Linear Parameter-Varying Control of a Ducted Fan Engine

Bobby Bodenheimer
Microsoft
One Microsoft Way
Redmond, WA 98052

Pascale Bendotti
Electricité de France
Direction des Etudes et Recherches
6 Quai Watier
78401 Chatou, France

Michael Kantner
Mail Code 116-81
California Institute of Technology
Pasadena, CA 91125

SUMMARY

Parameter-dependent control techniques are applied to a thrust vectored ducted fan engine. The synthesis technique is based on the solution of Linear Matrix Inequalities 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. The controller designed performs extremely well, and is compared to other controllers.

Keywords: parameter-varying control, thrust vectoring, experimental application

1 Introduction

Recent advances in optimal control theory use linear matrix inequalities (LMIs) extensively in an attempt to develop theoretical and computational machinery for gain scheduling.[17] In this paper, we present the first application of this technique to an actual physical system, a ducted fan engine. By applying this technique to a specific nonlinear system, we hope to gain insight into its utility and, more generally, generate new ideas and insights into the problem of nonlinear robust control.

A picture of the experimental system, a thrust-vectored ducted fan engine, is shown in Figure 1. It consists of a high-efficiency electric motor with a six inch diameter blade, capable of generating up to nine Newtons of thrust. In Choi *et al.*,[9] 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.

In designing controllers for plants which operate over a wide dynamic range, a common technique is to schedule various fixed-point designs. Unfortunately, there are no known methods for scheduling such controllers which provide an *a priori* guarantee on the resulting performance or stability of the closed-loop system. Additionally, large and often unacceptable transients can occur when switching between controllers. Recent advances in optimal control theory provide a design technique which avoids these difficulties by producing an optimal parameter-dependent controller; i.e., the controller is already scheduled depending on parameter values which are not known beforehand.[1, 17] The controller is optimized to provide performance against the worst-case time variation of the parameters. Such a controller is called

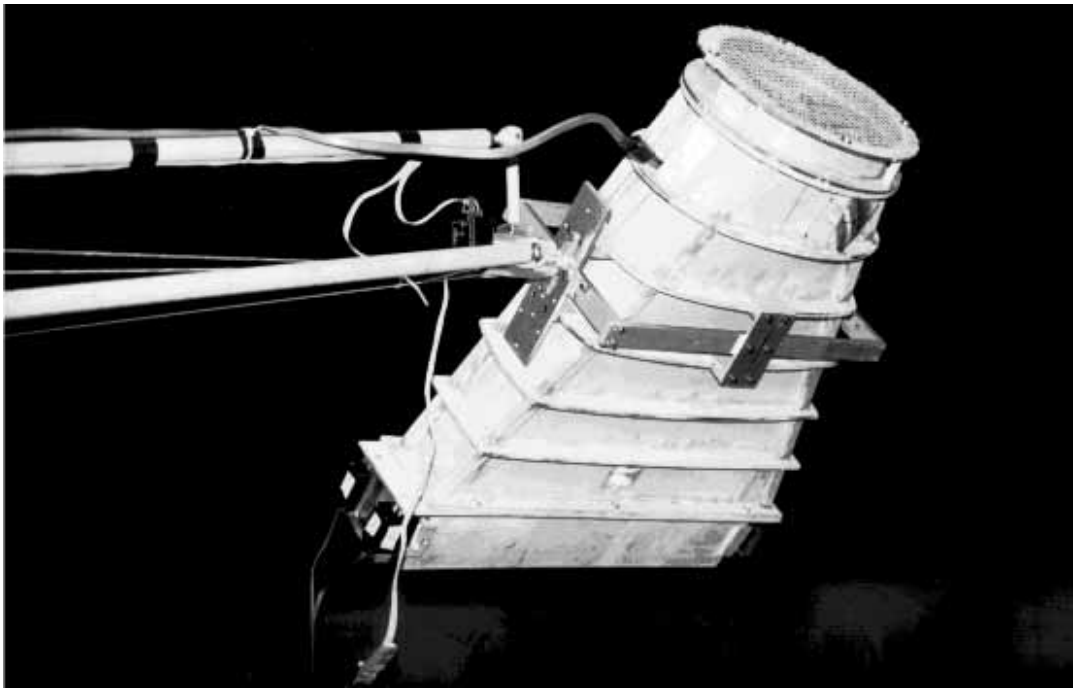

Figure 1: Ducted fan apparatus.

a linear parameter-varying (LPV) controller. The controller is synthesized via an iteration which involves solving three LMIs at each step.

A possible approach to control the ducted fan is to design a parameter-dependent controller with various measurements, such as velocity and pitch, as the parameters. 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 designing several controllers and then scheduling them. The potential drawback of LPV synthesis is that an LPV optimal controller is optimized against a worst-case time variation of parameters, and this time-variation may be so unrealistic that the controller has no performance. One of our goals in this paper is to determine if LPV synthesis can produce controllers which have reasonable performance, or better performance than linear controllers.

