Journal of Applied Intelligence 2, 299–320 (1992)

The University of Texas at Arlington Autonomous Aerial Vehicle—An Overview

JEFFREY O. SMITH
Computer Science Engineering Department, The University of Texas at Arlington, P.O. Box 19015, Arlington, TX 76019-0015

KLIFFTON M. BLACK
Mechanical Engineering Department, University of Texas at Arlington, Arlington, TX 76019-19023

FAHRAD A. KAMANGAR
Computer Science Engineering Department, University of Texas at Arlington, Arlington, TX 76019

JACK FITZER
Electrical Engineering Department, University of Texas at Arlington, Arlington, TX 76019

Received November 27, 1991, Revised February 24, 1992

Abstract. This paper reports on the status of The University of Texas at Arlington student effort to design, build and fly an Autonomous Aerial Vehicle. Both the 1991 entry into the First International Aerial Robotics Competition as well as refinements being made for 1992 are described. Significant technical highlights include a real-time vision system for target objective tracking, a real-time ultrasonic locator system for position sensing, a novel mechanism for gradually moving from human to computer control, and a hierarchical control structure implemented on a 32-bit microcontroller. Detailed discussion about the design of multivariable automatic controls for stability augmentation is included. Position and attitude control loops are optimized according to a combined $\mathcal{H}_2$ and $\mathcal{H}_\infty$ criteria. We present a modification of a recently published procedure for recovering a desired open-loop transfer function shape within the framework of the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ problem. This work has led to a new result that frees a design parameter related to imposing the $\mathcal{H}_\infty$-norm constraint. The additional freedom can be used to improve upon the performance and robustness characteristics of the system.

Key words: Autonomous vehicle, robust control, optimal control, inertial navigation, microprocessor control

1. Introduction

The Association for Unmanned Vehicle Systems has organized an International Aerial Robotics Competition open to student led teams from universities worldwide. The challenge is to design and build a completely autonomous aerial vehicle (AAV) to transfer six randomly placed metal spools or disks one at a time from a pick-up area, to a drop-off area in less than six minutes. The nominal distance to be traversed is 80 feet while clearing 3 foot high centrally located barrier. The 6 foot diameter pick-up and drop-off rings are defined by walls 3 inches in height.

In 1991, students from the University of Texas at Arlington (UTA), College of Engineering, entered the only vehicle that proved capable of autonomous flight at the first annual competition (Figure 1). The aircraft was designed, developed and constructed at a total cost of $8,500 (including the retail value of loaned or donated parts and supplies). Although UTA had no formal industrial partner, technical assistance, encouragement, and donations were provided by local in-

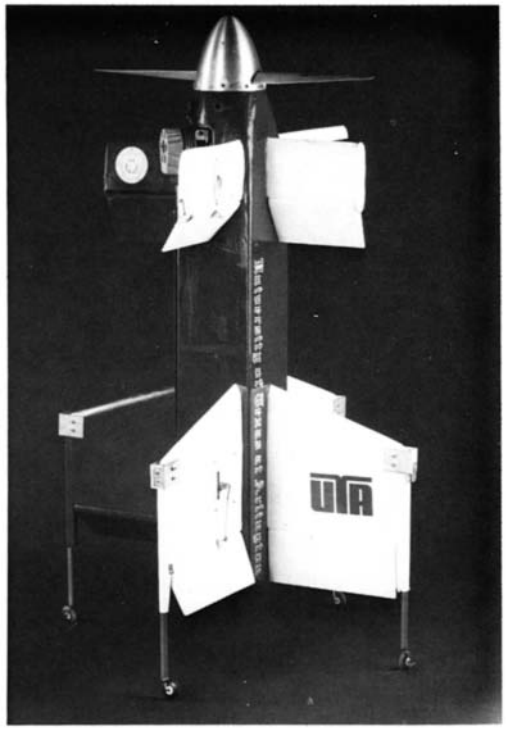

Fig. 1. The 1991 aircraft

dustry. At the competition, the vehicle lifted-off and flew to the center of the pick-up area, but a landing boom contacted the ring surrounding the pick-up area and the vehicle crashed.

