

An Experimental Comparison of Controllers for a Vectored Thrust, Ducted Fan Engine

Michael Kantner Bobby Bodenheimer
Pascale Bendotti Richard M. Murray

Division of Engineering and Applied Science
California Institute of Technology
Pasadena, CA 91125

Abstract

Experimental comparisons between four different control design methodologies are applied to a small vectored thrust engine. Each controller is applied to three trajectories of varying aggressiveness. The control strategies considered are LQR, \mathcal{H}_∞ , gain scheduling, and feedback linearization. The experiments show that gain scheduling is essential to achieving good performance. The strengths and weaknesses of each methodology are also examined.

1 Introduction

This paper is concerned with the application of linear and nonlinear control design methods to high-performance jet aircraft performing aggressive maneuvers. Very few design methods are available for building robust, nonlinear control laws for this class of systems. We are currently concentrating on the control of a ducted fan engine. By focusing on a specific nonlinear system, we hope to generate new ideas and techniques for applying the tools of robust linear control theory to systems with strongly nonlinear behavior.

There is a large literature on vectored propulsion systems and they are gaining in 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 [5] (see also the survey article [4]). 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 improving existing (linear) control methodologies. Our goal is to explore the nonlinear nature of flight control systems in a laboratory setting. A similar experiment has been constructed by Hauser at the University of Colorado [6].

A picture of our experimental system is shown in Figure 1. The ducted fan consists of a high-efficiency electric motor with a 6-inch diameter propeller, capable of generating up to 9 Newtons of thrust. Flaps on the fan allow the thrust to be vectored from side to side and even reversed. It is mounted on a three degree of freedom stand which allows horizontal and vertical translation as well as unrestricted pitch angle. A detailed description of the fan, including models for the thrust as a function of flap angle and motor current is given in [3].

We present experimental comparisons between several different methodologies which are used to design controllers: traditional operating point linearization, feedback linearization, and linearization about multiple operating points. These experimental comparisons comprise the first step towards our eventual goal of synthesizing robust nonlinear controllers for high-performance operation over the full flight envelope. While our experimental system is considerably simpler than a real vectored thrust aircraft, it has many of the characteristics of existing systems—strongly nonlinear dynamics, limited actuator bandwidth, and input saturation, to name a few.

One of our main goals in presenting these comparisons is to indicate how well currently available design tools work on a specific physical system and to gain insight into the types of problems which limit the applicability of those tools. In




Figure 1: Ducted fan apparatus

this sense, the ducted fan experiment provides a testbed for nonlinear controller designs which can be used to evaluate and verify both different underlying control theories and the different underlying control software. We hope our experience will help motivate future development of theoretical and software tools which can be used for designing real-time controllers for physical systems.

This paper is organized as follows: In Section 2 we introduce the experimental setup and briefly describe the nonlinear model. We describe the implementation of the controllers in Section 3 and give a detailed evaluation of their performance on a set of three different trajectories in Section 4. Section 5 summarizes our main conclusions and indicates a number of directions for future work. More details can be obtained electronically at <http://avalon.caltech.edu/~dfan>.

2 Experimental Setup

2.1 Hardware

The design philosophy was to have a relatively simple ducted fan which could provide two dimensional vectored and reverse thrust. The ducted fan is mounted on a rotating arm, as shown in Figure 2, which limits the ducted fan's 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. An optical encoder with an angular resolution of $\pi/1000$ radians is mounted on each axis. The system is described in detail in [3].

The experiment is interfaced to an 80486 computer running an MS-DOS-based real-time kernel called Sparrow [7]. Custom hardware is used to read in joint angles via encoders and generate PWM signals for the R/C servos. A voltage-to-current amplifier powers the fan engine. Joint angles are read in at 200 Hz and their rates are estimated using a four state filter. The PWM signals are output at 50 Hz, which is the standard update rate for R/C servos.

Controllers are designed and simulated using Matlab on Sun workstations. Sparrow loads linear controllers directly from Matlab data files. Gain scheduled controllers, which are im-




Figure 2: Overview of the experimental apparatus

plemented as a set of linear controllers, are also loaded from Matlab data files. Nonlinear controllers implemented as Matlab S-functions require a small amount of revision before they are linked to Sparrow.

2.2 Modeling

The dynamics of the system are modeled using standard rigid body mechanics with the angles, α_i , and their velocities, $\dot{\alpha}_i$, as the states. The equations of motion for the system have the functional form

$$M(\alpha)\ddot{\alpha} + C(\alpha, \dot{\alpha})\dot{\alpha} + N(\alpha, \dot{\alpha}) = T(\alpha, F),$$

where $M(\alpha)$ is the generalized inertia matrix, $C(\alpha, \dot{\alpha})$ is the Coriolis matrix, $N(\alpha, \dot{\alpha})$ is the vector of gravity terms and drag forces, and $T(\alpha, F)$ is the vector of applied joint torques [3].

