JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 19, No. 3, May–June 1996

Design of an Autoland Controller for an F-14 Aircraft Using $\mathcal{H}_\infty$ Synthesis

R. J. Niewoehner* and Isaac I. Kaminer†
U.S. Naval Postgraduate School, Monterey, California 93943

Our purpose is to present the results of a design effort to incorporate direct lift control into the F-14 aircraft automatic carrier landing system. This was accomplished using a new design methodology whereby typical single input/single output design specifications were translated into the weighting functions of an $\mathcal{H}_\infty$ output-feedback synthesis problem. Next, robustness of the $\mathcal{H}_\infty$ controller was analyzed using the structured singular value concept (μ). A simplified uncertainty model for air vehicles based on flight test data was developed for this analysis. Finally, the resulting controller was successfully validated on the nonlinear simulation.

I. Introduction

CARRIER approach and landing is a challenging multivariable control problem where the aircraft states must all be carefully controlled to comply with multiple structural and safety-of-flight constraints. Automatic landing systems currently in service on carrier-based aircraft (including the F-14) incorporate nested single input/single output (SISO) controllers that generally seek to regulate the angle of attack with the engines in the inner loop, while aerodynamic surfaces such as elevators or stabilators provide altitude control. Since neither the engines nor the aerodynamic surfaces can control the altitude state directly, their influence is indirect through a combination of other states. The F-14 has an unusual aerodynamic configuration that includes direct lift control (DLC), thereby providing an aero surface with substantial authority to control altitude directly. Moreover, the DLC is driven by actuators whose bandwidth exceeds that of the other control surfaces. Regrettably, this powerful control surface is not used by the automatic landing system currently in service. In this paper, a multivariable feedback controller is presented that seeks to exploit this powerful but dormant capability.

To achieve this objective, the design methodology presented here was developed to enable the control engineer to translate the design requirements into weighting functions for $\mathcal{H}_\infty$ synthesis.[1–5] It has been our experience when dealing with these problems that most design requirements are SISO in nature, whereas the $\mathcal{H}_\infty$ synthesis technique is truly a multivariable tool. Thus, the main feature of this methodology is a simple procedure for translating the SISO requirements into the various weighting functions for $\mathcal{H}_\infty$ synthesis. Moreover, once the SISO requirements have been satisfied, the $\mathcal{H}_\infty$ framework offers a natural way to expand the weighting functions to satisfy multivariable stability and performance robustness requirements.

This methodology has been applied to the design of control systems for commercial airplanes, autonomous underwater vehicles, flexible structures, and, most recently, unmanned aerial vehicles.[6–10] In Refs. 6–10, this technique was used to synthesize state-feedback controllers. In Ref. 9, where a controller for a flexible structure was designed, the methodology was extended to include an output feedback case. The methodology outlined here expands this previous work to include compliance with closed-loop sensor bandwidth requirements.

Specifically, the methodology offers a simple and effective way to design feedback controllers that satisfy specified command-loop bandwidths, control-loop bandwidths, closed-loop damping, and closed-loop sensor bandwidths. Each of these objectives is pursued through a specific formulation of the synthesis model and selection of the $\mathcal{H}_\infty$ synthesis weighting functions. Furthermore, it has been observed that an additional benefit of this methodology is that the resulting controllers do not cancel the undesirable modes of the open-loop plant. This is attributed to the suitable choice of weights to satisfy the closed-loop damping requirements.

The appeal of this methodology is that the control designer is provided with a straightforward framework to implement $\mathcal{H}_\infty$ controllers in pursuit of typical design requirements without a detailed understanding of the theoretical basis for these tools. Moreover, the results of the design effort are assessed using familiar SISO figures of merit. The availability of good commercial software utilities and ever improving computational resources only enhance the viability of iterative design methods such as the one presented here. This methodology is suggested as one means of placing these tools into the hands of practicing control designers.

The controller design methodology we propose necessarily has a heuristic component. This heuristic component has been influenced by our experience in solving the practical problems already mentioned. Although we were successful in applying the design methodology to this problem, we cannot offer any guarantees as to whether its application to an arbitrary control problem will yield a satisfactory solution.

Lastly, a method is proposed for accommodating structured model uncertainty in flight dynamics problems. Specifically, the gross uncertainties in lift, drag, and pitching moment in the nonlinear model are considered rather than the uncertainties in individual parameters. This approach is suggested for use when flight test is the origin of the aerodynamic stability data. The robustness of the resulting closed-loop nonlinear system is then analyzed using established structured singular value (μ) methods.[11,12]

This paper is organized as follows. In Sec. II, the carrier landing problem and the design requirements are described. Section III presents the details of the design process, and Sec. IV details the uncertainty modeling and analysis results. Section V presents the results of the nonlinear simulation. We present some concluding remarks in Sec. VI.

II. Problem Statement

The objective of the controller design is to provide precise automatic control of the approach and landing of a carrier-based aircraft in the longitudinal axis. In this section we describe both the plant to be controlled and the desired performance specifications. These specifications are classical in nature and are representative of those posed to control designers in industry. The notational convention adapted in this paper was to use uppercase letters of denote total values of the variables introduced, whereas lower case letters denoted the perturbations of these variables around their trim values.

Received July 5, 1994; revision received June 12, 1995; accepted for publication Nov. 11, 1995.

*Commander, U.S. Navy.
†Assistant Professor, Department of Aeronautics and Astronautics. Member AIAA.

A. Airplane and Model Description

The design problem to be solved here deals with the longitudinal motion of a fighter airplane and the control of the longitudinal rigid body dynamics. A complete description of an airplane's equations of motion can be found in many available references: see, for example, Ref. 13. Consequently, we will not describe the equations of motion in detail. Rather, we will present a brief qualitative description of the key features.

The longitudinal equations of motion of an airplane are described by two force equations (longitudinal force F_x and vertical force F_z) and one angular moment equation (pitching moment M). The state variable associated with the F_x equation is the forward velocity U (along airplane's body-fixed x direction). The state variable associated with the F_z equation is the angle of attack α (the angle between the body-fixed x direction and the true total velocity). The state variable associated with the M equation is the pitch rate Q. The integral of Q for a typical approach and landing condition is the pitch attitude Θ (the angle between the body fixed x direction and the horizon). Other motion variables of interest are the airplane's airspeed V_t (generally not aligned with the body-fixed x direction), flight-path angle γ (the angle between V_t and the horizon), and the airplane's altitude above sea level H. All of the angles used here are expressed in radians, angle rates in radian per second, position variables in feet, position rates in feet per second, and accelerations in gravity.