Out of ten university teams that registered for the competition, four besides UTA were successful in completing enough of the task to attend the fly-off at the Georgia Tech campus in July 1991. These included Cal Poly, University of Dayton, Georgia Tech, and M.I.T.

This paper is organized as follows. Section 2 gives an overview of the early planning that led to the choice of vehicle configuration and specifications. The airframe and power-plant are discussed in Section 3. Sections 4 details the sensor package including the on-board.vision system, acoustic position sensors, Doppler rate of climb sensor, altimeter, accelerometers, and rate gyros. Section 5 discusses the servos and the disk pick-up mechanism. In Section 6, the hierarchical control system consisting of a high level state machine, navigation control, and stability augmentation is discussed. The stability augmentation design is covered in detail. Since this design is based upon some new results, a full theoretical development is included in a supporting Appendix A. The final section provides some conclusions and a description of on-going research.

2. Performance Objectives

This section presents some of the early planning and analysis that led to major design decisions. We begin with an analysis of the mission requirements and then develop a set of minimum performance requirements.

2.1. Scoring Analysis

The 1992 rules outline a scoring formula that includes static, performance, and subjective categories of judging. The elapsed time of flight, vehicle weight, and vehicle volume all serve to reduce the maximum possible score. All other items contribute positively to the score as indicated in Table 1. The last column is based upon successful transfer of a single disk. The percentages indicate the relative importance of each scoring category. The next to last column is based upon successful transfer of all six disks. It is clear that to acquire and partially transfer (to within 25 ft of the Drop-off ring) a single disk on a flight lasting 30 seconds or more far outweighs all other scoring categories combined in relative importance.

Deductions in score for vehicle weight and volume will be assessed at one point per pound and one point per cubic foot, respectively. Extrapolating from the 1991 vehicle, UTA can expect a total deduction of about 40 to 50 points for these two categories combined, so the potential for improvement in these areas is minimal.

A deduction for flight time will be assessed at a rate of $1/(1+f)$ points per second where f is the number of disks acquired and partially transferred. Thus, the probable total deduction for time of flight is difficult to predict, and scenarios can be formulated whereby it is better to land early than to continue to attempt to pick up disks. For example, a successful flight of 30 seconds with zero transfers would result in a time deduction of 30 points and a net score of 69 points for items $f+c+e$ less the time deduction.

Table 1. Scoring analysis.

Item	Points		Maximum	% for 6	% for 1
f. Disk acquisition w/partial transfer	200 each × 6	=	1200	70.6	44.4
c. Successful disk transfer	50 each × 6	=	300	17.6	11.1
e. Fly 30 seconds or more	99	=	99	5.8	22.0
g. Successful landing	30	=	30	1.8	6.7
j. Safety of design	0–40	=	40	2.4	8.9
h. Elegance of design/craftsmanship	0–10	=	10	0.6	2.2
i. Innovation in design	0–10	=	10	0.6	2.2
k. Innovation in disk acquisition	0–10	=	10	0.6	2.2
l. Best team tee shirt	0–1	=	1	0.0	0.2
	Maximum Possible		1700	100.0	

Continuing the same flight for 360 seconds with zero transfers would result in a time deduction of 360 points and a similar net score of minus 261 points. So, a 30 second hover "bad weather" flight plan should be prepared and implemented should it be deemed unlikely that a disk can be acquired.

As disks are transferred, the maximum potential deduction for a continued flight becomes less significant. Assume one disk is successfully acquired and partially transferred giving a 200 point addition to the score. Then the time deduction becomes 90 points for a 3 minute flight and 180 points for a 6 minute flight. But if 6 disks are partially transferred, then the plus score is 1200 points with time deductions of 26 points for a 3 minute flight or 51 points for a 6 minute flight. It appears most advantageous not to "rush" the flight, but to ensure successful acquisitions with at least partial transfers.

2.2. *Proposed Mission Profile*

A reasonable target for total disk acquisition and transfer time that is well under the 6 minute rule must be specified. If a specification of 4 minutes is selected (leaving 2 minutes for take-off, landing, and unforeseen circumstances), then each trip must be completed in an average of 40 seconds. If we allow 18 seconds on station over the pick-up ring to locate and acquire a disk then the round-trip traverse time must be held to 22 seconds.

