon data, it is unnatural for discrete time finite horizon data, where an implicit assumption in the transformation is that the time domain data repeats periodically forever. Frequently, through careful experiment design, one may collect data that appears fairly consistent with this periodicity assumption, and the errors induced by going to the frequency domain are reasonably small. Often, however, it is not possible to collect data that seems suitable for transformation. In these cases, it would be a great benefit to perform the model validation computation in the time domain. Although this is no harder conceptually, the coupling of the problems from one time to the next makes the computation much more expensive. Effective computation of such problems is an area of current research.

In summary, the model validation has proven to be a capable tool for indicating improvements in our models of the ducted fan which would have been difficult to ascertain otherwise. In particular, better models should incorporate either a different type of uncertainty or more structure in the uncertainty.

6.3 Stabilizability Issues

In this section we discuss stabilizability issues associated with the ducted fan model when we move into the operator LPV framework. This work arose while investigating the applicability of model reduction methods of Beck [5] to the ducted fan. The basic work on stabilizability of systems in the operator framework was done by Lu [30].

Introduction

In this section, let $\mathbf{\Delta}$ refer to the set

$$\{\mathrm{diag}\,[\delta_1 I_{n_1}, \ldots, \delta_s I_{n_s}] : \delta_i \in \mathcal{L}(\ell_2)\},$$

and let P be the appropriately partitioned system realization matrix

$$P = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right].$$

In Section 2.2 we defined when the system $\Delta \star P$ was stable, stabilizable, and detectable.

Parameter-Varying Models of the Ducted Fan

Recall from Chapter 4 that the basic parameter-varying model of the ducted fan is

$$\left[\begin{array}{c|c} A & B \end{array}\right] = \left[\begin{array}{cccccc|cc} 1 & 0 & 0 & T & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & T & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & T & 0 & 0 \\ 0 & a_{42}(\alpha_3, \dot{\alpha}_1) & a_{43}(\alpha_3) & 1 & 0 & 0 & b_{41}(\alpha_3) & b_{42}(\alpha_3) \\ 0 & a_{52}(\dot{\alpha}_1) & a_{53}(\alpha_3) & 0 & 1 & 0 & b_{51}(\alpha_3) & b_{52}(\alpha_3) \\ 0 & a_{62}(\alpha_3, \dot{\alpha}_1) & a_{63}(\alpha_3) & 0 & 0 & 1 & b_{61}(\alpha_3) & b_{62}(\alpha_3) \end{array}\right] \qquad (6.9)$$

where T is the sampling rate. Associated with the model is a block structure, as the parameter variations are extracted into an LFT. Our block structure contains $\delta_0 = z^{-1}$, $\delta_1 = \alpha_3$, and $\delta_2 = \dot{\alpha}_1$, and has the form

$$\mathbf{\Delta} = \{\mathrm{diag}\,[\delta_0 I_{n_0}, \delta_1 I_{n_1}, \delta_2 I_{n_2}]\}.$$

The simplified model for design, $(P_{des}, \mathbf{\Delta}_{des})$ has $n_0 = 6$, $n_1 = 2$, and $n_2 = 2$. Recall also that the parameter variations corresponded to a change in α_3 from 0 to 1.5 radians, and in $\dot{\alpha}_1$from 0 to 1.5 rad/s.

In the following discussion we will refer to a model as being unstable or not stabilizable. As it turns out, all the models considered here are detectable. The instability of a model is not difficult to determine, since there are well-known conditions for it, both necessary and sufficient. To say that a model is not stabilizable is a much more difficult assertion to prove, since we must show that there does *not* exist a positive definite matrix X satisfying Equations 2.4. In general for an LPV system, we don't know how to do this. Thus when we say that a model is not stabilizable, we will mean that using the best numerical optimization routines available to us, we were unable to find a feasible solution to the stabilizability LMI. Recall that the LPV synthesis method requires the model to be stabilizable to compute solutions to the LMIs.

If parameter-variations from $(P_{des}, \mathbf{\Delta}_{des})$ are allowed to correspond to a full-range of operation for the ducted fan, for example from -1.5 to 1.5 radians for α_3 and -1.5 to 1.5 radians/second for $\dot{\alpha}_1$, then the model is not stabilizable. Limiting the parameter variations from -0.2 to 0.2 radians and radians/second does result in a stabilizable model. Since our trajectories typically drive the fan such that α_3 and $\dot{\alpha}_1$are positive and since the gyroscopic effects of the fan blade, which are completely neglected in this model, become significant in the opposite direction, the reason for our choice of parameter ranges is apparent.

Now consider parameter variations corresponding to a range of 0 to 1.5. If every parameter variation in Equation 6.9 is considered, the resulting $\mathbf{\Delta}$ set has $n_0 = 6$, $n_1 = 9$, and $n_2 = 2$. This model is not stabilizable. Restricting our parameter variations to a_{52}, b_{41}, b_{52}, and b_{61} results in a stabilizable model, the simplified pitch-velocity model of Chapter 4. A reason for choosing the variation of the cross-coupling terms instead of the direct terms (b_{42}, b_{51}, and b_{62}) is that the cross-coupling terms vary from zero, whereas the direct terms are non-zero. It seemed intuitively that something changing from zero to non-zero would be more important than a certain percentage change in a parameter.

Conclusions

In the time-varying LPV framework, the ducted fan is completely controllable from the two inputs. When we move into the operator framework we lose this property of the model. This is a place where the conservatism of the operator framework markedly appears. Currently, we would like to determine how important the various parameter variations are and if, in fact, some of them are unnecessary. If we are able to obtain a stabilizable model with a large number of parameter variations, we will attempt model reduction techniques to reduce the order of our models [5].

6.4 Summary

In this chapter we have presented several theoretical techniques and applied them to the ducted fan. Each gave a particular insight into the ducted fan that we did not have before and justifies the relevance of the techniques. In particular, the identification of a model and the consequent application of model validation provides key information on the development of a robust model for the fan. Also, we have demonstrated again how the ducted fan serves as a testbed for the practical application of new theoretical techniques.

6.A Appendix: Transfer Functions

This appendix shows the remaining the transfer function data for the linearized model.

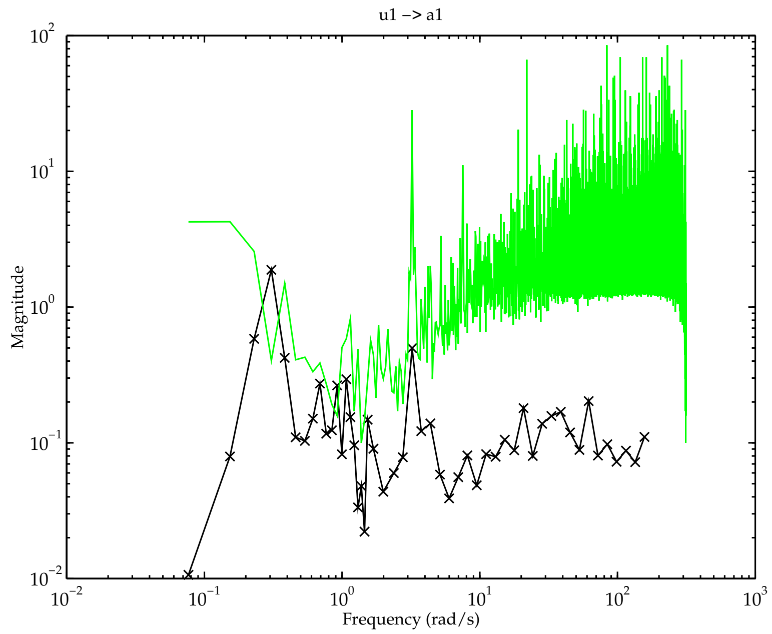

Figure 6.18: Transfer function from u_1 to α_1 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