A simplified model for the system which accounts for many of the important nonlinearities is given by

$$\begin{aligned} (\eta - \delta s_2^2)\ddot{\alpha}_1 &= 2\delta s_2 c_2 \dot{\alpha}_1 \dot{\alpha}_2 + f_1 c_3 - f_2 s_3 \\ \beta \ddot{\alpha}_2 &= -\beta \gamma c_2 - \delta s_2 c_2 \dot{\alpha}_1^2 + f_1 s_3 + f_2 c_3 \\ J \ddot{\alpha}_3 &= -mgl s_3 + r f_1 \end{aligned} \quad (1)$$

where η , β , and δ are moments of inertia for the stand plus fan, γ is the effective gravitational constant (taking the counterweight into account), m is the mass of the fan, J the inertia of the fan about the α_3 axis, r is the distance from the flap assembly to the α_3 axis, and l is the offset of the center of mass from the α_3 axis. The terms c_i and s_i are shorthand for $\cos \alpha_i$ and $\sin \alpha_i$, respectively. The forces f_1 and f_2 are the forces perpendicular and parallel to the axis of the fan. Specific numerical values for the parameters in this model can be derived from the values given in [3].

Identification experiments have shown that this model is correct near hover and successful control designs also support this conclusion. For the control designs in this paper, we ignore 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, but are of little importance during hover and at slower speeds. The model also omits the gyroscopic terms that result from the angular momentum of the propeller. This term can be significant at moderate pitch (α_3) and high $\dot{\alpha}_1$.

There are several additional sources of known uncertainty in this model. A key simplifying assumption is 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, the forces generated by the flaps are not directly commanded. Instead, an experimentally determined lookup table maps desired forces to motor current and flap angle. This table is computed during fan calibration and assumed to be constant, even though the

actual mapping between motor current and flap angle varies as a function of operating conditions.

The strongly nonlinear features of the model include the effect of the pitch α_3 on the input forces and also the centrifugal forces ($\dot{\alpha}_1^2$ terms in $\ddot{\alpha}_2$ equation). In particular, the latter forces can be quite high when the fan is flying quickly, and they tend to push the altitude, α_2 , towards zero.

3 Control Schemes

Four different control schemes were examined, ranging from static state feedback to gain scheduled control. Initial controller designs were performed on the model described in [3]. While the controllers' performances are a function not only of the method, but also the designers' skills, meaningful comparisons can still be made.

Two of the control designs are based solely upon the plant's linearization about hover.

$$\begin{bmatrix} \alpha_1(s) \\ \alpha_2(s) \\ \alpha_3(s) \end{bmatrix} = \begin{bmatrix} \frac{-0.2102s^2 + 0.8134}{s^4 + 10.4814s^2} & 0 \\ 0 & \frac{0.19416}{s^2 + 0.6637} \\ \frac{-5.7085}{s^2 + 10.4814} & 0 \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \end{bmatrix} \quad (2)$$

At hover, the α_1 and α_2 axes are decoupled and the $\alpha_1(s)/f_1(s)$ transfer function is non-minimum phase. However, when $\alpha_3 \neq 0$, the system is not decoupled.

The control objectives were to track both α_1 and α_2 trajectories. As minimal requirements on the controllers, the closed loop transfer function from α_1 reference to α_1 error must be below -20 dB at frequencies below 0.05 radians/second and the closed loop transfer function from α_2 reference to α_2 error must be below -20 dB at frequencies below 0.1 radians/second.

3.1 Linear Quadratic Regulator

The linear quadratic regulator (LQR) design is performed on the error system of plant linearized about hover. The control authority provided by the amplifier and servos leads to unity input weighting. State weightings are chosen to allow moderate α_3 swings while zeroing the α_1 error because a large penalty on α_3 drastically slows the α_1 response. Some penalty is placed on α_3 since large pitch swings lead to large changes in the plant's dynamics. A large weighting is also placed on the α_2 error. The state derivatives have zero weights.

This initial controller gives unacceptable performance. The system naturally has little gain from f_2 to α_2 , as can be seen in (2). To meet the performance requirement, a constant gain of 30 is needed. Due to unmodeled dynamics, a gain this large drives the system unstable.