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[12] (see also the survey article[11]). 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.[14]

Although we are aware of LPV control being applied to examples and simulations,[2–4, 7] 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

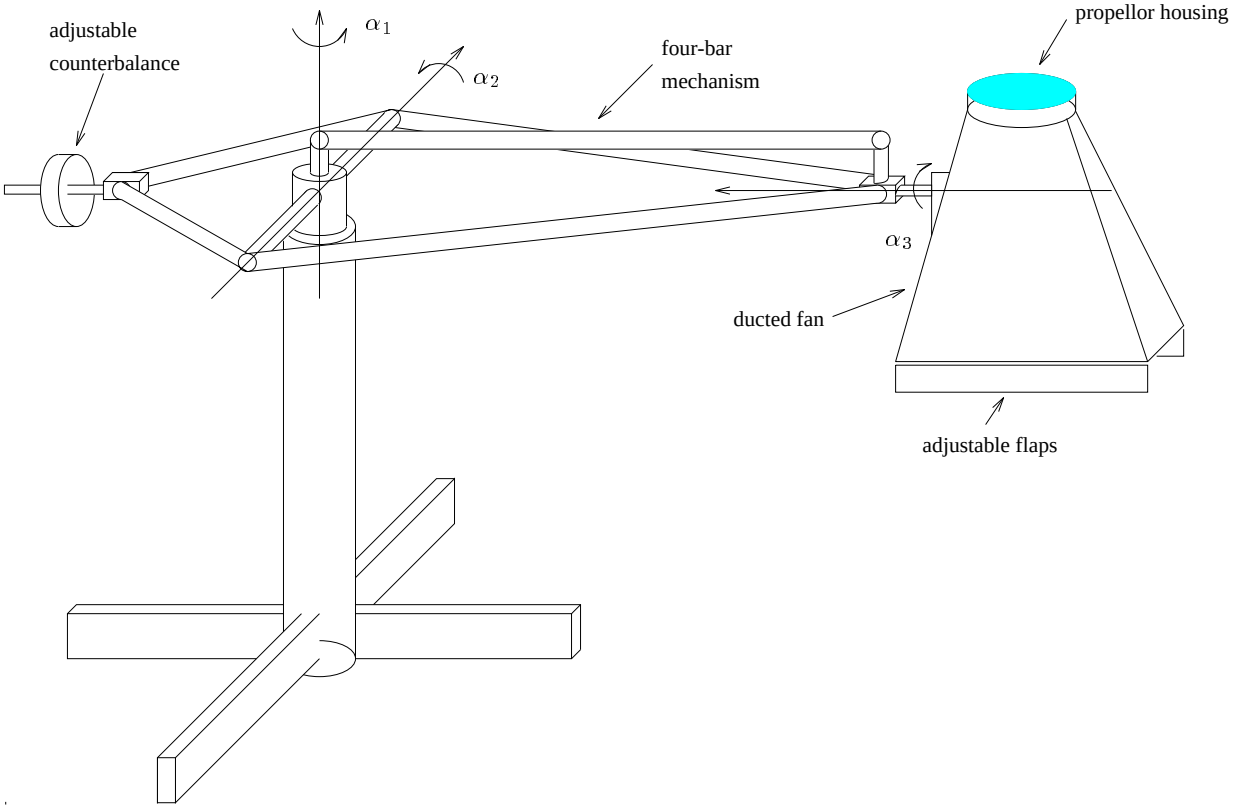

Figure 2: Ducted fan attached to stand.

previously designed for the ducted fan.[13]

In Section 2, we describe the configuration of the ducted fan and discuss its dynamics. In Section 3, the derivation of the parameter-varying models is presented. Section 4 provides a theoretical background on the parameter-varying synthesis. The experimental results of implementing the parameter-varying controllers are presented with analysis and comparisons in Section 5. We conclude with a discussion of our findings and present avenues for future work.

2 Description of the Fan Engine

2.1 Hardware

Overall, the experimental setup consists of the ducted fan attached to a three degree of freedom stand, as shown in Figure 2. The different thrust modes available are shown in Figure 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.[9]

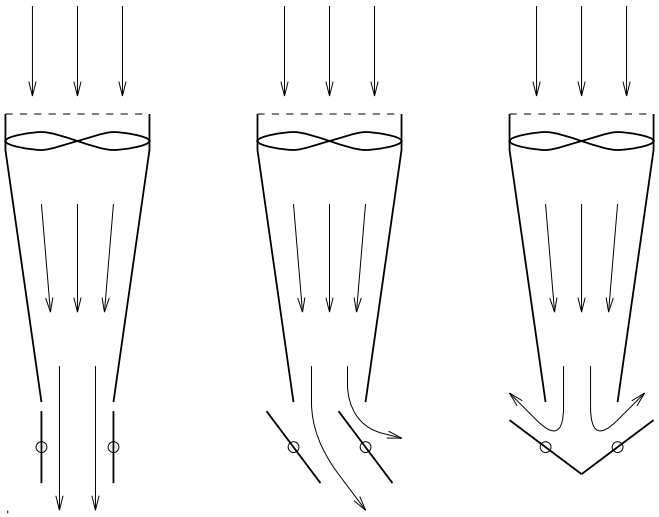

Figure 3: Different thrust modes for the ducted fan.

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.

2.2 Software interface

The experiment is interfaced to an 80486 computer running an MS-DOS-based real-time kernel called Sparrow.[16] Custom hardware is used to read in joint angles via the encoders and generate PWM signals for the R/C servos and a speed controller at a software selectable rate. 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. The R/C servos control the flap angles and the speed controller controls the motor speed, and consequently 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. Further discussion on the implementation of LPV controllers is found in Section 5.