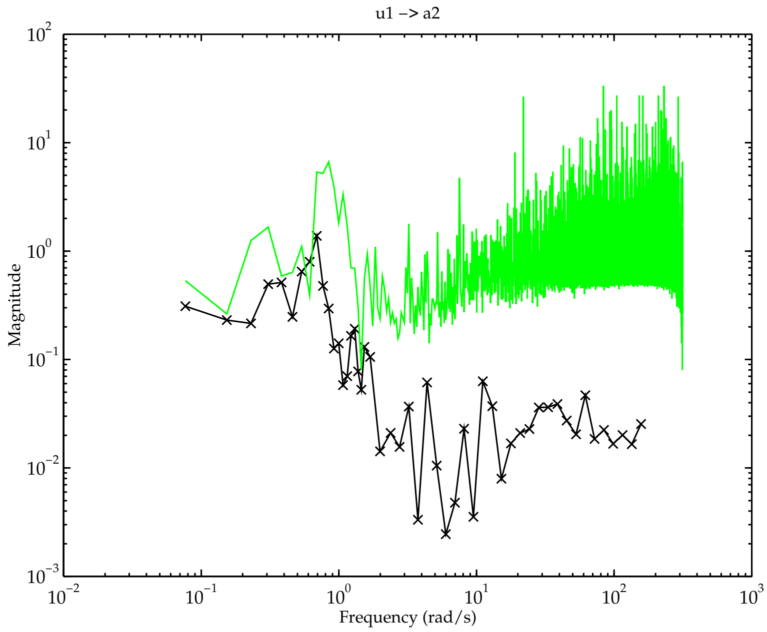

Figure 6.19: Transfer function from u_1 to α_2 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

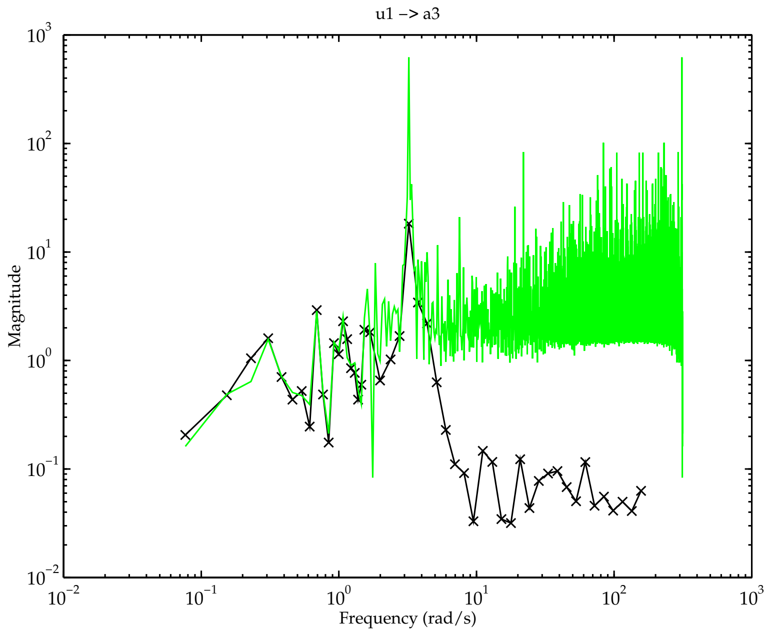

Figure 6.20: Transfer function from u_1 to α_3 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

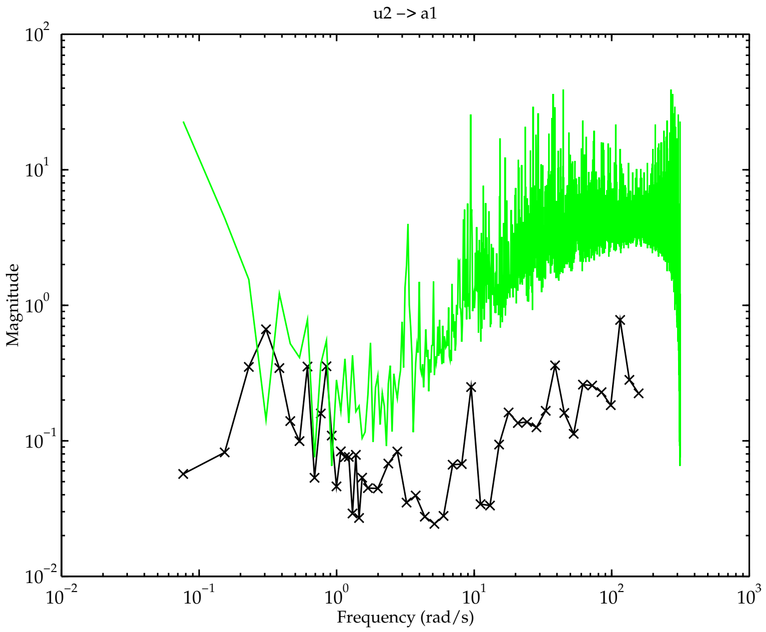

Figure 6.21: Transfer function from u_2 to α_1 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

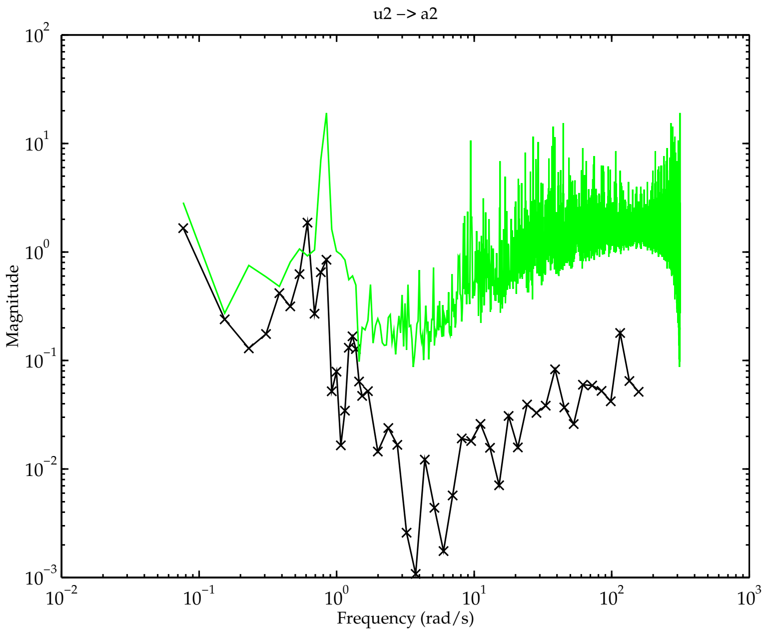

Figure 6.22: Transfer function from u_2 to α_2 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

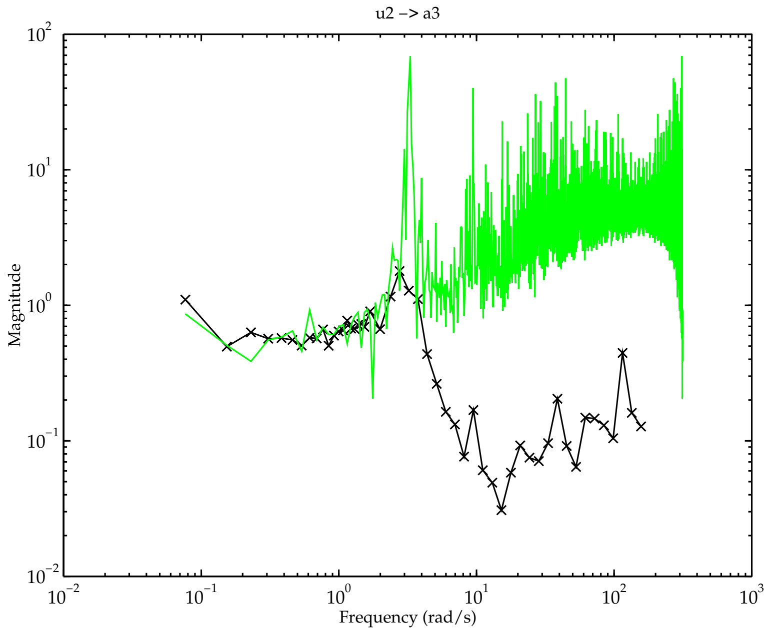

Figure 6.23: Transfer function from u_2 to α_3 from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