To achieve the necessary low frequency gain, an integrator is added to the α_2 axis and a second LQR design is performed on the augmented plant. The weights are chosen using similar guidelines as before, and the α_2 weight is chosen to be seven times greater than the weight on its integral.

3.2 \mathcal{H}_∞ Control

\mathcal{H}_∞ synthesis is a disturbance rejection methodology. To track trajectories, the performance goal is to reject low frequency components of the error signal. The tracking becomes faster as higher frequencies are rejected.

As compared with the LQR design, the \mathcal{H}_∞ structure is almost identical. A constant but large performance weight is placed on α_1 . Integral action is needed on α_2 since the natural gain of the channel is not large enough to track satisfactorily. Finally, a weight on input uncertainty is added to ensure controller roll off at high frequencies. The performance weight on α_2 seeks to make the dynamics of the response of α_2 much faster than the dynamics on α_1 . This was done in an effort to reduce the α_2 error induced when the fan is tilted at an angle. Sensor noise is explicitly considered in all measurements. The resulting controller has nine states.

3.3 Gain Scheduled Control

Since linear design techniques are well developed, one common practice is to linearize a plant about several different operating points, design a controller for each operating point, and switch between them. From (1), it is clear that the ducted fan's linearization is sensitive to changes in α_3 .

Three operating points potentially encountered during the trajectories of interest were chosen. They correspond to hover, acceleration, and deceleration for flight in the positive α_1 direction. At each operating point, an LQR controller was designed. The hover controller is identical to the previously described LQR controller.

Since the controllers have no states, switching transients due to improper initial conditions are avoided. Many types of switching were examined, including switching based upon nominal state, switching based upon actual state, and blending of control outputs (smooth switching).

3.4 Feedback Linearization Based Control

It is tedious but straightforward to check that the dynamics for the ducted fan are not full state feedback linearizable using static state feedback. However, the system is I/O linearizable with respect to any pair of outputs (α_i, α_j). It is thus possible to feedback linearize the dynamics and decouple the control of the pitch (α_3) and altitude (α_2) axes. In this case, the zero dynamics for the fan are the α_1 dynamics.

Using the same notation as the simplified model (1), applying the feedback linearizing law

$$\begin{bmatrix} f_2 \\ f_1 \end{bmatrix} = \begin{bmatrix} \frac{\psi}{c_3} \\ \frac{mgl s_3}{r} \end{bmatrix} + \begin{bmatrix} \frac{\beta}{c_3} \\ 0 \end{bmatrix} v_1 + \begin{bmatrix} \frac{-J \tan \alpha_3}{r} \\ \frac{1}{r} \end{bmatrix} v_2, \quad (3)$$

where $\psi = \beta \gamma c_2 + \delta s_2 c_2 \alpha_1^2 - \frac{mgl s_3^2}{r}$, yields the output response

$$\begin{bmatrix} \ddot{\alpha}_1 \\ \ddot{\alpha}_2 \\ \ddot{\alpha}_3 \end{bmatrix} = \begin{bmatrix} \frac{\tan \alpha_3}{\eta - \delta s_2^2} (v + \beta v_1 + \frac{1}{r} (c_3^2 / s_3 - s_3) v_2) \\ v_1 \\ v_2 \end{bmatrix}, \quad (4)$$

where $v = 2\delta s_2 c_2 \alpha_1 \dot{\alpha}_2 / \tan \alpha_3 + mgl c_3^2 / r - \psi$.

Standard I/O linearization based control design uses pole placement to stabilize the system. This design method was not successful for the ducted fan since I/O linearization is essentially a plant inversion and due to unmodeled dynamics, the inversion is not exact. The system behaves as two poles near the origin whose position changes as the plant's state varies. To stabilize each axis, a second order lead filter was used to add phase margin at the desired crossover frequency.

The α_1 control loop is designed around the previous closed loop system using a separation of time scales argument and some simplifying assumptions. First, it is assumed that α_3 can be commanded directly and used as an input to control α_1 . The second assumption, $\alpha_2 = 0$, removes many of the stand's effects from the equations of motion. Third, it is assumed that the ducted fan is operating at constant α_2 and α_3 , so v_1, v_2 , and $\dot{\alpha}_2 = 0$. With these assumptions, the $\ddot{\alpha}_1$ equation from (4) becomes

$$\ddot{\alpha}_1 \approx \frac{1}{\eta} \left(\frac{mgl}{r} - \beta \gamma \right) \tan \alpha_3. \quad (5)$$

Finally, using the approximation $\tan \alpha_3 \approx \alpha_3$, which is valid for small α_3 , the simplified system is

$$\begin{bmatrix} \ddot{\alpha}_1 \\ \ddot{\alpha}_2 \end{bmatrix} = \begin{bmatrix} k \alpha_3 \\ v_1 \end{bmatrix}, \quad k = \frac{1}{\eta} \left(\frac{mgl}{r} - \beta \gamma \right) \quad (6)$$

For this experiment, $k \approx 0.52$.