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

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

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 with numerical values are derived in detail elsewhere[6] (see also the web page at http://avalon.caltech.edu/~dfan).

The model is accurate enough for control design, although it does have limitations. Identification experiments show it is reasonably accurate near hover.[6] Initial step responses on single axes compared favorably with experimental measurements, and a PID control test gave expected results.[9] A decoupling controller, essentially a plant inversion, worked well.[13] Nonetheless, the model omits many effects: all actuator dynamics, sensor limitations, friction, and aerodynamic effects. Static friction about the lateral axis (α_1) 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 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 fix the altitude (α_2) at 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

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

where T is the sampling rate.

Figures 4 and 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 can be found in.[6] The dependence is obtained by linearizing the nonlinear model at various equilibrium operating points for different values of pitch (α_3) and lateral velocity ($\dot{\alpha}_1$). Note the dependence of a_{42} and a_{62} on both pitch and lateral velocity, 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 pitch. In Figure 5, the actual values are shown as asterisks, and the least-squares fit described below is the solid line.

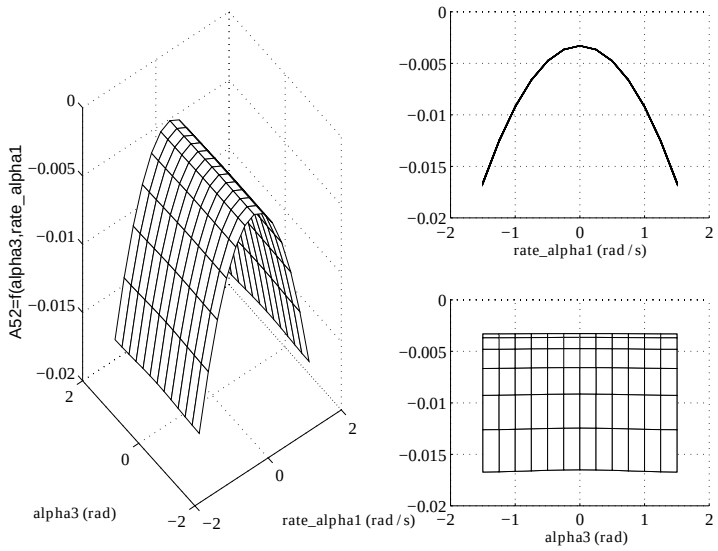

Figure 4: Dependence of a_{52} upon pitch (α_3) and lateral ($\dot{\alpha}_1$).

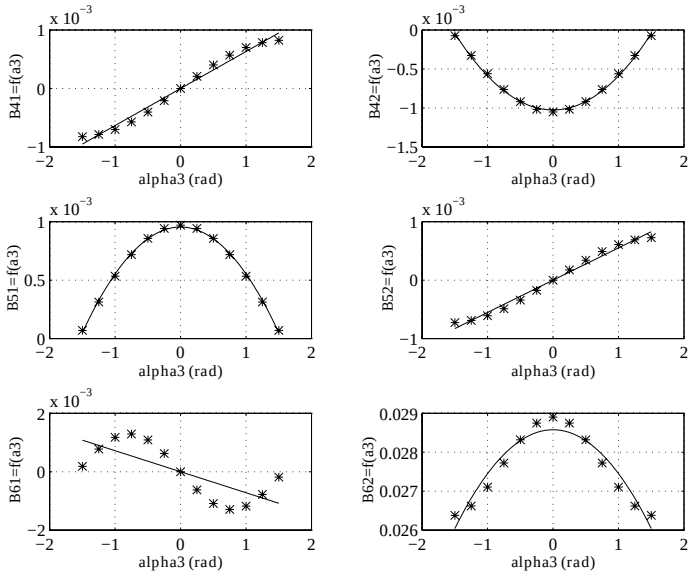

Figure 5: Dependence of the entries of the B matrix upon pitch (α_3).

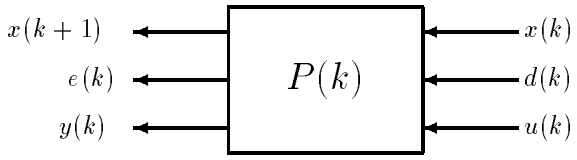

Figure 6: Time-Varying System.

To construct a model of the ducted fan suitable for LPV control, we assume the system has the general time-varying structure shown in Figure 6, 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

$$P(k) = P_{22} + P_{21}\Delta(k)\left(I - P_{11}\Delta(k)\right)^{-1} P_{12} \tag{2}$$

where

$$\Delta(k) = \begin{bmatrix} \delta_1(k) I_{n_1} & & \\ & \ddots & \\ & & \delta_m(k) I_{n_m} \end{bmatrix} \tag{3}$$

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. 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 7.

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 lateral velocity ($\dot{\alpha}_1$) was neglected, making them depend only on pitch (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

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

The resulting model with δ-dependence, $P(\delta)$, is shown in Figure 7, where n_0 is the size of the block corresponding to δ_0.