Although the preceding discussion is germane to the flight dynamics of all aircraft, the F-14 has a distinctive aerodynamic configuration because of the environment in which it was designed to operate. Mission requirements dictated a variable-sweep wing for low drag at high speeds, with full-span flaps and slats, for very high lift at low carrier takeoff and landing speeds. Control along the longitudinal axis is provided by two afterburning turbofan engines (thrust). Pitch control is provided by symmetric deflection of two stabilators (stab). In the landing configuration (gear down, wings fully forward, and flaps fully extended), DLC is provided by symmetric deflection of wing mounted spoilers. The neutral position of the spoilers is approximately 40% of full deflection so as to provide for both positive and negative contributions to the lift.

For this design study, the sensors available included onboard accelerometers and gyros that provide pitch attitude Θ, longitudinal acceleration N_x, vertical acceleration N_z, and pitch rate Q. Total velocity V_t was provided by the aircraft's air data system. Lastly, for automated approaches and landings, the altitude H was determined by a shipboard tracking radar.

The model used for the design process was a linear model obtained from a nonlinear simulation built using aerodynamic coefficient data. This simulation model was nonlinear in that although the aerodynamic derivatives were held constant, the equations of motion included the nonlinear influence of airspeed and gravity, as well as the dynamic coupling terms. The flight condition was a nominal approach point of 230 fps, at sea level, with a gross weight of 54,000 lb. The linearized longitudinal model included five states, u, α, q, θ, and h, and three control inputs, δ_{stab}, δ_{thrust}, and δ_{dlc}, where all of the linear states and inputs are small perturbations around the nominal operating point. At this condition, the longitudinal rigid body motion of the F-14 is characterized by two second-order stable modes, the phugoid and short period, and an altitude integrator. The phugoid involves perturbations in V_t and H with nearly constant α, whereas the short period mode involves perturbations in α and Q, with V_t and H remaining constant. The short period mode had a natural frequency of 1.04 rad/s and a damping ratio of 0.45. The phugoid had a natural frequency of 0.18 rad/s and a damping ratio of 0.06. In addition to the plant's five states, the actuators were modeled by three first-order transfer functions. These were appended to the control inputs and had bandwidths of 20, 2.5, and 50 rad/s for the stabilators, engine, and DLC, respectively. As a result the complete system was represented by an eighth-order linear model.

B. Problem Description

The general problem was to design a feedback controller that would satisfy the operational constraints imposed by the mission. The challenge of landing an aircraft at sea requires very precise control of the aircraft states. Glideslope, which is the desired spatial trajectory of the aircraft, must be tightly controlled to provide for safety and to achieve the precise touch-down necessary to be arrested on the ship. The glideslope is an imaginary ramp oriented three degrees above the horizon, moving with the ship and terminating in the center of the landing area. In the case of manual landings, deviations from glideslope are detected visually by the pilot with a visual reference to a shipboard optical system. For automated landings, a precision tracking radar onboard the ship compares the aircraft's position with an internally calculated glideslope and transmits an error signal to the aircraft's flight control system via data link. Tight control of aircraft total velocity is driven by the competing requirements of providing for adequate aerodynamic performance, while minimizing the kinetic energy that the airframe and arresting gear must absorb on landing. Tight control of the aircraft attitude is required to prevent tailstrike. Both of these objectives can be achieved by controlling angle of attack with total velocity and the pitch attitude as dependent variables. Angle of attack is also an attractive control variable as consistent aerodynamic performance is achieved for wide ranges of gross weights.

The approach-to-landing problem can be fully characterized by several combinations of the variables in the state vector because of the mathematical and aerodynamic relationships between these variables. Likewise the control objectives can be achieved by tracking any one of these combinations. Systems in fleet use today incorporate nested SISO controllers, with the engine controlling angle of attack in the inner-loop and the stabilators controlling either sink rate or flight path. The remaining variables are then dependent functions of the gross weight and the two controlled parameters. In this paper we propose to use DLC to provide independent altitude control of the F-14 in approach and landing. Currently, the F-14's DLC is engaged for approach, such that the spoilers are deployed to their neutral DLC position. DLC is not utilized as part of the control system, however, and serves only to increase both the drag and the trimmed power setting (this is done to keep the engines in a more responsive range of operation). Neither the engines nor the stabilators provide direct control over the flight path, but instead indirectly control it through airspeed and pitch attitude. Performance may thereby be sacrificed, as DLC is the only control effector that couples directly into altitude, the most critical of the three parameters, and DLC actuator is the fastest of the three available actuators. Given three independent control effectors with sufficient control power, F-14 has the resident capability to track three independent command signals. A multivariable approach to the control design would permit inclusion of the DLC in the control system resulting in both an enhanced capability and an improvement in performance. Therefore our control strategy was to track altitude h and angle of attack α using stabilators, engines, and DLC. Since the number of control inputs exceeds the number of command variables, we have the flexibility to wash out one of the inputs. We chose to wash out the DLC. (Note that this is equivalent to require that DLC be driven to zero in steady state.) The desired effect was that the thrust would control the glideslope in steady state, whereas the DLC would provide dynamic glideslope control. Thus, the design problem was to synthesize a feedback controller that tracked a glideslope signal, while controlling angle of attack in steady state given the available sensor suite.

C. Design Requirements

In view of the problem description, the controller was required to satisfy the following design requirements.

1) The controller had to achieve zero steady-state values for DLC and for the errors in h and α in response to step commands in angle of attack and ramp commands in altitude (this was necessary for glideslope signal tracking and wind disturbance rejection).

2) The bandwidth requirements were as follows: a) The input-output command response bandwidth for all three command loops (altitude, angle of attack, and DLC) was not to exceed 1 rad/s. b) The control loop bandwidth was not to exceed 40 rad/s for the DLC actuator, 20 rad/s for symmetric stabilator, and 2 rad/s for the engine. These numbers represented 80% of the corresponding actuator bandwidths to ensure that the actuators were not driven beyond their linear operating range. c) The sensor response bandwidths were not

to exceed 100 rad/s for the gyros and accelerometer and 5 rad/s for the altitude, angle of attack, and airspeed data.

3) The closed-loop eigenvalues associated with physical states were to have the damping ratio of at least 0.6. (This permits controller modes to have damping ratios less than 0.6.)

4) The robustness requirements were as follows: a) The controller could not cancel the lightly damped open-loop poles of the plant. b) Simultaneous gain and phase margins of ±6 dB and 45 deg were to be achieved in all control and sensor loops. c) Stability was to be guaranteed for simultaneous variations of 20% in the perturbed lift and drag forces and pitching moment.

III. Controller Design

In this section, we will describe the key features of the controller design process that we followed. The section is organized into a number of subsections that emphasize some of the important engineering issues that arose in the controller design.

A. Synthesis Model

The first step in the controller design process was the development of the synthesis model that served as an interface between the designer and the $\mathcal{H}_\infty$ controller synthesis algorithm.