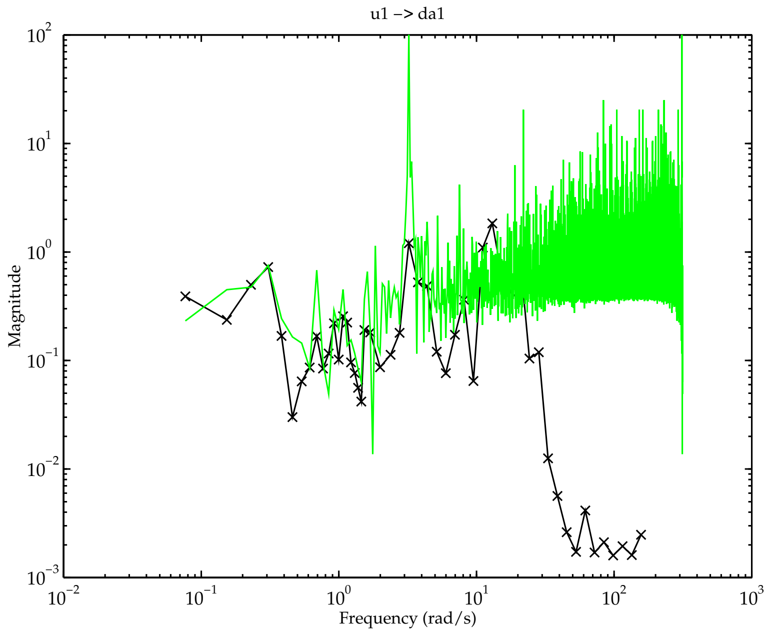

Figure 6.24: Transfer function from u_2 to $\dot{\alpha}_1$ from measured data (solid with an 'x' at the actual data point) and simulated with linearized model about hover (shaded).

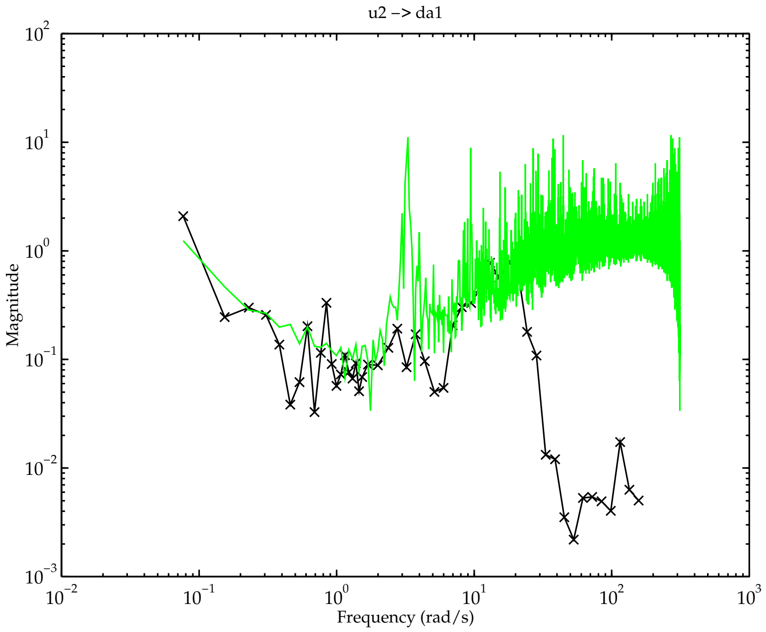

Figure 6.25: Transfer function from u_2 to $\dot{\alpha}_1$ from measured data (solid with an 'x' at the actual data point and simulated with linearized model about hover (shaded).

6.B Appendix: Validation Results

This section shows the results of other validation runs, both for the identified and linearized models. An ‘o’ in a figure shows a the value of the upper bound at that point, and an ‘x’ shows the lower bound.

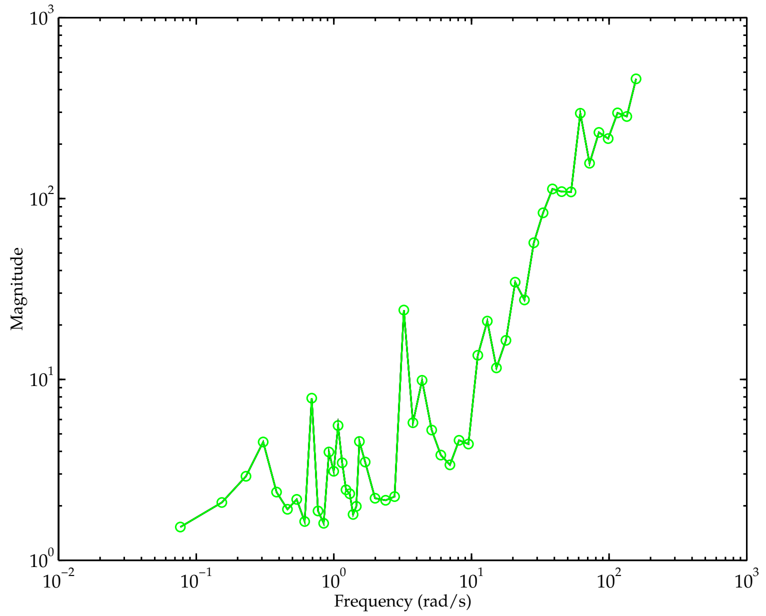

Figure 6.26: Model validation upper bound for the identified model on α_2.

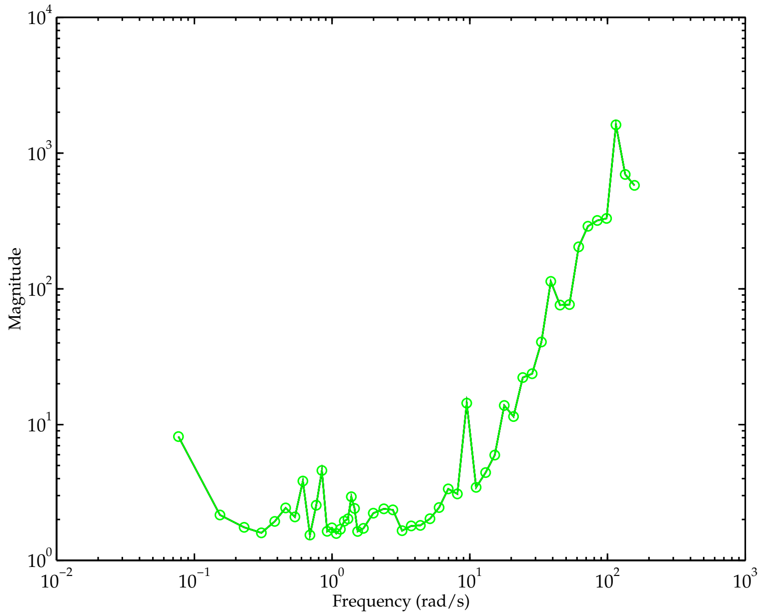

Figure 6.27: Model validation upper bound for the identified model on $\dot{\alpha}_1$.