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 $(P_{des}, \mathbf{\Delta}_{des})$. 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. Nonetheless, they provide an important property to the model: when moving into the operator LPV framework, a more complicated model can lose stabilizability. This issue is discussed more thoroughly elsewhere.[5,6]

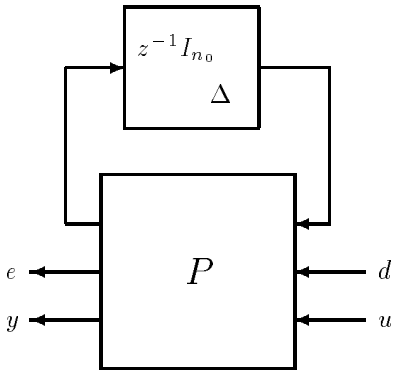

Figure 7: Parameter-Dependent Plant. The $z^{-1}I_{n_0}$ term represents the states of P, and the Δ represents the time variation of Equation 3.

4 Controller Synthesis

This section gives an introduction to the theoretical background for the synthesis methodology. All the results are due to Packard,[17] so the discussion is of a brief and general nature. We then discuss the synthesis for the ducted fan in detail.

4.1 Theoretical Background

Recall that the plant has the structure given in Figure 7. 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 8. 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 9, 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 $\ell_2 \to \ell_2$ induced norm from d to e is small for all allowable parameter variations $\Delta(k)$ (see Equation 3). 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 $\mathbf{\Delta}$ 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).[18] This condition now becomes (compare Lemma 3.1[17] and Theorem 10.4[18]):

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*

$$\left\| \begin{bmatrix} T & 0 \\ 0 & I \end{bmatrix} (R \star K) \begin{bmatrix} T^{-1} & 0 \\ 0 & I \end{bmatrix} \right\|_\infty < 1 \tag{4}$$

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 9 is internally exponentially stable, and for zero initial conditions, if $d \in \ell_2$, then $\|e\|_2 \le \gamma \|d\|_2$.

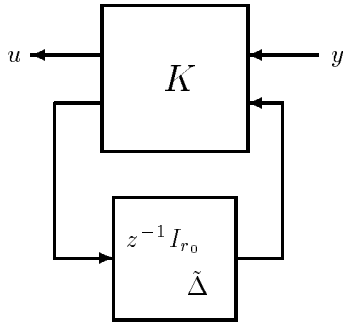

Figure 8: Parameter-dependent controller; $z^{-1}I_{r_0}$ represents the states of the controller and $\widehat{\Delta}$ the time variations.

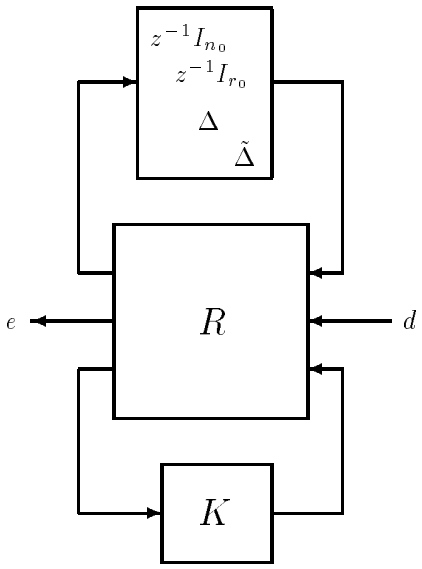

Figure 9: Parameter-dependent closed-loop system.

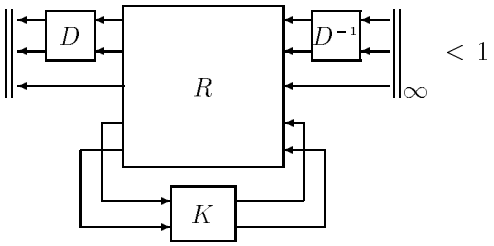

Figure 10: Diagram of Theorem 1.

Pictorially, this theorem is shown in Figure 10. 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 of Lu[15] that such a K will exist when R is stabilizable and detectable with respect to the block structure

$$\mathbf{\Delta} = \text{diag}\left(z^{-1}I_{n_0}, z^{-1}I_{r_0}, \Delta, \widehat{\Delta}\right).$$

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

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

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

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

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.[17] 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 3, 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 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.

4.2 Weight Selection

The design process parallels classical $\mathcal{H}_\infty$ design and LTI weights were used. Using weights depending on parts of the Δ-block can enhance performance,[7] but was not needed here. 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 11. 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 (altitude) 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. The response of α_1 (lateral position) has double integrator slope at low frequencies and a mode at 3.5 rad/s, the same frequency as the mode for pitch. 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 12, 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. The parameter variations Δ have been absorbed into P in this diagram. W_u is a penalty on the control, and W_m is multiplicative actuator uncertainty. W_p is a diagonal performance weight on the lateral position (α_1), altitude (α_2), pitch (α_3), and lateral velocity ($\dot{\alpha}_1$). Note that LPV synthesis deals with one full block, analogous to $\mathcal{H}_\infty$ synthesis Implicitly, there are two blocks in Figure 12 — 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 LPV 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.