The α_1 controller generates an α_3 command based upon the α_1 error and $\dot{\alpha}_1$ error. To ensure that $\tan \alpha_3 \approx \alpha_3$, a saturation element at ± 1.2 radians is placed on the command. The saturation element limits performance, but improves stability. For a large step disturbance, a linear controller would generate a large command. However, commands larger than $\pi/2$ are

Criteria	Controllers			
	LQR	\mathcal{H}_∞	Gain S	I/O D.
α_1 rise time	3.34 ²	3.42 ³	3.24 ¹	3.68 ⁴
α_1 90% delay	0.62 ²	1.16 ³	0.48 ¹	1.62 ⁴
α_1 % overshoot	7.24 ²	3.15 ¹	8.81 ³	18.87 ⁴
steady state e_1	0.03 ³	0.02 ¹	0.02 ¹	0.07 ⁴
$\ e_1\ _\infty$	0.26 ²	0.34 ³	0.25 ¹	0.41 ⁴
$\ e_1\ _2$	2.53 ²	3.96 ³	2.35 ¹	6.09 ⁴
$\ e_1\ _1$	53.59 ²	73.51 ³	48.91 ¹	136.64 ⁴
$\ e_2\ _\infty$	0.05 ²	0.05 ²	0.07 ⁴	0.03 ¹
$\ e_2\ _2$	0.51 ³	0.44 ²	0.68 ⁴	0.37 ¹
$\ e_2\ _1$	10.42 ³	8.11 ¹	14.87 ⁴	8.79 ²
$\ \alpha_3\ _\infty$	1.28 ⁴	1.12 ²	1.20 ³	0.80 ¹
$\ \alpha_3\ _\infty$	4.14 ⁴	2.29 ²	3.35 ³	2.05 ¹

Table 1: Comparison of controllers for a ramp in α_1 . The difference between α_i commanded and α_i measured is denoted by e_i .

obviously not desired. Furthermore, the decoupler is singular when $\alpha_3 = \pm \pi/2$.

4 Experimental Comparisons

The controllers were tested on three trajectories. Two of the trajectories are simple, commanding changes to one axis at a time. The third trajectory commands rapid changes to both the α_1 and α_2 axes.

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 demonstrate the controllers' abilities to track each axis independently.

The third trajectory is more complex, commanding the fan to fly rapidly in the positive α_1 direction while α_2 is varied. During forward flight, the fan achieves $\dot{\alpha}_1$ of 0.628 radians per second, more than three times greater than during the first trajectory.

The results for each of the controllers are presented in a figure with 4 plots. The upper left plot presents the α_1 trajectory (called "X" in the plot) and the system's response. The lower left plot shows the α_2 trajectory (called "Y" in the plot) and the system's response. The upper right plot shows the α_3 behavior. For the decoupled controller results, Figures 6, 10, and 14, this plot also contains the α_3 command. The final plot shows the commanded forces. The axial force, f_2 , has a DC component of 2.65 Newtons to maintain hover.

Quantitative results are summarized in three tables, one for each trajectory. The controllers' rankings with respect to the performance measures are indicated by superscript numbers in each table entry. The 10-90% rise time is a standard figure of merit for step responses. For the α_1 ramp, it provides a measure of how closely the ramp follows the signal. The 90% delay factor is computed by measuring the difference, in seconds, between when the trajectory reached 90% of its final value, and when the system reached 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 remaining figures of merit are self-explanatory. For all measures, e_i is defined as the difference between the α_i commanded and α_i measured.

4.1 Ramp in α_1 (horizontal position)

This trajectory consisted only of a slow change to one axis, and all controllers performed well. As shown in Table 1, the responses are similar. The most notable differences are in the delay to 90% of the final value and the α_3 measures.

Since the LQR controller has no filtering states, any step changes at the inputs are passed directly to the plant. In Figure 3, the ramp corresponds to a step change on $\dot{\alpha}_1$ and abrupt changes are seen in the commanded forces. This causes α_3 to have larger variations than the other controllers.