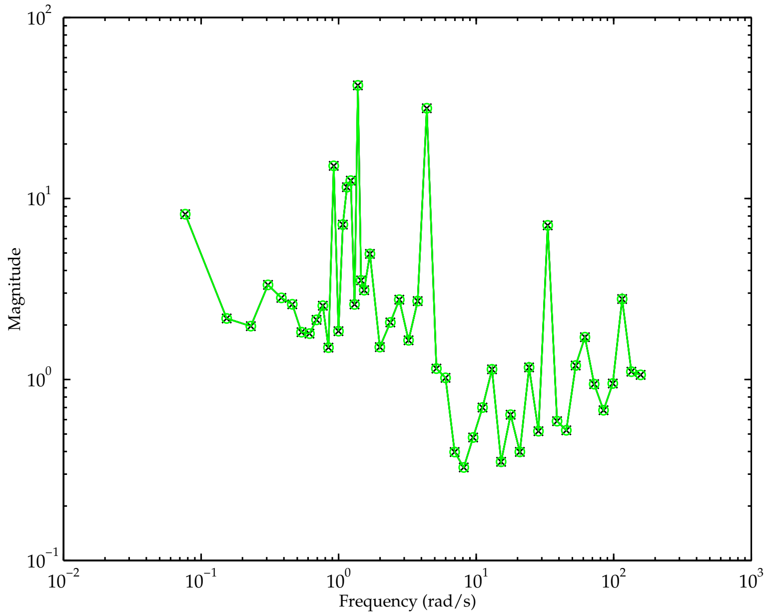

Figure 6.28: Model validation bounds for the linearized model about hover on α_1.

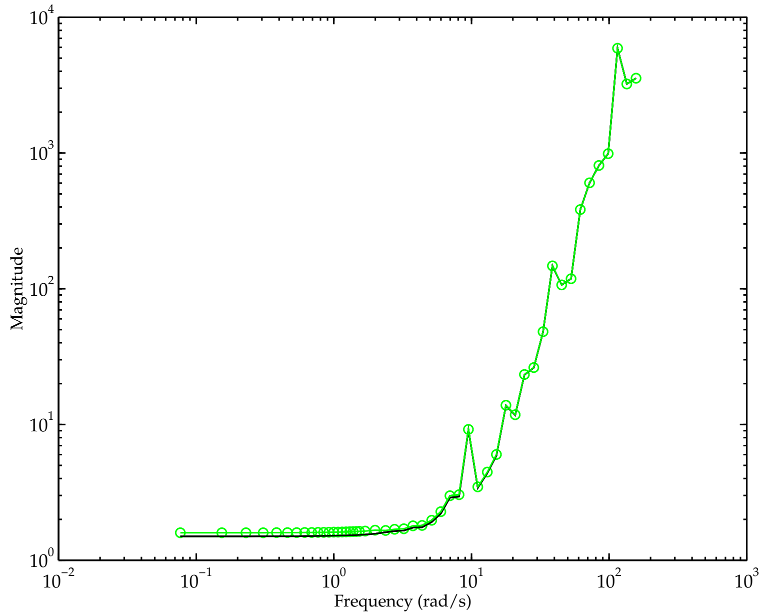

Figure 6.29: Model validation upper bound for the linearized model about hover on α_2 and α_3.

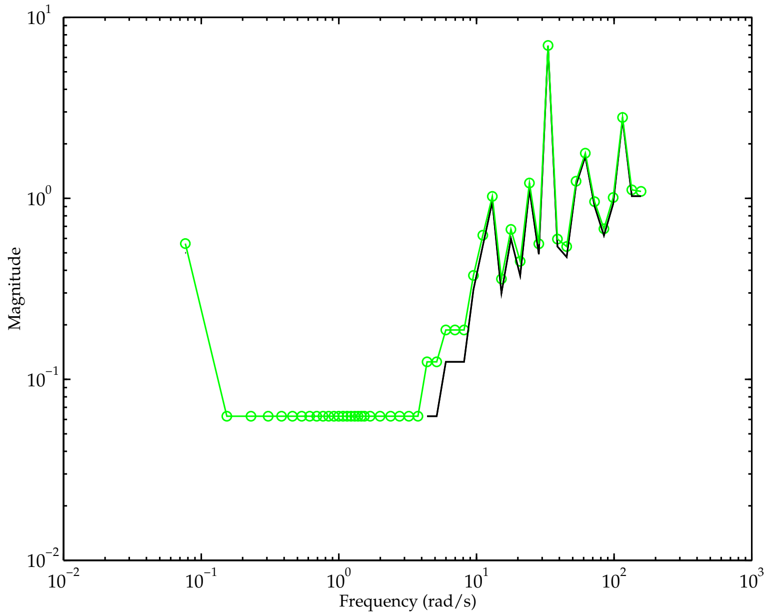

Figure 6.30: Model validation upper bound for the linearized model about hover on α_1, α_2, and α_3.

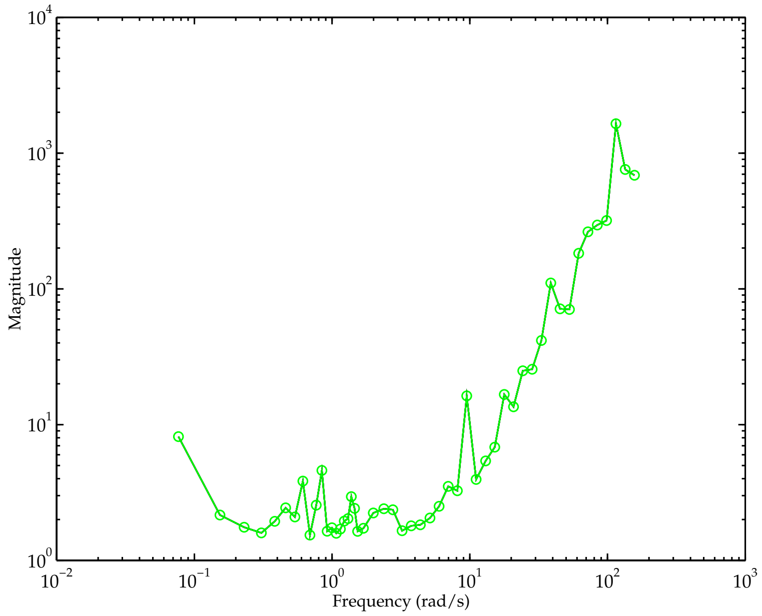

Figure 6.31: Model validation upper bound for the linearized model about hover on $\dot{\alpha}_1$.

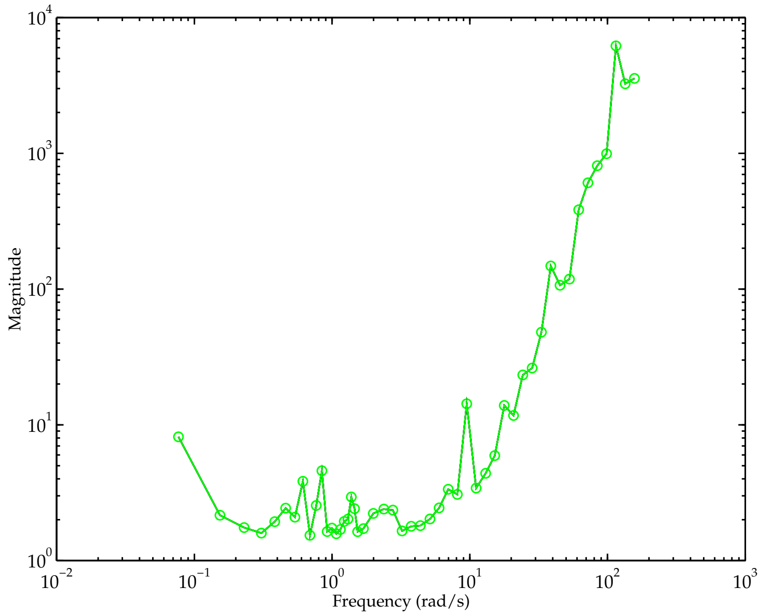

Figure 6.32: Model validation upper bound for the linearized model about hover on α_3.