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. W_m is the frequency-dependent weight on this uncertainty. It resembles a high pass filter and is shown in solid lines in Figure 11. 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 pitch, but some small performance weight on this output was found helpful. To provide some damping on the mode without sacrificing performance on lateral position (α_1) or altitude (α_2), the weight is set to a small constant value of 0.1. A constant weight is sufficient on lateral position since it has high gain in the low frequency. A similar weight was placed on lateral velocity ($\dot{\alpha}_1$). Notice that the gain of the α_2 channel is small. Then

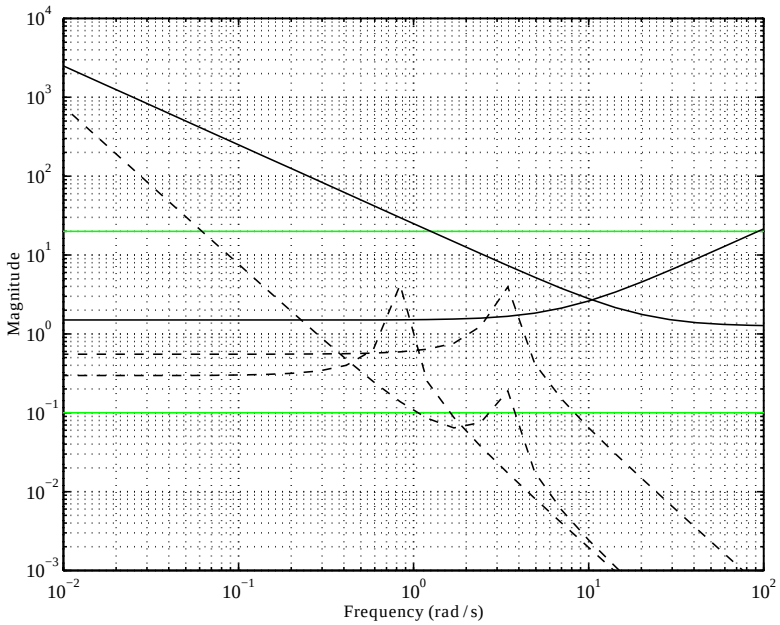

Figure 11: Performance and uncertainty weights for the $\mathcal{H}_\infty$ controller. The solid black lines are the performance weights on α_2 and the multiplicative uncertainty weight W_m. The shaded lines are performance weights on α_1 and α_3. The dashed lines are a Bode singular value plot of the linearized model.

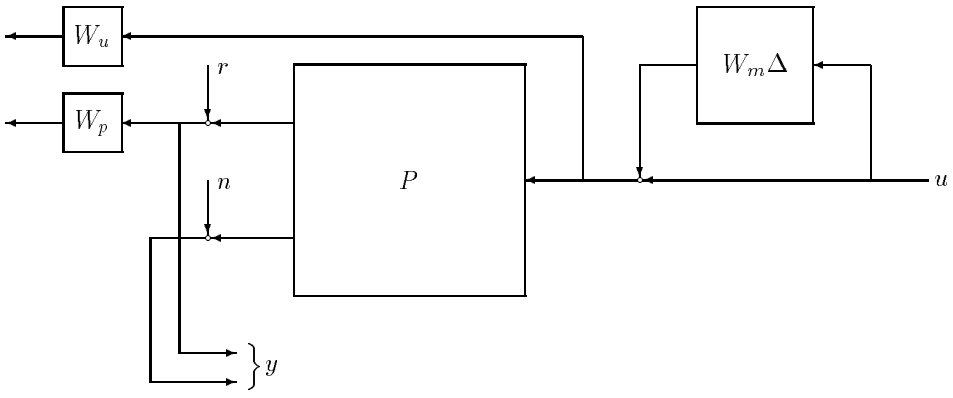

Figure 12: Synthesis structure used for designing the $\mathcal{H}_\infty$ and LPV controllers. See the text for an explanation of how this diagram fits into the $\mathcal{H}_\infty$ framework.

our controllers must have a large gain on α_2 to give good altitude performance, and so an integral-like weight is used. This is shown in solid lines as a low pass filter in Figure 11, and adds one state to the controller. This performance weight seeks to make the closed loop dynamics for altitude faster than those for lateral position, in an effort to reduce the altitude error caused when the fan is tilted at an angle.

5 Experimental Results

The LPV controllers were implemented using Sparrow, as mentioned previously. The LPV package 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. Since an LPV controller is an LFT on a Δ of measured parameters, implementing it requires a real-time matrix inversion, as in Equation 2. To reduce computations, though, the LFT can be eliminated using past values of Δ. This involves implementing the algorithm as shown pictorially in Figure 7, where a past value of Δ becomes an input to the plant. 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.

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 results for implementing the LPV controller based on the simplified pitch-velocity model are presented in Section 6.1. For comparison, we also present results obtained by locking the LPV controller at hover and in forward flight, as well as for other controllers designed previously.[13] The results shown are quite representative of the behavior of the ducted fan, but each one represents the "best" result of a series of three to five tests. The results are repeatable. An analysis of the repeatability of experiments on the ducted fan has been done.[6]

5.1 Trajectory Description

The controllers were tested on three trajectories. Two of the trajectories are simple and command changes on only one axis. The third trajectory is demanding and commands rapid changes in lateral position and altitude simultaneously. These are called the "α_1 ramp" and "α_2 step," respectively.

The first trajectory is a one radian change in lateral position over 5 seconds. The second is a 0.1 radian step change in altitude. 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 a lateral velocity of 0.628 rad/s, over three times greater than in the first trajectory. While in forward flight, a sinusoidal variation is commanded in altitude, 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 "rolling" trajectory.

6 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.[13] These results are summarized in Tables 1 through 3.

The 10-90% rise time is a standard figure of merit for step responses. For the ramp in lateral position (α_1-ramp) 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 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 13. Boyd and Barratt[8] have a 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 $i = 1, 2$.