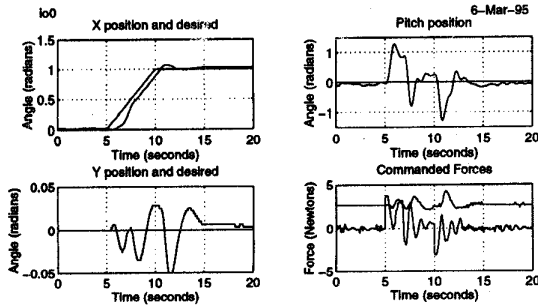


Figure 3: Closed loop response of LQR controller for α_1 ramp

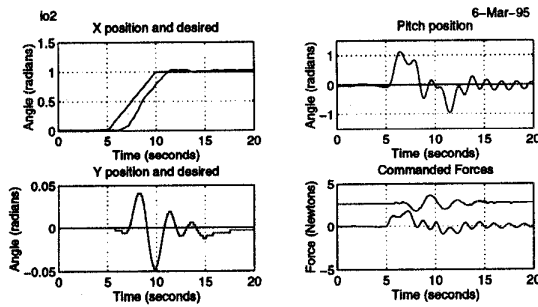


Figure 4: Closed loop response of \mathcal{H}_∞ controller for α_1 ramp

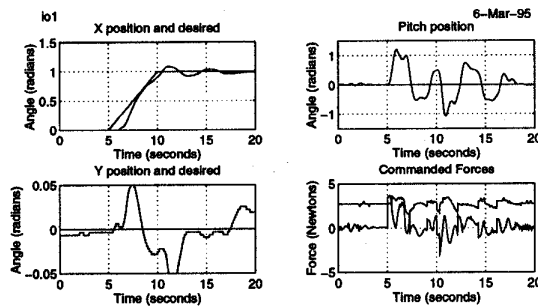


Figure 5: Closed loop response of gain scheduled controller for α_1 ramp

The response of the \mathcal{H}_∞ controller for the ramp on α_1 is shown in Figure 4. The \mathcal{H}_∞ controller has a much larger delay than the LQR controller, but is otherwise identical.

Various switching techniques were examined on the gain scheduled controller. The differences between smooth blending of output and abrupt switching were small. For all results in this paper, switching occurred at $|\alpha_3| > 0.4$. The gain scheduled controller's performance, shown in Figure 5, is similar to the LQR controller. This is expected, since the gain scheduled controller at hover is identical to that of the LQR controller.

The I/O decoupling controller tracks the ramp slowly, but with a very smooth α_3 response, as shown in Figure 6.

4.2 Step in α_2 (altitude)

All controllers were able to track this trajectory well. The design requirement ensured steady state errors would be small, and this was the case. Table 2 summarizes the results. Due to encoder quantization, the minimum detectable error for this trajectory is a 3.1% error. Values less than this should be considered zero.

With the addition of an integrator on α_2 , the LQR controller tracks the step with essentially zero steady state error, as shown in Figure 7. The performance is not drastically different from that of the other controllers. The response of the \mathcal{H}_∞ controller to the step in α_2 is shown in Figure 8 and the

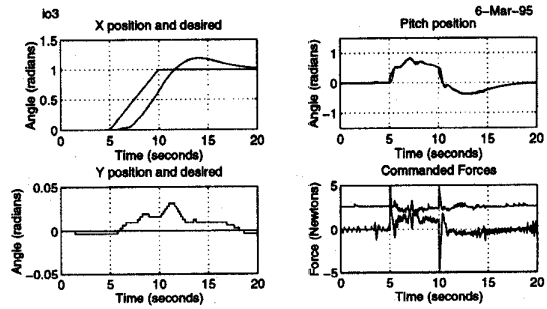


Figure 6: Closed loop response of I/O decoupled controller for α_1 ramp

Criteria	Controllers			
	LQR	\mathcal{H}_∞	Gain S	I/O D.
α_2 rise time	1.12 ⁴	0.96 ²	1.10 ³	0.64 ¹
α_2 90% delay	1.90 ⁴	1.72 ²	1.76 ³	1.02 ¹
α_2 % overshoot	9.96 ¹	22.52 ³	19.38 ²	25.66 ⁴
steady state e_2	0.00 ¹	0.00 ¹	0.00 ¹	0.01 ⁴
$\ e_2\ _2$	0.75 ⁴	0.74 ³	0.71 ²	0.58 ¹
$\ e_2\ _1$	10.03 ³	10.26 ⁴	8.64 ¹	9.78 ²
$\ e_1\ _\infty$	0.01 ²	0.04 ⁴	0.02 ³	0.00 ¹
$\ e_1\ _2$	0.18 ²	0.82 ⁴	0.23 ³	0.07 ¹
$\ e_1\ _1$	4.84 ²	22.33 ⁴	5.90 ³	1.58 ¹
$\ \alpha_3\ _\infty$	0.09 ²	0.21 ⁴	0.09 ²	0.05 ¹
$\ \alpha_3\ _2$	0.35 ²	0.51 ⁴	0.41 ³	0.28 ¹

Table 2: Comparison of controllers for a step in α_2 . The difference between α_i commanded and α_i measured is denoted by e_i .

response of the gain scheduled controller is shown in Figure 9.