Consider the feedback system in Fig. 1. The synthesis model was derived from the linear model of the airplane by appending the depicted weights. The weights became the knobs that the designer adjusted to achieve the specified performance requirements. Here C denoted the controller to be designed, P was the linear model of the F-14, and the block $\mathcal{G}$ within the dashed line was the synthesis model. The vector $\boldsymbol{w}_1$ represented the commanded inputs to be tracked,

$$\boldsymbol{w}_1 = \begin{bmatrix} h_{\text{cmd}} & \alpha_{\text{cmd}} & \delta_{\text{dlc}_{\text{cmd}}} \end{bmatrix}'$$

The vector $\boldsymbol{w}_2$ consisted of the measurements noise inputs for each of the aircraft sensors and of the process noise inputs for each of the aircraft control effectors. The vector $\boldsymbol{u}_c$ represented the control inputs to the aircraft and was composed of the stabilator control input, the thrust control input, and the DLC control input. The signals $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ were

$$\boldsymbol{x}_1 = (h \quad \alpha \quad \delta_{\text{dlc}})' \qquad \boldsymbol{x}_2 = (\dot{u} \quad \dot{\alpha} \quad \dot{q} \quad \dot{\theta} \quad \dot{h})'$$

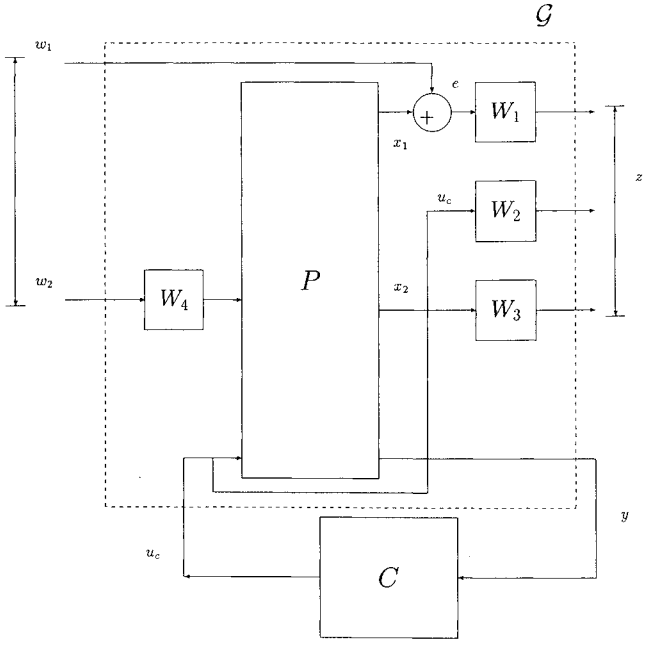

Fig. 1 Synthesis model.

The signal $\boldsymbol{e}$ represented the vector of the tracking errors ($\boldsymbol{e} = \boldsymbol{w}_1 - \boldsymbol{x}_1 =: [h_e, \alpha_e, \delta_{\text{dlc}_e}]'$).

The outputs of W_1, W_2, and W_3 comprised the vector $\boldsymbol{z}$. To achieve zero steady-state errors in tracking a ramp altitude command, and step α and DLC commands, the weighting function W_1 was chosen to have the following form:

$$W_1 = \begin{pmatrix} c_1/s^2 & 0 & 0 \\ 0 & c_2/s & 0 \\ 0 & 0 & c_3/s \end{pmatrix}$$

where the constants c_1, c_2, and c_3 were adjusted to get the desired command-loop bandwidths. Thus, W_1 provided weighted integrators on the error channels of the regulated variables.

The choice for the weighting function W_2 was

$$W_2 = \begin{pmatrix} c_4 & 0 & 0 \\ 0 & c_5 & 0 \\ 0 & 0 & c_6 \end{pmatrix}$$

where c_4, c_5, and c_6 were adjusted to achieve the desired control-loop bandwidths.

Next, notice that the elements of the vector $\boldsymbol{x}_2$ were the rate terms on the principal states of the aircraft. Selection of $\boldsymbol{x}_2$ was an important element of our methodology. Applying the weighting matrix W_3 to $\boldsymbol{x}_2$, and including these signals in $\boldsymbol{z}$, penalized activity in their corresponding states. As a result, W_3 had the following form:

$$W_3 = \text{diag}(c_i), \qquad i = 7, \ldots, 11$$

where c_i were used to improve damping in a particular mode, as discussed in Sec. III.B.

Therefore, we obtained

$$\boldsymbol{z} = \left(\frac{c_1 h_e}{s^2} \quad \frac{c_2 \alpha_e}{s} \quad \frac{c_3 \text{dlc}_e}{s} \quad c_4 \delta_{\text{stab}} \quad c_5 \delta_{\text{thrust}} \quad c_6 \delta_{\text{dlc}} \quad c_7 \dot{u} \quad c_8 \dot{\alpha} \quad c_9 \dot{q} \quad c_{10} \dot{\theta} \quad c_{11} \dot{h} \right)'$$

This selection of the regulated output vector $\boldsymbol{z}$ guaranteed that the appropriate integral control was applied to achieve steady-state requirements and that a mechanism existed for providing adequate damping on excessively active states.

The vector $\boldsymbol{y}$ included the aircraft's sensor outputs. Furthermore, $\boldsymbol{y}$ had to include the integral error states to satisfy the detectability assumption in Theorem 3 of Ref. 1:

$$\boldsymbol{y} = \left[h \quad \alpha \quad v_t \quad \theta \quad n_z \quad n_x \quad q \quad \frac{h_e}{s^2} \quad \frac{\alpha_e}{s} \quad \frac{\delta_{\text{dlc}_e}}{s} \right]'$$

Finally, from the definition of the disturbance input vector $\boldsymbol{w}_2$ and the measurement output vector $\boldsymbol{y}$, it followed that

$$W_4 = \text{diag}(c_i), \qquad i = 12, \ldots, 25$$

where the first three elements of W_4 corresponded to the process noise input and were used for loop transfer recovery procedure similar to linear quadratic Gaussian/loop transfer recovery (LQG/LTR),[14,15] whereas the remaining elements corresponded to the sensor noise inputs and were adjusted to obtain desired sensor loop bandwidth (see Secs. III.B.2 and III.B.3).

Next we present a procedure for adjusting the weights W_1, W_2, W_3, and W_4 to achieve the design specifications.

B. Design Procedure

The design process used to obtain a feedback controller C is now summarized, followed by a detailed discussion of how each of the design steps were applied to the F-14 design example.

1) Set W_3 to zero. Use $\mathcal{H}_\infty$ state-feedback synthesis to determine W_1 and W_2 to satisfy the command and control-loop bandwidth requirements.