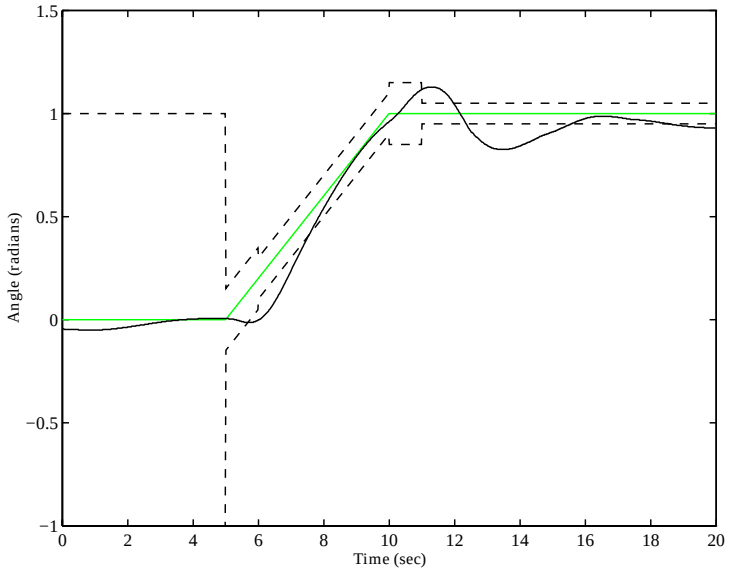

Figure 13: An example of the ramp envelope specification, with the response from the LPV controller. The shaded line is the desired response and the dashed lines are the allowable error bounds. The solid line is the actual response.

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. Finally, for the rolling trajectory, we present a normalized ranking scheme described later.

The controllers presented in the tables are the ones discussed previously, namely the LPV, and locked LPV controllers, in addition to some new ones. The design of these new controllers can be found in Kantner *et al.*,[13] and nothing beyond a few superficial remarks will be made here. The $\mathcal{H}_\infty$ controller is a straightforward simplification of the LPV design presented 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.

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

6.1 Results and Evaluations

Table 1 shows the figures of merit for the ramp in α_1, Table 2 shows the results for the step in α_2, and Table 3 shows the results for the 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$_{fl}$."

To conserve space, plots for each controller on each trajectory are not shown. Complete plots for the controllers on all trajectories may be found elsewhere.[6,13] These plots show the lateral, altitude, and pitch axes, with desired and actual values, in addition to the commanded forces. Recall that the "X" direction is lateral and the "Y" direction is vertical.

For the α_1 trajectory data shown in Table 1, several results are apparent. The LPV is the fastest controller, but has problems settling. This is shown in Figure 14. 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 altitude 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 15. 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

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

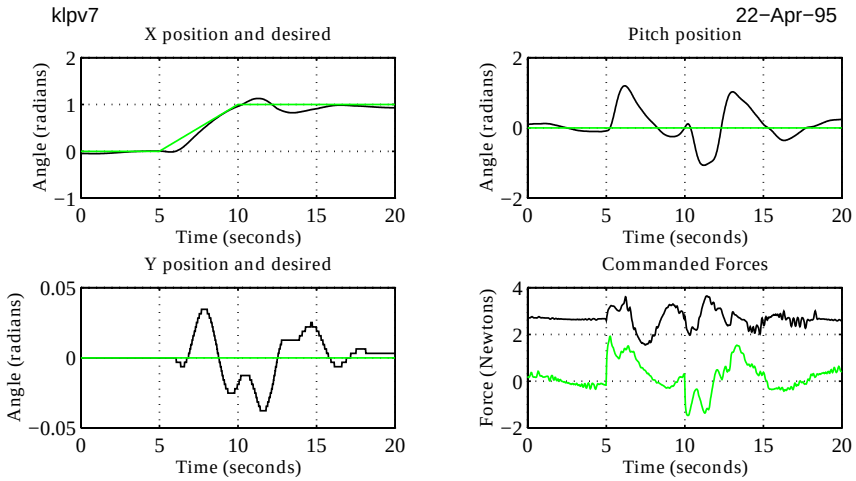

Figure 14: Closed loop response of the LPV controller for an α_1 ramp.

follow the ramp in α_1.

The performance results for the step in α_2 are shown in Table 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 specification.

The results for the rolling trajectory are tabulated in Table 3. The gain-scheduled and LPV controllers are close in performance, as is illustrated in Figures 17 and 18. The gain-scheduled controller is slightly better in lateral position while the LPV controller is slightly better in altitude. The LPV controller is better overall on tracking position and altitude, 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.