Even without an integrator in the I/O decoupling controller, there is still sufficient gain to complete the step, shown in Figure 10. However, since there is no integrator in the loop, there is a small steady state error. The I/O decoupling is essentially a plant inversion. Due to plant uncertainty, the decoupled system does not behave as an integrator chain. In order to achieve zero steady state error, an integrator must be added to the controller.

4.3 Complex Trajectory

This trajectory demonstrated the differences between the linear and nonlinear controllers. The trajectory takes the fan far from hover, the point about which all linear controllers were designed. Table 3 summarizes the results.

Figure 11 shows the response of the LQR controller for the complex trajectory. This controller tracks the α_1 revolutions very well. However, it is not able to track the α_2 oscillations. The α_2 changes occur while the plant is not at $\alpha_3 \approx 0$, the point for which the controller was designed. Due to the structure of




Figure 7: Closed loop response of LQR controller for α_2 step

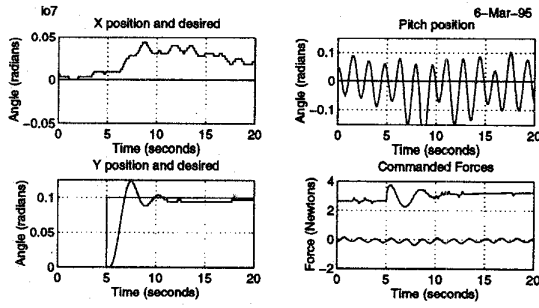


Figure 8: Closed loop response of \mathcal{H}_∞ controller for α_2 step

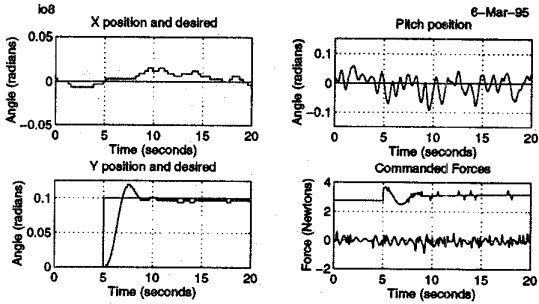


Figure 9: Closed loop response of gain scheduled controller for α_2 step

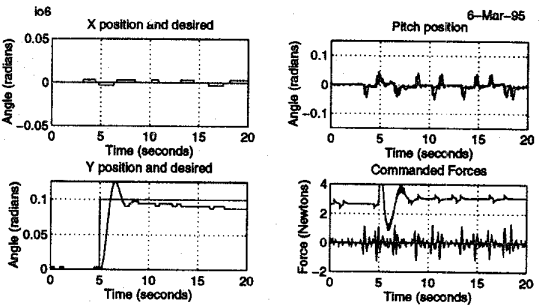


Figure 10: Closed loop response of I/O decoupled controller for α_2 step

Criteria	Controllers			
	LQR	\mathcal{H}_∞	Gain S	I/O D.
e_1 at $t=25$	0.39 ³	0.60 ⁴	0.18 ¹	0.37 ²
$\ e_1\ _\infty$	0.41 ²	0.73 ³	0.26 ¹	0.74 ⁴
$\ e_1\ _2$	11.88 ²	19.42 ⁴	6.33 ¹	13.50 ³
$\ e_1\ _1$	404.06 ²	658.47 ⁴	213.93 ¹	441.81 ³
$\ e_2\ _\infty$	0.25 ⁴	0.14 ¹	0.20 ³	0.16 ²
$\ e_2\ _2$	4.35 ⁴	1.60 ¹	2.15 ²	2.63 ³
$\ e_2\ _1$	122.95 ⁴	43.45 ¹	48.16 ²	72.59 ³
$\ \alpha_3\ _\infty$	1.49 ⁴	1.42 ³	1.08 ²	1.04 ¹
$\ \dot{\alpha}_3\ _\infty$	2.03 ³	3.00 ⁴	1.49 ²	0.98 ¹

Table 3: Comparison of controllers for the complex trajectory. The difference between α_i commanded and α_i measured is denoted by e_i .