2) If closed-loop damping was unsatisfactory, refine the state-feedback design by adjusting W_3 weights to include derivatives of the lightly damped states in the output z. These states were identified by examining the eigenvectors associated with lightly damped eigenvalues of $A + B_2K_{sfb}$, where A and B_2 were the state and control input matrices of $\mathcal{G}$ and K_{sfb} was the state-feedback gain determined in step 1. The maximum element of such eigenvectors corresponded to the state contributing most to the lightly damped mode. Increasing the weight in the corresponding W_3 entry had the effect of damping the dynamic activity of that state. Readjust W_1 and W_2 weights to maintain the previously achieved bandwidth specifications.

3) Given $W_{1,2,3}$ determined as described, use measurement feedback design to determine the sensor noise weights in W_4 necessary to satisfy sensor response bandwidths.

4) Determine the process noise weights in W_4 by analyzing the broken-loop system responses, adjusting the weights as necessary to match the crossover frequencies that were observed for the state feedback design. This step is similar to LTR technique developed for linear quadratic methods.[14,15]

5) Readjust $W_{1,2,3}$ as required to maintain previously achieved specifications.

6) Evaluate resultant controller using linear and nonlinear simulation. Readjust weights if necessary.

7) Confirm compliance with other specifications: robustness, damping, and cancellation of lightly damped open-loop poles.

1. State Feedback Design: Determining the W_1 and W_2 Weights

The objective of this step was the determination of the appropriate W_1 and W_2 weighting functions to achieve the specified command and controller bandwidths. As a general rule, increasing the weight on the integral errors in W_1 increased the bandwidth of the respective command loop, and increasing the weight on the controller commands in W_2 decreased the crossover frequency in the corresponding control loop. Although adjusting the weight on a given term in W_1 or W_2 consistently had the desired effect on the corresponding command or control loop bandwidth, the multivariable nature of the problem meant that the changing a given weight occasionally resulted in wild variations in the bandwidths of the other loops, invariably in the most undesirable direction. Generally, once a desired crossover or corner frequency was attained, weights were then readjusted in response to adverse coupling effects. Table 1 depicts both the nominal performance and the refined performance for several iterations as the weights were varied. Here the numerator of each entry represents the value of the weight c_i applied to the identified term in z, whereas the denominator provides the resulting bandwidth in the corresponding loop. The column containing damping ratios represents the minimum damping ratio for all complex closed-loop poles at each iteration.

As can be seen from Table 1, by the third iteration, the bandwidths had been balanced close enough to the desired specifications to attempt improvement of closed-loop damping. The objective here was to identify the principal states contributing to the under-damped modes and to increase the weight on their respective rates in the output vector z such that activity in that mode was penalized. After the third iteration, the closed-loop system matrix $(A + B_2K_{sfb})$ had two complex pairs of eigenvalues with damping ratios of 0.55 and 0.43. The eigenvectors corresponding to these under-damped modes pointed at the h_e/s and q states, respectively. This indicated that h_e/s and q were the dominant participants in each of the two under-damped modes. Since h_e/s was a state internal to the controller, it was disregarded and attention was focused on enhancing the damping on q. To improve the damping of this mode, the weighted output $\dot{q}$ was included in z, and the value of the corresponding term in W_3 was set to a nonzero value.

Initially a very small weight was introduced relative to the other weights. Each time a damping weight was adjusted, the W_1 and W_2 weights were readjusted in response, to maintain the bandwidth characteristics previously achieved prior to further refinement of the damping weights. The eigenanalysis was performed after each adjustment of the damping weights to identify the most active state of each under-damped mode and to ensure that another mode had not become under damped. In fact, penalizing q made α and θ the dominant participating states in the same under-damped mode, requiring inclusion of a penalty on $\dot{\alpha}$ and $\dot{\theta}$ in W_3. A total of 24 iterations were required to achieve the design that approached both bandwidth and damping specifications. At the conclusion of this phase the weighting functions W_1, W_2, and W_3 had the following values:

$$W_1 = \begin{pmatrix} 10/s^2 & 0 & 0 \\ 0 & 30/s & 0 \\ 0 & 0 & 5/s \end{pmatrix}, \qquad W_2 = \begin{pmatrix} 5 & 0 & 0 \\ 0 & 0.01 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$W_3 = \text{diag}(0, 10, 5, 1, 0)$$

Figure 2 includes the broken-loop responses for the elevator, thrust, and DLC control loops for the given values of W_1, W_2, and W_3. Similarly, Fig. 3 includes Bode magnitude plots for h/h_{cmd}, α/α_{cmd}, and $\delta_{dlc}/\delta_{dlc_{cmd}}$ command loops.

2. Output Feedback Controller Design: Selecting the Measurement Noise Weights

The next step was to determine the terms in the weighting function W_4 that correspond to the measurement noise signals. The objective of this phase was to tune the controller to filter each sensor input as not to exceed the expected reliability of the sensor. For example, if

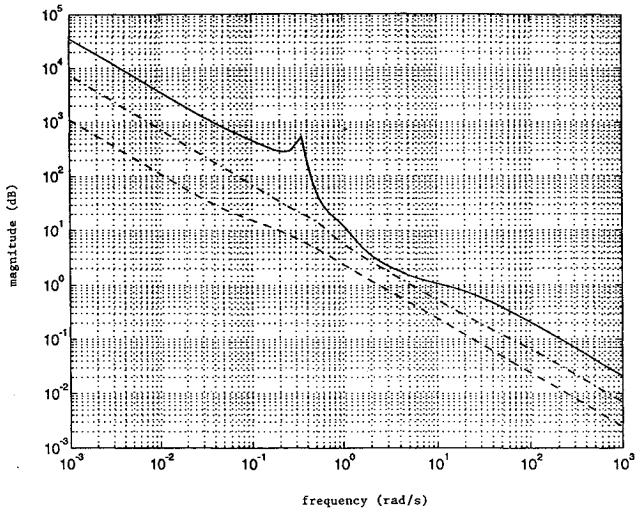

Fig. 2 Bode magnitude plots of the broken-loop control-loop responses.