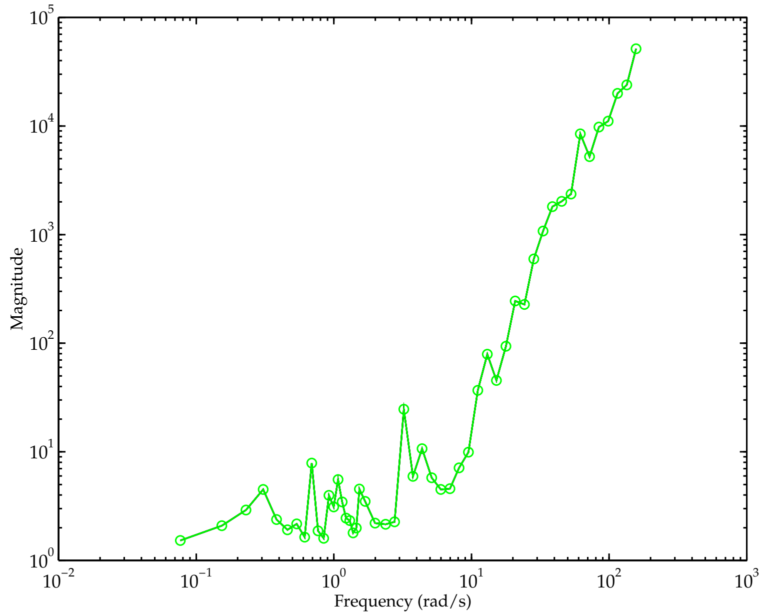

Figure 6.33: Model validation bounds for the linearized model about hover on α_3.

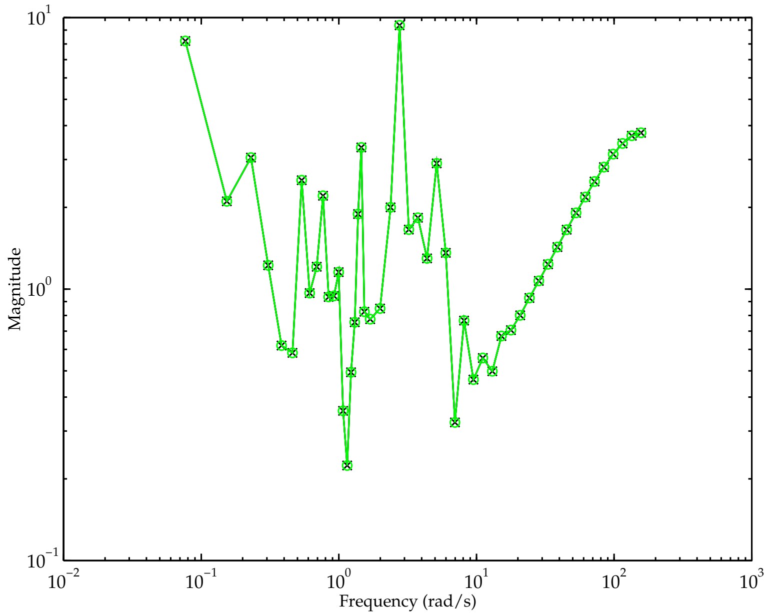

Figure 6.34: Model validation bounds for the linearized model with simulated data for α_3.

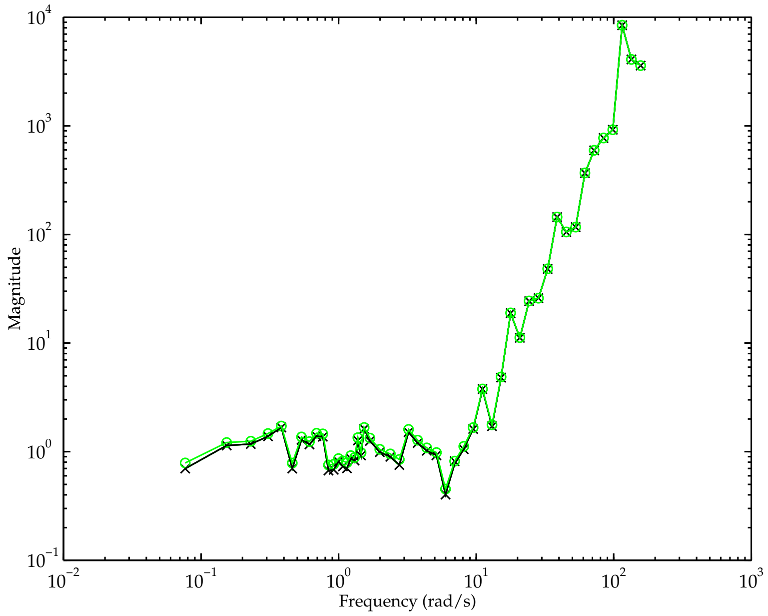

Figure 6.35: Model validation bounds for the linearized model with simulated data for α_2 and α_3 .

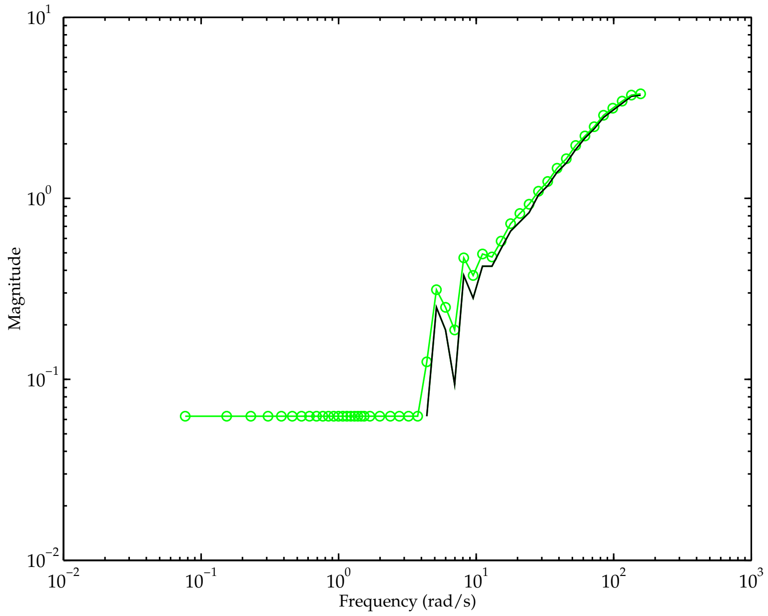

Figure 6.36: Model validation upper bound for the linearized model with simulated data for α_1, α_2 and α_3.

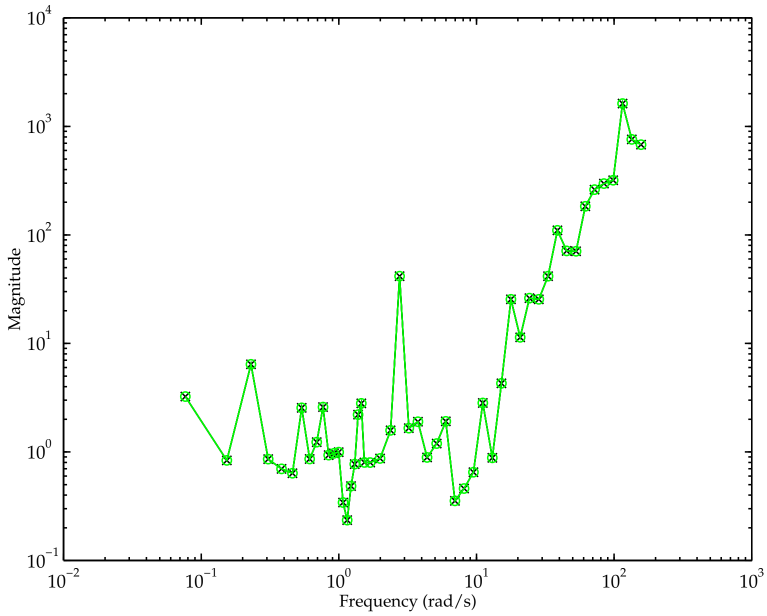

Figure 6.37: Model validation bounds for the linearized model with simulated data for $\dot{\alpha}_1$.