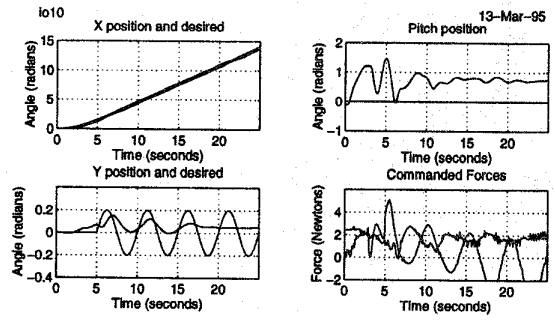


Figure 11: Closed loop response of augmented LQR controller for complex trajectory

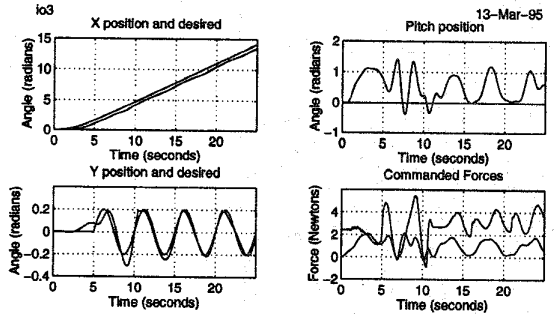


Figure 12: Closed loop response of \mathcal{H}_∞ controller for the complex trajectory.

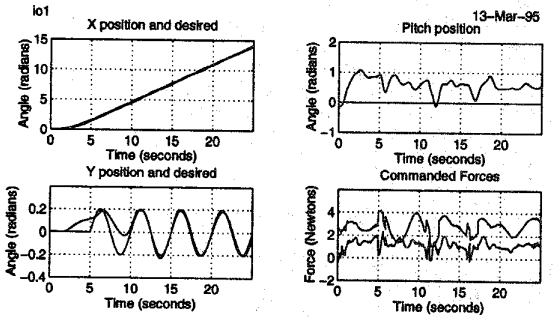


Figure 13: Closed loop response of gain scheduled controller for complex trajectory

(2), any error in α_2 is controlled by changing only the motor current. However, when $\alpha_3 \neq 0$, this induces both an α_3 and an α_1 change. At large α_3 , the motor current change has little effect on α_2 . In addition, the integrator on the α_2 error winds up, and the commanded forces quickly saturate the actuators.

The response of the \mathcal{H}_∞ controller, is shown in Figure 12. Due to different weight selection, the controller performs better on α_2 , but worse on α_1 .

The gain scheduled controller performs extremely well on all measures. Once the initial transients die, it has the best performance of all the controllers.

One notable item in Figure 14 is the relatively smooth α_3 response. Due to the structure of the controller, α_3 does not have large oscillations. The saturation element also ensures that excessive α_3 values are not achieved. This type of behavior is desirable, since it corresponds to a smoother set of commanded forces.

4.4 Results Summary

All controllers are able to follow the ramp in α_1 . The gain scheduled controller performs the best of all the controllers.

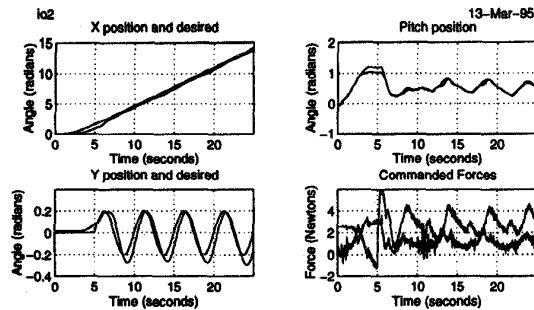


Figure 14: Closed loop response of I/O decoupled controller for complex trajectory

The \mathcal{H}_∞ and LQR designs have similar performance, with the \mathcal{H}_∞ design performing slightly better than the LQR design at the expense of nine states. The I/O decoupling controller performs the worst overall, although its α_2 performance over the trajectory is better than the other controllers.

All controllers do a good job on the step in α_2 as well. The I/O decoupling controller performs best. The gain scheduled controller has the second best performance, followed quite closely by the LQR controller. The \mathcal{H}_∞ controller performs the worst. It is interesting that the I/O decoupler performs the best, since no decoupling is needed for a step on α_2 . The fact that the gain scheduled and LQR designs perform quite similarly is not surprising either, as no scheduling occurs on this trajectory.

Three controllers perform reasonably well on the complex trajectory. The exception is the LQR controller, which has an integrator wind-up problem, and thus is unable to follow the α_2 changes in the trajectory. This is clearly evidenced by the Figure 14. The best controller overall is the gain scheduled controller, which has the best α_1 tracking performance while maintaining good α_2 performance. The \mathcal{H}_∞ controller tracks the α_2 portion of the trajectory very well, but suffers on the α_1 response. The I/O decoupling controller achieves consistently good results overall rather than being best at one particular performance criterion.