Table 1 State-feedback design: weights/resulting bandwidths

	W_1			W_2			W_3					
Iteration: spec.	h_e >1 rad/s	α_e >1 rad/s	DLC >1 rad/s	$\delta_{stab_{cmd}}$ <20 rad/s	$\delta_{thrust_{cmd}}$ <2 rad/s	$\delta_{dlc_{cmd}}$ <40 rad/s	$\dot{u}$ ——	$\dot{\alpha}$ ——	$\dot{q}$ ——	$\dot{\theta}$ ——	$\dot{h}$ ——	ζ >0.6
1	1/1.5	1/0.01	1/1	1/2.5	1/0.01	1/1	0	0	0	0	0	0.43
2	1/0.8	10/2	1/0.01	1/4.5	1/0.01	1/1	0	0	0	0	0	0.43
3	1/0.8	5/1	1/1	1/4	1/0.01	1/1	0	0	0	0	0	0.45
5	1/0.9	5/1	1/1	1/100	1/0.01	1/1	0	0	5	0	0	0.53
7	1/0.9	5/1	1/1	1/100	1/0.01	1/1	0	5	5	0	0	0.62
8	1/0.8	5/1	1/1	5/8	1/0.01	1/1	0	5	5	0	0	0.44
24	10/1	30/3	5/20	5/10	0.01/2	1/6	0	10	5	1	0	0.60

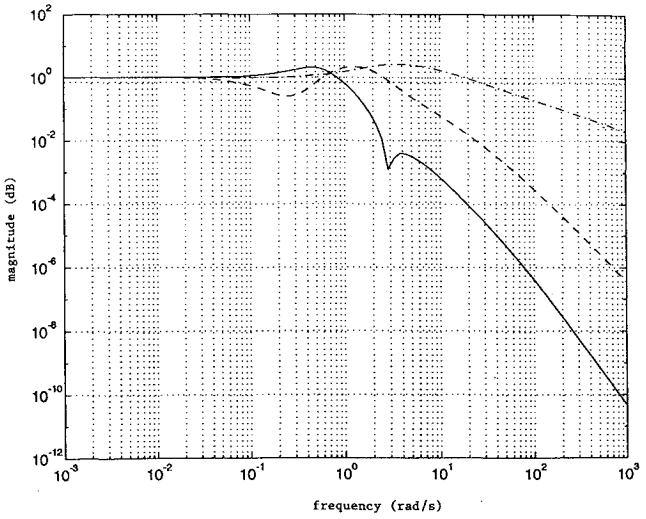

Fig. 3 Bode magnitude plots of the command-loop responses.

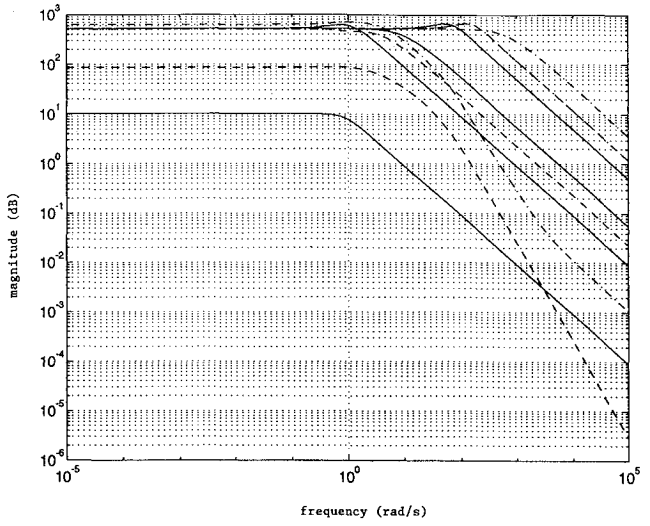

Fig. 4 Bode magnitude plots of the sensor loop responses.

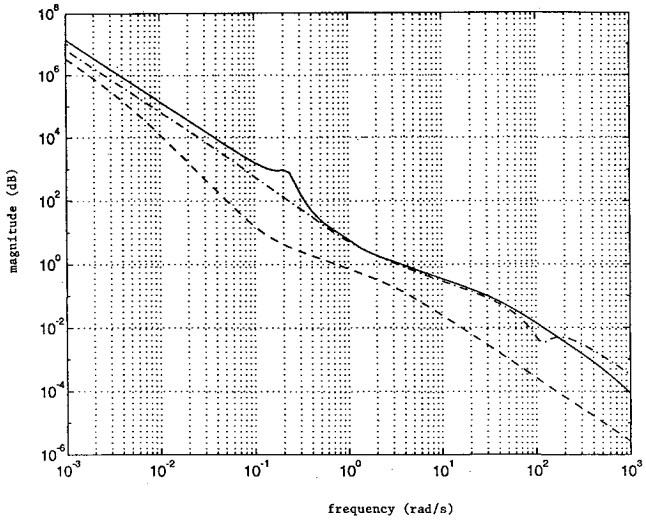

Fig. 5 Output feedback broken-loop controller responses.

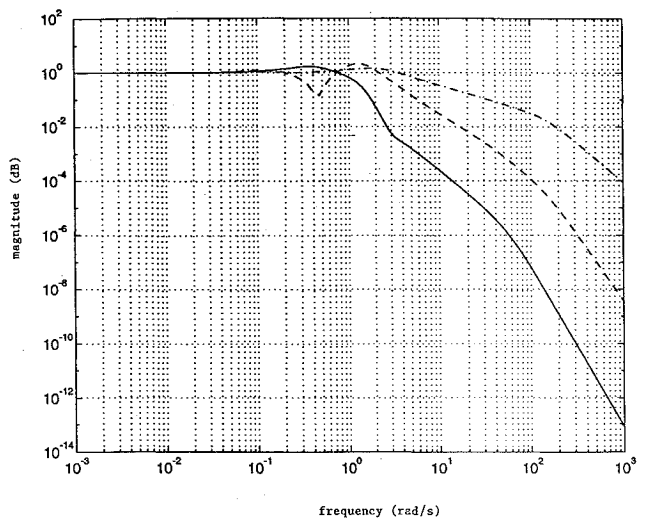

Fig. 6 Output feedback closed-loop command responses.

a sensor could be regarded as reliable at frequencies up to 10 rad/s, then the controller response to that sensor channel should roll off at a frequency less than or equal to 10 rad/s.

Consider the following representation of the controller:

$$\mathcal{C} = \begin{cases} \dot{\zeta} = A_c\zeta + B_c y \\ u = C_c\zeta \end{cases}$$

The frequency responses of the diagonal terms of the transfer function matrix $C_2(sI - A_c)^{-1}B_c$ were plotted to evaluate the influence of each channel on the controller (Fig. 4). The outputs of this transfer matrix represented the estimates of the states of $\mathcal{G}$ in the presence of the worst-case disturbance.[1] By examination of the corner frequency for each of the 10 channels of $C_2(sI - A_c)^{-1}B_c$, the weights could be adjusted to achieve the desired sensor response bandwidth. Increasing the weight on a given term resulted in the bandwidth for that channel being decreased. This is similar to the results encountered in LQG design[15] and is consistent with intuition. Very little effort was required to find the appropriate set of weights since little coupling was observed between the sensor channels.