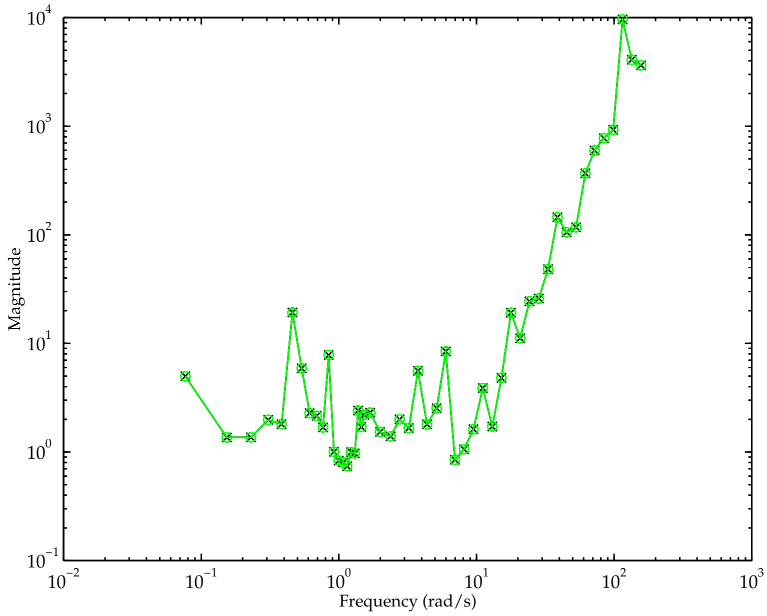

Figure 6.38: Model validation bounds for the linearized model with simulated data for α_3.

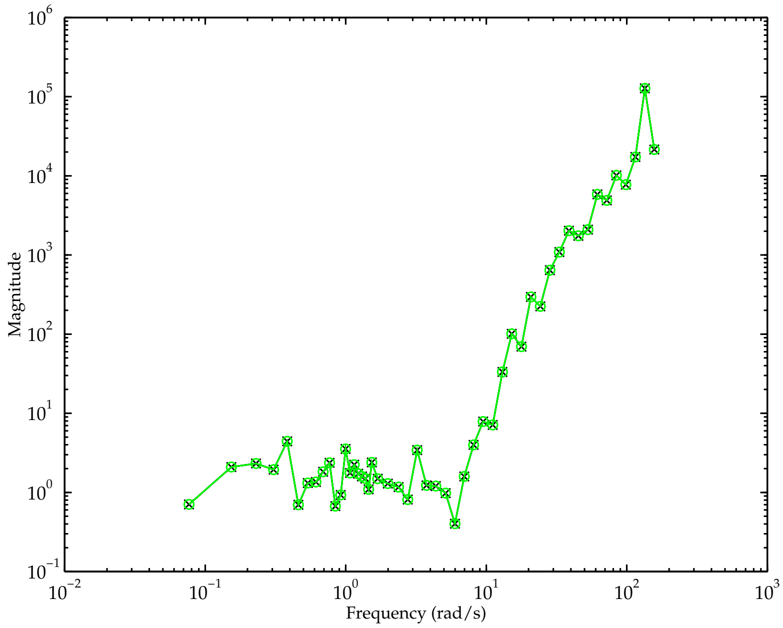

Figure 6.39: Model validation bounds for the linearized model with simulated data for α_2.

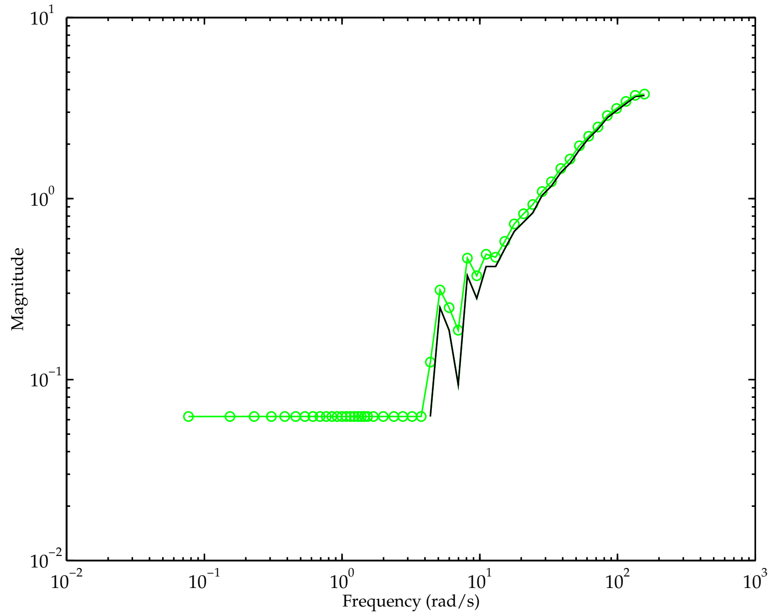

Figure 6.40: Model validation upper bound for the linearized model with simulated data for α_1 and α_3.

Chapter 7

Conclusions and the Future

> We are at the end of our enquiry, but as often happens in the search after truth, if we have answered one question, we have raised many more; if we have followed one track home, we have had to pass by others that opened off it and led, or seemed to lead, to far other goals than the sacred grove at Nemi. Some of these paths we have followed a little way; others, if fortune should be kind, the writer and the reader may one day pursue together.
>
> —James George Frazer

7.1 Summary of Contributions

In this dissertation we have demonstrated, on two applications, that synthesizing controllers to provide performance against worst-case time variations is not an inherently bad idea. The LPV technique produced a better design for the nuclear power plant than standard linear techniques. Also, the LPV techniques produced the best controller yet synthesized for the ducted fan, on most trajectories. These represent the first two applications of this particular technique.

On neither of the applications were standard linear techniques able to do as well as the scheduled ones. Note however, that the $\mathcal{H}_\infty$ design for the ducted fan performs well; the performance is not quite as good as the scheduled designs. An advantage of the LPV techniques over conventional methods of gain scheduling is that they design a controller of fixed order that works reasonably well for all plants in the operating regime.

For the PWR, the major drawback is that they do not switch control strategy between low and high power. On the nonlinear simulation of the PWR, the LPV controllers performed well even for small operating changes, where the assumptions on the uncertainty are extremely conservative.

Additionally, we showed that parameter variations can be placed in weights with beneficial effect. This is the first time such a design has been attempted that does this, and it is important because for many systems, one can expect dynamics to change so much that "frozen" time-invariant specifications will not yield adequate performance. This is of concern even in adaptive control.

Model validation provided insights into the model for the ducted fan which were not readily obtainable otherwise. It convincingly demonstrates that our robust control model is inadequate to describe not only the physical system, but the full nonlinear model. This is particularly surprising and is the subject of current investigation. Additionally, model validation provided a useful and easy means of comparing the identified model and the linearized model. We were able to determine that the identified model is superior to the linearized model. This is not a surprising result, but would be difficult to verify by other means.

Of perhaps more general interest than the specific applications, we have presented a simple methodology useful for designing future controllers which need to be gain scheduled. The closeness of linear $\mathcal{H}_\infty$ synthesis and LPV design techniques lead us to attempt an $\mathcal{H}_\infty$ design first. This allows us to exploit our experience and intuition at weight selection for linear plants. Once an $\mathcal{H}_\infty$ design has been synthesized, the weights are iterated on in the context of the LPV synthesis.

We showed how to construct parameter-varying models from linearizations about equilibrium points and from models identified at various operating points. This in itself is not a difficult task, but perhaps one under-appreciated, since, as we saw in Chapter 4, the models constructed can often be unsuitable for use in design because of stabilizability or detectability concerns.