Finally, since ranking is obviously not the best way of evaluating a controller, we present an improved scheme for the rolling trajectory, shown in Table 4 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 4. 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

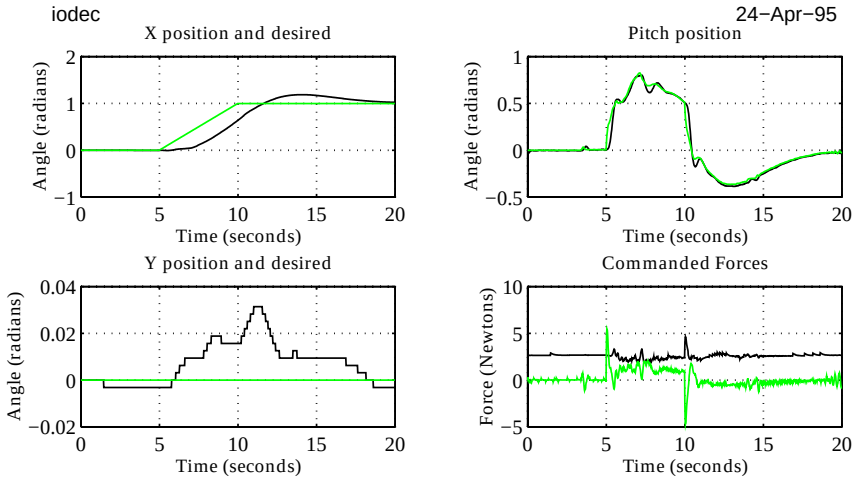

Figure 15: Closed loop response of the I/O decoupled controller for an α_1 ramp.

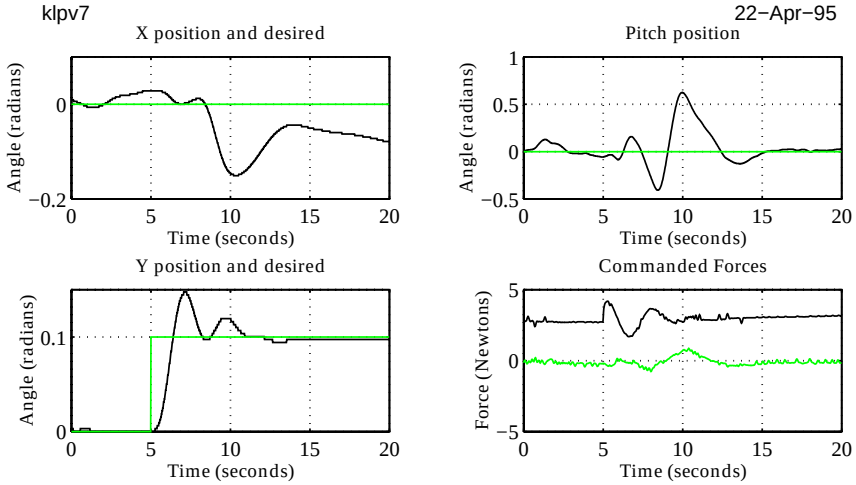

Figure 16: Closed loop response of the LPV controller for an α_2 step.

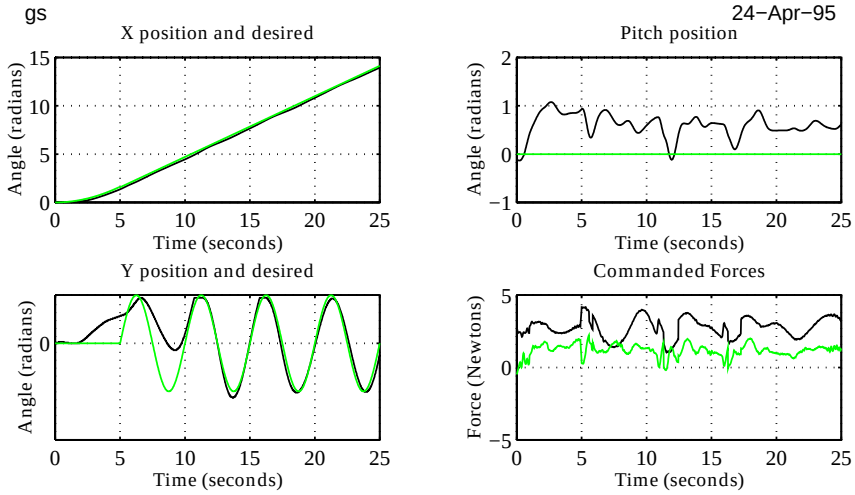

Figure 17: Closed loop response of the gain-scheduled controller for the rolling trajectory.

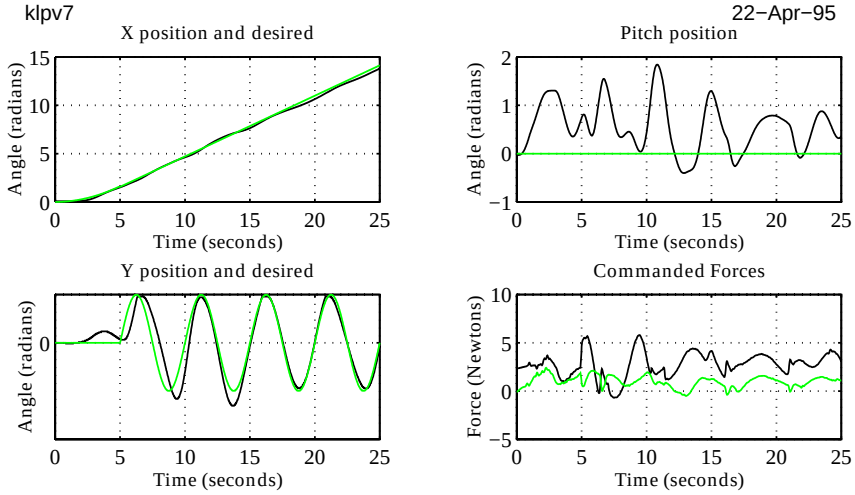

Figure 18: Closed loop response of the LPV controller for the rolling trajectory.