3. Output Feedback Controller Design: Selecting the Process Noise Weights

The next step in the design process was the determination of the terms in the weighting function W_4 corresponding to the process noise inputs in the vector $\boldsymbol{w}_2$. This step was similar in execution to the loop recovery part of the LQG/LTR design procedure, with a slightly different purpose. Whereas the objective of LQG/LTR is principally the recovery of state-feedback robustness properties, the objective here was the recovery of the performance characteristics of the state-feedback controller, as reflected by the various bandwidth and damping specifications. Consequently, these adjustments were made to recover the control- and command-loop responses of the state-feedback controller obtained in step 2 of the design procedure. Figures 5 and 6 contain the Bode magnitude plots of the control and command loops for the output feedback controller. Comparison with the Figs. 2 and 3 points to a successful recovery of the state-feedback controller performance. As a result the final value of W_4 was

$$W_4 = 0.00001 * \text{diag}(0.1, 4, 4, 0.01, 0.01, 0.01$$

$$0.1, 10, 10, 1, 5, 5, 5)$$

4. Closed-Loop Analysis: Assessing the Controller Structure

At this point the closed-loop linear system was analyzed to determine whether reasonable actuator deflections were being used and to ensure that an altitude ramp could be tracked while controlling α to a desired trim value and DLC to zero. The closed-loop system was initialized to level flight and then expected to intercept and track an altitude ramp. Although the altitude ramp was successfully intercepted and tracked, both α and δ_{dlc} stabilized at values other than their respective set points. Examination of the transfer functions from altitude command to α and δ_{dlc} revealed only a single zero at the origin within each numerator. In both cases there was an additional zero numerically close to the origin but insufficient to provide the desired washout characteristics. To achieve these characteristics, an additional integrator was added to α and δ_{dlc} terms in W_1. A new controller was then obtained using this modified synthesis model with

the weights already determined. This controller was then evaluated using the identical simulation, with the result that both α and δ_{dlc} stabilized at their desired values. Finally, slight changes in W_1 and W_2 resulted in the final determination of control-and command-loop bandwidths. Therefore, the final values of the weighting matrices were

$$W_1 = \begin{pmatrix} 10/s^2 & 0 & 0 \\ 0 & 20/s^2 & 0 \\ 0 & 0 & 6/s^2 \end{pmatrix}, \qquad W_2 = \begin{pmatrix} 5 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$W_3 = \text{diag}(0, 10, 5, 1, 0)$$

Figures 5 and 6 depict the resulting broken-loop controller responses and the closed-loop command responses. The Nyquist plots of the broken-loop responses for each control input are shown in Fig. 7.

5. Specification Compliance

Figures 5 and 6 show all of the control and command loops as having satisfied their respective specifications, with the exception of the α command response, which was slightly low at approximately 0.3 rad/s. This was because of the presence of a complex zero at 0.45 rad/s in the α command loop. No objectionable properties resulted since this channel was used to regulate a constant command rather than respond to changes in the command signal. Figure 7 clearly depicts the simultaneous phase and gain margin requirements as having been satisfied, with all three controller traces staying in the right-half plane for all frequencies. Eigen decomposition of the closed-loop system revealed a complex pair of eigenvalues

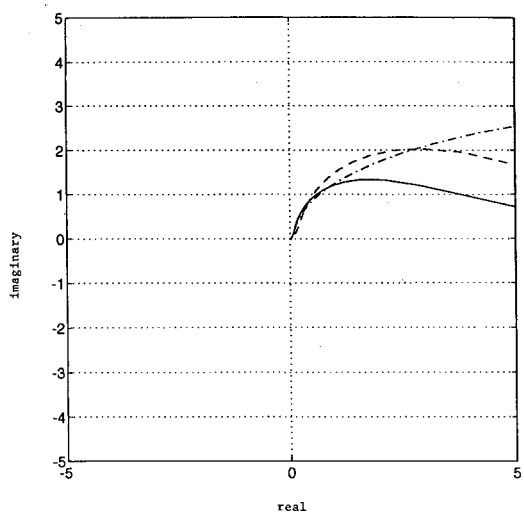

Fig. 7 Broken-loop Nyquist plot of the output-feedback system.

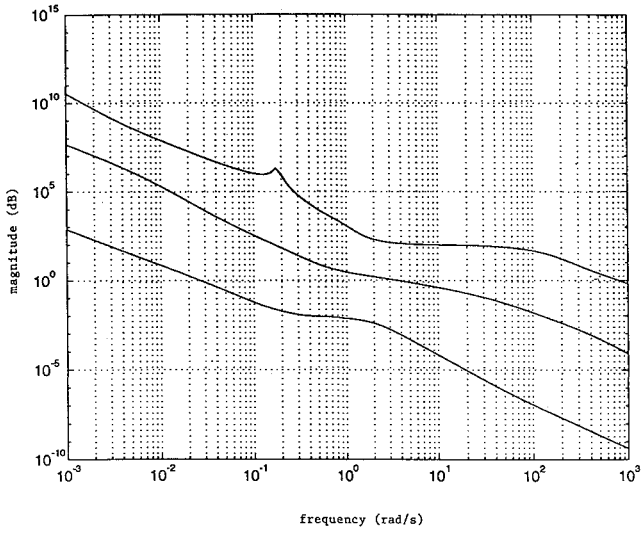

Fig. 8 Singular values of the output-feedback controller.

that failed the damping ratio requirement (0.45). Examination of the corresponding eigenvectors, however, revealed that no vehicle states and only controller states were participating in this mode.

The next step was to ensure that the cancellation of the plant's lightly damped modes by the controller had not occurred. The controller is a 3×10 matrix of transfer functions, and although cancellation may occur in one or several channels, complete cancellation would only occur if all of the channels had a common numerator term. See Ref. 16 for a discussion of multivariable transmission zeros. Figure 8 depicts Bode plots of the singular values of the open-loop controller. The presence of a clear spike at the frequency of the lightly damped phugoid mode and a small bump at the frequency of moderately damped short period mode indicated that the controller had not canceled these open-loop poles. Furthermore, numerical analysis revealed no transmission zeros in the controller, confirming the absence of pole-zero cancellations.

IV. Robustness Analysis

A. Uncertainty Modeling

This section describes the methodology used for analyzing robustness of the closed-loop system consisting of the aircraft model and the $\mathcal{H}_\infty$ controller. The aircraft model used in this paper was obtained from the flight-test data. As a result, some of the terms in the model were poorly known. These terms included the airplane's lift, drag, and pitching moment. On the other hand, terms such as gravity and aircraft dynamic coupling were well known. Furthermore, lift and drag were measured in the so-called stability axis, whereas the aircraft model was derived in the body-fixed coordinate system. These considerations indicated that the uncertainties in lift, drag, and pitching moment must be modeled exactly where they occurred. Therefore, this needed to be done using the aircraft's nonlinear equations of motion while taking into account the coordinate systems where these forces and moments were measured during the flight test. A detailed discussion of this process is given next.