A variety of performance measures were employed in evaluating the controllers for the ducted fan. Rather than overwhelming the reader with a barrage of statistics, our intent was to present a variety of measures and, for a particular application, select those which seem important and evaluate the controllers based on that. This should aid in comparing controllers when there are no specific performance objectives.

Based on this body of experimental results with the ducted fan, it seems that the performance specifications which coincide best with our intuition about what constitutes good performance are the infinity and one norms of the error signals. Unfortunately, none of our controllers are optimized to perform on those particular specifications. Moreover, there are some aspects of the ducted fan performance we have not been able to capture in a specification yet. Most of these are related to contrasting how the fan looks when under the control of a particular controller. For example, most of the controllers generate fairly wild oscillations on α_3. A notable exception here is the I/O decoupled controller, which doesn't oscillate since it actually tracks α_3. Thus the I/O decoupled controller appears much smoother on large motions than most of the others. Further investigation of performance specifications in an effort to capture these qualities seems warranted.

The conclusions of [27] have been reinforced. In particular, we have demonstrated that scheduled designs achieve better tracking performance than linear designs. Thus, pure linear techniques seem inadequate to control the ducted fan. Additionally, we have demonstrated that very good performance is achievable using a good first principles model. More importantly, a clear advantage the LPV and gain scheduling techniques have over nonlinear methods is that they allow the application of a variety of knowledge and intuition derived from the linear setting. Most current nonlinear techniques rely at some point on the use of pole placement, which has clearly understood drawbacks in the linear setting, and this puts them at a disadvantage for ducted fan applications.

7.2 Future Work

A complete controller design for the primary circuit of the PWR was not presented in this dissertation. A complete control system would, for example, account for saturation nonlinearities in the input signals and usually has a dead-band built-in to minimize the movement of control rods to small variations in operating conditions. Accounting for saturations to prevent wind-up is certainly an essential component of any realistic design and should be considered in future efforts towards design of a complete system. Another goal is to re-examine the temperature reference to determine if a reference derived differently leads to better minimization of the axial offset. Understanding why the LPV controller fails to switch strategies should be made.

Another interesting avenue to explore would be to try the model reduction techniques of Beck [5] on the parameter-varying model for the PWR. The PWR is stable so the theory is directly applicable, as it is not for the ducted fan. Some preliminary steps in this direction have already been made. One way of checking whether the plant is reducible in the size of the Δ-block is to treat the state as an input and output, and the Δ-block as the state, then look at the Hankel singular values of the system. For the PWR they are: 2.5448, 0.1031, 0.0325, 0.0187, 0.0152, and 0.0035. This indicates the size of the Δ-block could probably be reduced by at least one.

Regarding the ducted fan, experimental goals are further study of nonlinear robust control using this fan or a successor. Currently, a new ducted fan with a more aerodynamic shape is being designed and modelled using a wind tunnel at Caltech. This fan will be much more powerful and maneuverable than the rather heavy model used currently. There are also plans to add a wing to the ducted fan, so that aerodynamic effects become more significant.

More work is needed to investigate aggressive trajectories, similar to Herbst maneuvers [24], for the ducted fan. Currently work is being done on developing control strategies which couple with real-time trajectory generation. It would be extremely interesting to compare such techniques with the LPV techniques. As mentioned in the previous section, further work on performance specifications should also be done.

The LPV models should be expanded to the limits of stabilizability, to see how the achievable performance changes. Strictly from the point of view of applying the LPV techniques to the ducted fan, there are many more designs and experiments that can be done. Extending the parameter range of interest, and comparing how the achievable performance varies as the model becomes closer to not being stabilizable would be an extremely interesting set of tests. Moreover, we should determine what parameters and ranges have the most impact on stabilizability. Additionally, robustness properties of the LPV controllers could be further investigated.

A major avenue to explore in the context of real physical systems is to determine if, and by how much, practical performance can be improved by using synthesis methods which are theoretically less conservative. The area of reducing the theoretical conservatism of LPV techniques is an area of study attracting a good deal of attention in the control community now, but some assessment of its relevance seems in order. Additionally, most of the methods involve contorting the model so that all the parameter variations fall in specific places, usually not in the B, C, or D matrices of a system, but in the A matrix [8, 2]. The process for doing this involves application of a filter or some other method of expanding the size of the problem, and its impact on the computability of the problem is not currently understood.

The nonlinear model of the ducted fan should be improved based on data obtained from the identified model (parameters in the nonlinear model can be backed out from the parameters of the identified model). Additionally, the overall robust control model used for linear designs should be redesigned based on the results of model validation and investigations into any first principles' sources of uncertainty pursued. When this is done, the benefits of employing a synthesis methodology capable of exploiting the new model, such as μ-synthesis, should be undertaken.

Additionally, further work on identification and modelling is needed. The identified model presented here did not behave as accurately as we would like on α_1. Accounting for the variation of the distance from the center of rotation of the fan to the point at which the forces act, r, should be done in future robust control models of the ducted fan. Measurements have shown that r can vary by as much as 20% during operation of the fan. Dealing with this variation seems important.

Additionally, we have seen a need for better LMI solvers. The LPV synthesis routines experienced various anomalies when attempting to solve the LMIs. In the synthesis case, the LMI solver seemed very sensitive to the value of γ being used, and would occasionally find feasible solutions at lower levels of γ than where it previously could not find feasible solutions. This experience leads us to conclude that more work on developing reliable tools for solving LMIs is still needed, and that simply reducing a problem to an LMI is not yet equivalent to solving it.

The model validation problems we would like to solve are probably better handled in the time-domain. This is a more difficult problem than handling the data in the frequency domain, but handling the data in the frequency domain implies that, for the ducted fan at least, the data be processed in some way. This signal processing usually involves making it look nice in the time domain, so that the Fourier transform has reasonable behavior. Unfortunately, the effect of this processing on the model validation problem is not well understood. Additionally, with time-domain techniques, the nature of the perturbation required, e.g., time-varying or time-invariant, can be explored. Preliminary work by Smith and Dullerud [47] has been done on this problem, but more work is needed.

Bibliography

> Vengamos ahora a la citación de los autores que los otros libros tienen, que en el vuestro os faltan. El remedio que esto tiene es muy fácil, porque no habéis de hacer otra cosa que buscar un libro que los acote todos, desde la A hasta la Z, como vos decís. Pues ese mismo abecedario pondréis vos en vuestro libro; que, puesto que a la clara se vea la mentira, por la poca necesidad que vos teníades de aprovecharos dellos, no importa nada; y quizá alguno habrá tan simple que crea que de todos os habéis aprovechado en la simple y sencilla historia vuestra, y largo catálogo de autores a dar de improviso autoridad al libro.
>
> —Miguel de Cervantes Saavedra

[1] P. Apkarian and P. Gahinet. A convex characterization of parameter dependent $\mathcal{H}_\infty$ controllers. *IEEE Transactions on Automatic Control*, 1995. To appear.

[2] P. Apkarian, P. Gahinet, and G. Becker. $\mathcal{H}_\infty$ control of linear parameter-varying systems: A design example. *Automatica*, 1995. Submitted.

[3] P. Apkarian, P. Gahinet, and J. Biannic. Self-scheduled $\mathcal{H}_\infty$ control of a missile via LMIs. In *Proceedings of the 33rd Conference on Decision and Control*, pages 3312–3317, Florida, 1994.

[4] Gary J. Balas and Andy Packard. Design of robust, time-varying controllers for missile autopilots. In *Proceedings of the IEEE Conference on Control Applications*, 1992.

[5] Carolyn Beck. Minimality for uncertain systems and IQCs. In *Proceedings of the 33rd Conference on Decision and Control*, pages 3068–3073, Florida, 1994.

[6] Carolyn Beck, Bobby Bodenheimer, and Pascale Bendotti. LMI-based model reduction for a vectored-thrust ducted fan experiment. In *Proceedings of the Conference on Decision and Control*, 1995. Submitted.