5 Summary and Conclusions

The objective of this project is to study the application of different methodologies for flight control to a relatively simple nonlinear system with nonlinearities similar to those found on vectored thrust aircraft. We draw the following conclusions from the work described in this paper:

1. *Scheduled designs achieve better tracking performance than linear designs.* For aggressive trajectories over a wide operating envelope, the gain scheduled controller performs better than linear controllers. This, combined with other experience on the fan [2], suggests that pure linear techniques are not sufficient for high performance on this experiment.
2. *Good tracking performance is achievable using (good) first principles models.* The controllers designed in this study are based on first principles models of the dynamics. The use of system identification techniques is limited in this context because of the need for (nonlinear) models which hold over the entire operating envelope. When poor models are available, there is little difference in overall performance between the different linear and nonlinear techniques.
3. *Gain scheduling performs as well as or better than feedback linearization.* Simple gain scheduled controllers built from three pointwise linearized models performed better than much more complicated nonlinear controllers on most of the trajectories tested. This shows that local linearizations can often capture most of the nonlinear effects which appear in flight control systems. More

advanced scheduling techniques using linear parameter-varying synthesis (LPV) perform even better than the simple gain scheduled controllers studied here (see [2]), supporting this claim.

4. *Nonlinear techniques need to be more closely coupled with good linear control design.* Many nonlinear control techniques implicitly rely on the use of pole placement as part of the overall control law synthesis. Nonlinear designs based on pole placement do not perform well on the ducted fan and more advanced nonlinear techniques must take advantage of the wealth of good linear design techniques which are available for achieving small-signal performance around an operating point.

Future work on the ducted fan will improve upon the controllers presented here and continue to address the practical issues involved in good nonlinear control design. In particular, we will focus on the modeling of the system by including the gyroscopic effects of the blade, aerodynamic drag terms, and results from identification experiments in the model. With improved models, we expect to design controllers capable of achieving exceptional performance over aggressive trajectories.

Two specific tasks currently underway are the application of LPV synthesis techniques [2] and real-time trajectory generation with two-degree of freedom nonlinear control design [10]. These are complementary techniques for building nonlinear controllers which exploit the known nonlinearities of the system while giving guaranteed performance over a large range of operating conditions. The ducted fan also serves as an experimental testbed for new theoretical developments. This work includes the application of nonlinear robustness analysis [9, 8] and model reduction techniques for parameter-varying models [1].

References

- [1] C. Beck, B. Bodenheimer, and P. Bendotti. Lmi-based model reduction for a vectored thrust ducted fan experiment. In *Proc. IEEE Control and Decision Conference*, 1995. (submitted).
- [2] B. Bodenheimer, P. Bendotti, and M. Kantner. Linear parameter-varying control of a ducted fan engine. Technical Report CIT-CDS 95-004, California Institute of Technology, January 1995.
- [3] H. Choi, P. Sturdza, and R. M. Murray. Design and construction of a small ducted fan engine for nonlinear control experiments. In *Proc. American Control Conference*, pages 2618-2622, 1994. Available electronically from <http://avalon.caltech.edu/~dfan>.
- [4] B. Gal-Or. Fundamental concepts of vectored propulsion. *Journal of Propulsion*, 6(6):747-757, 1990.
- [5] B. Gal-Or. *Vectored Propulsion, Supermaneuverability and Robot Aircraft*. Springer-Verlag, 1990.
- [6] C. Lemon and J. Hauser. Design and initial flight test of the champagne flyer. In *Proc. IEEE Control and Decision Conference*, pages 3852-3859, 1994.
- [7] R. M. Murray and E. L. Wernhoff. *Sparrow 2.0 Reference Manual*. California Institute of Technology, 1994. Available electronically from <http://avalon.caltech.edu/murray/~sparrow>.
- [8] J. Tierno and R. M. Murray. Robust performance analysis for a class of uncertain nonlinear systems. In *Proc. IEEE Control and Decision Conference*, 1995. (submitted).
- [9] J. Tierno, R. M. Murray, and J. C. Doyle. An efficient algorithm for performance analysis of nonlinear control systems. In *Proc. American Control Conference*, 1995. (to appear).
- [10] M. van Nieuwstadt and R. M. Murray. Approximate trajectory generation for differentially flat systems with zero dynamics. In *Proc. IEEE Control and Decision Conference*, 1995. (submitted).