Let $\boldsymbol{X}$ represent the vector of the longitudinal states of the aircraft:

$$\boldsymbol{X} = [U, \alpha, Q, \Theta, H]$$

Let δ represent the vector of the aerodynamic control effectors, consisting of the stabilators and DLC, and let $\boldsymbol{F}$ represent the vector of body-axis forces and moments, where

$\boldsymbol{F}_{\text{grav}}(\boldsymbol{X})$:= influence of gravity
$\boldsymbol{F}_{\text{dyn}}(\boldsymbol{X})$:= influence of dynamic coupling
$\boldsymbol{F}_{\text{aero}}(\boldsymbol{X})$:= influence of aerodynamic forces on the body
$\boldsymbol{F}_{\text{thrust}}$:= influence of thrust
$\boldsymbol{F}_{\delta}$:= influence of aerodynamic forces on the control surfaces

The aircraft equations of motion can now be expressed as

$$\begin{bmatrix} \dot{U} \\ \dot{\alpha} \\ \dot{Q} \end{bmatrix} = F_{\text{grav}}(x) + F_{\text{dyn}}(x) + F_{\text{thrust}}$$

$$+ R_{\text{wb}}(x)[F_{\text{aero}}(x) + F(\delta)] \tag{1}$$

where $R_{\text{wb}}(x)$ is the wind to body axis rotation matrix

$$R_{\text{wb}} = \begin{bmatrix} -\cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & -\cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2}$$

We only consider uncertainties in the aerodynamic forces and moments, since the gravity, dynamic coupling, and thrust are all well known.

The aerodynamic model of the airplane is derived using stability derivatives that represent the contribution of each state and control input to the aerodynamic forces and moments acting on the airplane. This data can be obtained either in the wind tunnel, where isolation of individual contributions is possible, or flight test, where only macroscopic behavior is observed, and then numerically distributed among the aircraft's states and inputs. The stability derivatives obtained from the flight-test data therefore depend on very complex multivariable parameter indentification (PID) methods, which are executed in two steps. First, forces and moments are computed from

observed accelerations and rates. Second, the PID process tries to identify which control inputs and airplane states contributed to the observed change in forces and moments. This step clearly introduces errors not present in the first computation.

Our choice was, therefore, to consider either the contribution of the uncertainty in each parameter or the net uncertainty in our knowledge of the forces and moments. There are several reasons to choose the latter. First, the size of the uncertainty model is significantly reduced. Second, the uncertainties in each stability derivative were not independent. If they were, then large uncertainties in each stability derivative would have resulted in large uncertainties in the net forces or moments, leading to an unnecessarily conservative design. A similar approach was used in Ref. 17 in synthesizing an $\mathcal{H}_\infty$ controller for a missile autopilot. Consequently, we chose to model the uncertainty block Δ as a 3×3 diagonal matrix, where each diagonal element represented a percentage of the nominal perturbation in the aerodynamic forces and moments (drag, lift, and pitching moment). Incorporating this block in Eq. (1) yields

$$\begin{bmatrix} \dot{U} \\ \dot{\alpha} \\ \dot{Q} \end{bmatrix} = F_{\text{grav}}(x) + F_{\text{dyn}}(x) + F_{\text{thrust}} + R_{\text{wb}}(x)(I + \Delta)[F_{\text{aero}}(x) + F(\delta)] \tag{3}$$

Figure 9 depicts Eq. (3). Here signals w_δ and z_δ denote the uncertainty inputs and outputs.

B. μ Analysis

The robustness analysis was performed by determining the structured singular value of a linearization of the closed-loop system consisting of the nonlinear plant model and the $\mathcal{H}_\infty$ controller. The analysis was performed about the nonlinear closed-loop system to achieve the best possible fidelity. The linearization was performed about the trimmed operating condition, with the uncertainty inputs and outputs w_δ and z_δ included in the nonlinear equations of motion as shown in Fig. 9. Figure 10 depicts the resulting structured singular value, where the peak value of 0.6 indicates compliance with the stability robustness condition $\mu < 1$ and the robustness design specification. Note that had the design effort failed to yield the desired robustness, the design process could have been repeated with the signals w_δ and z_δ incorporated in the synthesis model, and using either $\mathcal{H}_\infty$ synthesis or D-K iteration.[12]

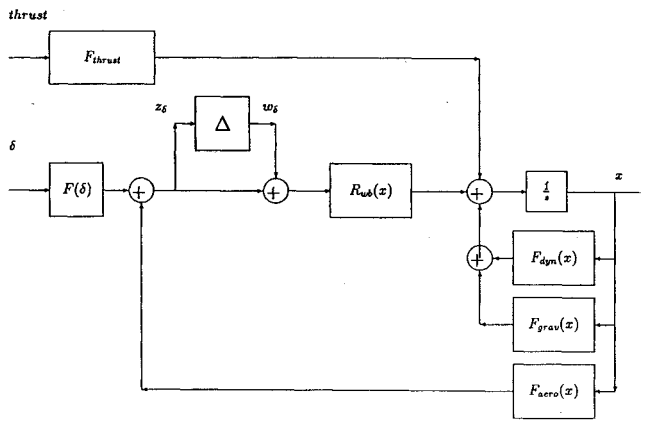

Fig. 9 Uncertainty model.

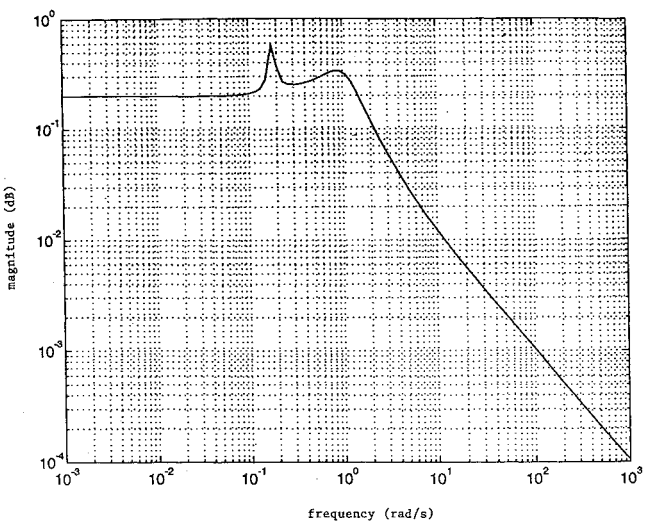

Fig. 10 Structured singular value plot.

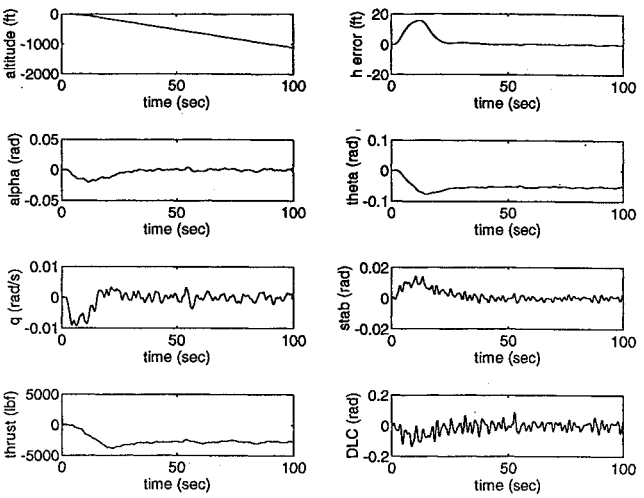

Fig. 11 Nonlinear simulation results.