[7] Carolyn Beck, John Doyle, and Keith Glover. Model reduction of multidimensional and uncertain systems. *IEEE Transactions on Automatic Control*, 1995. Submitted.

[8] Gregory Becker, Andy Packard, Doug Philbrick, and Gary Balas. Control of parametrically-dependent linear systems: A single quadratic Lyapunov approach. In *Proceedings of the American Control Conference*, pages 2795–2799, 1992.

[9] T. G. Beckwith, R. D. Marangoni, and J. H. Lienhard. *Mechanical Measurements*. Addison-Wesley, Reading, Massachusetts, 1993.

[10] Pascale Bendotti and Bobby Bodenheimer. Identification and $\mathcal{H}_\infty$ control design for a pressurized water reactor. In *Proceedings of the 33rd Conference on Decision and Control*, pages 1072–1077, Florida, 1994.

[11] Pascale Bendotti and Ernest Irving. Modeling and robust frequency-shaped LQG control of a pressurized water reactor. Technical Report EDF HI/20/93/009, Electricite de France, August 1993.

[12] Pascale Bendotti and Ernest Irving. Design of an $\mathcal{H}_2$ controller for a pressurized water reactor. In *Proceedings of the IFAC Symposium on Robust Control Design*, pages 382–387, September 1994.

[13] Bobby Bodenheimer and Pascale Bendotti. Optimal linear parameter-varying control design for a pressurized water reactor. Technical Report CIT-CDS 95-001, California Institute of Technology, January 1995.

[14] Bobby Bodenheimer, Pascale Bendotti, and Michael Kantner. Linear parameter-varying control of a ducted fan engine. Technical Report CIT-CDS 95-004, California Institute of Technology, January 1995.

[15] Stephen Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM, Philadelphia, 1994.

[16] Stephen P. Boyd and Craig H. Barratt. *Linear Controller Design: Limits of Performance*. Prentice-Hall, Englewood Cliffs, New Jersey, 1991.

[17] H. Choi, P. Sturdza, and R. M. Murray. Design and construction of a small ducted fan engine for nonlinear control experiments. In *Proceedings of the American Control Conference*, pages 2618–2622, 1994. Available electronically from `http://avalon.caltech.edu/~dfan`.

[18] John C. Doyle. Analysis of feedback systems with structured uncertainties. *IEE Proceedings*, 129, Part D(6):242–250, November 1982.

[19] John C. Doyle, Keith Glover, Pramod P. Khargonekar, and Bruce A. Francis. State-space solutions to standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems. *IEEE Transactions on Automatic Control*, AC-34(8):831–847, August 1989.

[20] John C. Doyle, K. Lenz, and A. Packard. Design examples using μ-synthesis: Space shuttle lateral axis FCS during reentry. In *Proceedings of the Conference on Decision and Control*, pages 2218–2223, 1986.

[21] Dale F. Enns. Model reduction with balanced realizations: an error bound and a frequency weighted generalization. In *Proceedings of the Conference on Decision and Control*, volume 2, pages 127–132, 1984.

[22] P. Gahinet, A. Nemirovskii, A. J. Laub, and M. Chilali. The LMI control toolbox. In *Proceedings of the 33rd Conference on Decision and Control*, pages 2038–2041, 1994.

[23] B. Gal-Or. Fundamental concepts of vectored propulsion. *Journal of Propulsion*, 6(6):747–757, 1990.

[24] B. Gal-Or. *Vectored Propulsion, Supermaneuverability and Robot Aircraft*. Springer-Verlag, New York, 1990.

[25] Keith Glover and John C. Doyle. A state space approach to $\mathcal{H}_\infty$ optimal control. In H. Nijmeijer and J. M. Schumacher, editors, *Three Decades of Mathematical Systems Theory: A Collection of Surveys at the Occasion of the 50th Birthday of Jan C. Willems*, volume 135 of *Lecture Notes in Control and Information Sciences*. Springer-Verlag, New York, 1989.

[26] R. A. Hyde and K. Glover. The application of scheduled $\mathcal{H}_\infty$ controllers to VSTOL aircraft. *IEEE Transactions on Automatic Control*, 38(7):1021–1039, July 1993.

[27] Michael Kantner, Bobby Bodenheimer, Pascale Bendotti, and Richard M. Murray. An experimental comparison of controllers for a vectored thrust, ducted fan engine. In *Proceedings of the American Control Conference*, 1995. To appear.

[28] C. Lemon and J. Hauser. Design and initial flight test of the champagne flyer. In *Proceedings of the 33rd Conference on Decision and Control*, pages 3852–3859, 1994.

[29] Lennart Ljung. *System Identification - Theory for the User*. Prentice-Hall, Englewood Cliffs, New Jersey, 1987.

[30] Wei-Min Lu. *Control of Uncertain Systems: State-Space Characterizations*. PhD thesis, California Institute of Technology, 1995.

[31] A. Megretskii and S. Treil. Power distribution inequalities in optimization and robustness of uncertain systems. *Journal of Mathematical Systems, Estimation, and Control*, 3, No. 3:301–319, 1993.

[32] Bruce Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Transactions on Automatic Control*, 26:17–31, 1981.

[33] J. Morris and M. Newlin. Model validation in the frequency-domain. In *Proceedings of the Conference on Decision and Control*, 1995. Submitted.

[34] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[35] R. M. Murray and E. L. Wemhoff. *Sparrow 2.0 Reference Manual*. California Institute of Technology, 1994. Available electronically from `http://avalon.caltech.edu/murray/~sparrow`.

[36] Yurii Nesterov and Arkadii Nemirovskii. *Interior Point Polynomial Methods in Convex Programming*. SIAM, Philadelphia, 1994.

[37] M. Newlin and R. Smith. A generalization of the structured singular value and its application to model validation. *IEEE Transactions on Automatic Control*, 1995. To appear.

[38] R. A. Nichols, R. T. Reichert, and W. J. Rugh. Gain scheduling for $\mathcal{H}_\infty$ controllers: A flight control example. *IEEE Transactions on Control Systems Technology*, 1(2):69–75, 1993.

[39] Andy Packard. Gain scheduling via linear fractional transformations. *Systems and Control Letters*, 22:79–92, 1994.

[40] Andy Packard and John Doyle. The complex structured singular value. *Automatica*, 29(1):71–109, January 1993.

[41] Fernando Paganini and John Doyle. Analysis of implicit uncertain systems Part I: Theoretical framework. Technical Report CIT-CDS 94-018, California Institute of Technology, December 1994.

[42] R. Redheffer. On a certain linear fractional transformation. *Journal of Mathematical Physics*, 39:269–286, 1960.

[43] Michael G. Safonov. Stability margins of diagonally perturbed multivariable feedback systems. *IEE Proceedings*, 129-D(6):251–155, 1982.

[44] Jeff S. Shamma. Robust stability with time-varying uncertainty. *IEEE Transactions on Automatic Control*, 39(4):714–724, April 1994.

[45] Jeff S. Shamma and Michael Athans. Guaranteed properties of gain scheduled control for linear parameter-varying plants. *Automatica*, 27(3):559–564, March 1991.

[46] Jeff S. Shamma and James R. Cloutier. A linear parameter varying approach to gain scheduled missile autopilot design. In *Proceedings of the American Control Conference*, pages 1317–1321, 1992.

[47] Roy Smith and Geir Dullerud. Validation of continuous-time control models by finite experimental data. In *Proceedings of the American Control Conference*, 1995. To appear.

[48] Roy S. Smith and John C. Doyle. Model validation: A connection between robust control and identification. *IEEE Transactions on Automatic Control*, 37(7):942–952, July 1992.

[49] Jan C. Willems. *The Analysis of Feedback Systems.* M.I.T. Press, Cambridge, Massachusetts, 1971.

[50] Kemin Zhou, John Doyle, and Keith Glover. *Robust and Optimal Control.* Prentice-Hall, Englewood Cliffs, New Jersey, 1995.