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 4: Normalized performance data for the rolling trajectory, with primary and secondary performance measures.

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.

7 Conclusions

In this paper we have applied an LMI based control technique to a practical application, a ducted fan engine. An LPV controller was synthesized which successfully controlled the ducted fan with adequate performance on all test trajectories and excellent performance on one particularly difficult trajectory. The LPV controller is the best controller for primary performance specifications yet designed for the rolling trajectory.

The actual difference between the gain scheduled controller and the LPV controller is not large, and the design of the LPV controller is much more involved. An advantage that the LPV controller possesses over the gain-scheduled design is that there is a theoretical justification for believing that the LPV controller will operate with adequate performance over the entire operating regime, whereas there is none for the gain-scheduled controller. Thus, we do not believe that the LPV designs presented here should supplant gain-scheduling, but careful consideration should be given as to whether the potential advantages of LPV design outweigh the complexity of the design process.

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 earlier 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.

Experimentally, our goals are oriented toward 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. Also, there are plans to add a wing to the ducted fan, so that aerodynamic effects become more significant.

Strictly from the point of view of applying the LPV techniques to the ducted fan, there are many more designs and experiments to do. Theoretically, we are investigating ways to improve the performance of the LPV synthesis. We would like to extend the parameter range in our designs, with the idea of eventually testing trajectories similar to Herbst maneuvers in aircraft. The robustness properties of the LPV controllers could be further investigated, and we would like to design controllers using more of the parameter variations than the reasonably simple model used in this paper.

Since we have explored only full LPV synthesis in this paper, we have not exploited any special properties of the ducted fan. In particular, the state is available for measurement, and the parameter variations are all states. When this is the case, there is an equivalence between the general synthesis problem and a full-information version of it. Solving the full information problem would result in a nonlinear constant-gain controller with the same performance level as the full LPV controller. Further investigation of this is currently in progress.

Acknowledgements

The authors have enormously benefited from the guidance and advice of Richard Murray. Carolyn Beck helped with many useful comments. Gary Fay constructed the force lookup tables used in the software, and has assisted throughout in hardware maintenance. The second author gratefully acknowledges financial support provided by Ministère de la Recherche et de l'Espace (MRE), France. We further thank the authors of the *LMI Control Toolbox*[10] for providing an early release which was used to solve the synthesis LMIs. The authors thank the anonymous reviewers for helpful comments. This work was performed while all the authors were at the California Institute of Technology.

References

[1] Apkarian, P., and P. Gahinet, 'A Convex Characterization of Parameter Dependent $\mathcal{H}_\infty$ Controllers', to appear in *IEEE Trans. on Automatic Control*, (1995).

[2] Apkarian, P., P. Gahinet, and J. Biannic, 'Self-Scheduled $\mathcal{H}_\infty$ Control of a Missile via LMIs', *Proc. 33rd Conf. on Decision and Control*, 3312–3317, Florida (1994).

[3] Apkarian, P., P. Gahinet, and G. Becker, '$\mathcal{H}_\infty$ Control of Linear Parameter-Varying Systems: A Design Example', submitted to *Automatica*.

[4] Balas, G. J., and A. Packard, 'Design of Robust, Time-Varying Controllers for Missile Autopilots,' *Proc. IEEE Conf. on Control Appl.*, Dayton, Ohio (1992).

[5] Beck, C., B. Bodenheimer, and P. Bendotti, 'LMI-Based Model Reduction for a Vectored-Thrust Ducted Fan,' *Proc. 34th Conf. on Decision and Control*, 871–872, New Orleans (1995).

[6] Bodenheimer, Jr., Robert E., *The Whirling Blade and the Steaming Cauldron*, Ph.D. thesis, California Institute of Technology, 1995.

[7] Bodenheimer, B., and P. Bendotti, 'Optimal Linear Parameter-Varying Control Design for a Pressurized Water Reactor', *Proc. 34th Conf. on Decision and Control*, 182–187, New Orleans (1995).

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

[9] Choi, H., P. Sturdza, and R. M. Murray, 'Design and Construction of a Small Ducted Fan Engine for Nonlinear Control Experiments,' *Proc. 1994 American Control Conf.*, 2618–2622, Baltimore (1994).

[10] Gahinet, P., A. Nemirovskii, A. J. Laub, and M. Chilali, 'The LMI Control Toolbox,' *Proc. 33rd Conf. on Decision and Control*, 2038–2041, Florida (1994).

[11] Gal-Or, B., 'Fundamental Concepts of Vectored Propulsion,' *Journal of Propulsion*, **6**(6), 747–757 (1990).

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

[13] Kantner, M., B. Bodenheimer, P. Bendotti, and R. M. Murray, 'An Experimental Comparison of Controllers for a Vectored Thrust, Ducted Fan Engine,' submitted to *American Control Conf.* (1995).

[14] Lemon, C., and J. Hauser, 'Design and Initial Flight Test of the Champagne Flyer', *Proc. 33rd IEEE Conf. on Decision and Control*, 3852–3853, Florida (1994).

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

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

[17] Packard, A., 'Gain Scheduling via Linear Fractional Transformations,' *System and Control Letters*, **22**, 79–92 (1994).

[18] Packard, A., and J. Doyle, 'The Complex Structured Singular Value', *Automatica*, **29**(1), 71–109 (1993).