V. Nonlinear Simulation

In this section we present the results of the nonlinear simulation of the $\mathcal{H}_\infty$ controller. First, the controller was implemented on the nonlinear plant using the $\mathcal{D}$-implementation methodology (see Ref. 18). This methodology is based on the observation that linear controllers are designed to act on the perturbations about the plant's nominal trajectory. Next, as with the linear simulation performed in Sec. III.B.4, the controller task was to intercept and track an altitude ramp, while appropriately controlling the other signals of interest. A vertical and horizontal gust field was simulated by passing white noise of moderate intensity (rms = 10 fps) through a zero-order hold. Results of the simulation are depicted in Fig. 11, where all of the variables are shown as deviations from the trimmed level flight condition. The simulation was initialized in steady level flight with all surfaces at their trim positions. In response to the altitude ramp, the DLC and stabilizer deflected immediately to establish the appropriate descent rate. The result is a decrease in both α and Θ. Both the average angle of attack and the average DLC wash out to their desired trim values as the altitude error is nulled and the thrust decreases to stabilize at the new steady-state condition. Deflections and deviations are all well within reasonable practical values. This simulation validated the results of the design effort.

VI. Conclusions

An output feedback controller was successfully designed to provide longitudinal control of an F-14 aircraft during automatic landing and was implemented on a nonlinear simulation. A key feature of the design was the exploitation of the aircraft's DLC to allow direct control of glideslope and altitude. Additionally, a methodology was detailed whereby SISO performance requirements were achieved using $\mathcal{H}_\infty$ synthesis. Next, a method was presented for the macroscopic consideration of parametric uncertainties arising from flight-test data to assess the robustness of the resulting closed-loop system. Finally, the resulting controller was validated on the nonlinear simulation.

Acknowledgments

This work was supported in part by the Research Initiation Grant at the Naval Postgraduate School and the Naval Air Systems Command under AIRTASK A546546TD.

References

[1]Doyle, J. C., Glover, K., Khargonekar, P. P., and Francis, B. A., "State-Space Solutions to Standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ and Control Problems," *IEEE Transactions on Automatic Control*, Vol. 34, 1989, pp. 831–847.

[2]Glover, K., and Doyle, J. C., "State-Space Formulas for All Stabilizing Controllers That Satisfy an $\mathcal{H}_\infty$ Norm Bound and Relations to Risk Sensitivity," *Systems and Control Letters*, Vol. 11, 1988, pp. 167–172.

[3]Khargonekar, P. P., Petersen, I. R., and Rotea, M. A., "$\mathcal{H}_\infty$ Optimal Control with State Feedback," *IEEE Transactions on Automatic Control*, Vol. AC-33, 1988, pp. 768–788.

[4]Khargonekar, P. P., Petersen, I. R., and Zhou, K., "Robust Stabilization of Uncertain Linear Systems: Quadratis Stabilizability and $\mathcal{H}_\infty$ Control Theory," *IEEE Transactions on Automatic Control*, Vol. AC-35, 1990, pp. 356–361.

[5]Petersen, I. R., "Disturbance Attenuation and $\mathcal{H}_\infty$ Optimization: A Design Method Based on Riccati Equation," *IEEE Transactions on Automatic Control*, Vol. AC-32, 1987, pp. 427–429.

[6]Kaminer, I., Khargonekar, P. P., and Robel, G., "Design of Localizer Capture and Track Modes for a Lateral Autopilot Using $\mathcal{H}_\infty$ Synthesis," *IEEE Control Systems Magazine*, Vol. 10, 1990, pp. 13–21.

[7]Kaminer, I., and Khargonekar, P. P., "Design of Flare Control Law for Longitudinal Autopilot Using $\mathcal{H}_\infty$ Synthesis," *Proceedings of the 29th Conference on Decision and Control* (Honolulu, HI), 1990, pp. 2981–2986.

[8]Kaminer, I., Pascoal, A. M., Silvestre, C., and Khargonekar, P. P., "Design of a Control System for an Underwater Vehicle Using $\mathcal{H}_\infty$ Synthesis," *Proceedings of the Conference on Decision and Control* (Brighton, England, UK), 1991, pp. 2350–2355.

[9]Sivashankar, N., Kaminer, I., and Khargonekar, P. P., "Design of a Controller for the Turret System Using $\mathcal{H}_\infty$ Synthesis," *Proceedings of the American Control Conference* (San Francisco, CA), 1993, pp. 1612–1616.

[10]Sivashankar, N., Kaminer, I., and Kuechenmeister, D., "Design, Analysis and Hardware-in-the-Loop Simulation of a MIMO Controller for a VTOL Unmanned Aerial Vehicle Using $\mathcal{H}_\infty$ Synthesis," *Proceedings of the American Control Conference* (Baltimore, MD), 1994, pp. 2506–2510.

[11]Doyle, J. C., "Analysis of Feedback Systems with Structured Uncertainties," *Proceedings of the IEE*, Pt. D, Vol. 129, No. 6, 1982, pp. 242–250.

[12]Doyle, J. C., "Structured Uncertainty in Control System Design," *Proceedings of the Conference on Decision and Control* (Fort Lauderdale, FL), 1984, pp. 260–265.

[13]Roskam, J., "Airplane Flight Dynamics and Automatic Flight Controls," Roskam Aviation and Engineering Corp., Lawrence, KS, 1979.

[14]Doyle, J. C., and Stein, G., "Robustness with Observers," *IEEE Transactions on Automatic Control*, Vol. 24, 1979, pp. 607–611.

[15]Anderson, B. D. O., and Moore, J. B., *Optimal Control: Linear Quadratic Methods*, Prentice–Hall, Englewood Cliffs, NJ, 1990.

[16]Kailath, T., *Linear Systems*, Prentice–Hall, Englewood Cliffs, NJ, 1980.

[17]Reichert, R. T., "Application of $\mathcal{H}_\infty$ Optimal Control to Missile Autopilot Design," *Proceedings of the AIAA Conference on Guidance, Navigation, and Control* (Boston, MA), AIAA, Washington, DC, 1989, pp. 1065–1072.

[18]Kaminer, I., Pascoal, A. M., Khargonekar, P. P., and Coleman, E., "A Velocity Algorithm for the Implementation of Nonlinear Gain Scheduled Controllers," *Automatica*, Vol. 37, 1995, pp. 1185–1191.