To verify the feasibility of this flight profile, assume the vehicle is capable of a ramp change in velocity. Clearly, the implied ability to make step changes in acceleration is an over-simplification, but this assumption will suffice for initial approximations. To traverse 80 feet in 11 seconds dictates an average velocity of 7.27 ft/s. If the maximum horizontal cruise velocity is 10 ft/s and the maximum acceleration is 4 ft/s^2 (approximately ⅛ g), then the one-way traverse time works out to be 10.5 seconds. A side force of about 3 lb applied to a 24 lb vehicle would be required.

2.3. *Pitch-Over versus Constant-Attitude Flight*

Two ways of achieving the accelerations and airspeeds required to meet the proposed mission profile are considered:

Method (A)—Pitch-over of the entire vehicle so that the engine thrust is vectored as desired;

Method (B)—Constant-attitude with deflection of flying surfaces to achieve lift in the desired direction of travel.

These two methods have been selected for this analysis because both may readily be applied to the basic plan-form of the tail-sitter which was used by UTA in the 1991 competition. While other possible aircraft configurations involving alternative philosophies (e.g., helicopter, twin-rotor, ducted fan) were considered, it was decided that to draw upon the experience gained by the 1991 team would yield the highest possible potential for success.

Using Method (A), a pitch-over of about 7 de-

grees will give the minimum required horizontal acceleration of 4 ft/s^2. With Method (B), a wing area of 1.0 ft^2 at a maximum loading of 3 lb/ft^2 would be needed. If a maximum coefficient of lift of 0.5 is assumed, then 3 lb/ft^2 can be achieved with a nominal air velocity of V_∞ = 72 ft/s. The '91 aircraft had a nominal prop-wash velocity of 80 ft/s in steady hover. Thus, either method appears capable of generating the minimum required translational acceleration needed for the proposed mission profile.

The only obvious advantage of employing Method (A) pitch-over for translational acceleration is that the use of this method in 1991 affords a certain level of experience that can be drawn upon in designing the 1992 vehicle. A second possible advantage is that the overall size of stabilizing fins and control surfaces might be significantly smaller than those required for Method (B) constant-attitude flight. Disadvantages of Method (A) include higher levels of complexity and difficulty associated with altitude and navigational determinations, precision hold over a target disk, target disk retrieval, and cross-wind landing. There will also be slower response to wind gusts (since the entire vehicle must pitch and yaw to compensate) and significant gyroscopic coupling effects due to propeller and engine inertia. Stability augmentation and control of the vehicle using either method can conceivably be accomplished using 5 independent control surfaces and a throttle. However, it is anticipated that for Method (A) the control and navigational algorithms will be more complicated.

2.3.1. Engine Gyroscopic Effect. The potential adverse yaw due to propeller gyroscopic effect during a pitch-over maneuver is significant. Per the definitions in Figure 2, positive pitch rate results in a positive yaw acceleration given by

$$\dot{r}_P \simeq \frac{M_{zP}}{I_{zz}} = \frac{I_P \cdot \omega_P \cdot q}{I_{zz}} \qquad (2.1)$$

where

$\dot{r}_P$ = vehicle adverse yaw acceleration induced by propeller gyroscopic effect;
I_{zz} = vehicle moment of inertia about z-axis;
M_{zP} = yawing moment due to propeller gyroscopic effect;
I_P = propeller moment of inertia;
ω_P = propeller rotational speed;
q = vehicle pitch rate.

Assume the propeller is 22 inches long, turning at 9167 rpm (960 rad/s), and approximated by a thin rod weighing 0.72 ounces (0.045 lb). Then,

$$M_{zP} = \frac{1}{12}\left(\frac{0.045}{32.2}\right)(1.833)^2(960)q$$
$$= 0.376q \qquad \text{lb.-ft.}$$

Further assume a 3 ft long 24 lb vehicle approximated by a thin rod ($I_{zz} = 0.56 \frac{\text{lb-s}^2}{\text{ft}}$ ft^2). Then the total adverse yaw produced by a command to pitch at 0.1 rad/s for 1 second is given by

$$\Psi = \int_0^1 \int \frac{M_{zP}}{I_{zz}}\, dt\, dt = 0.034 \text{ radians};$$

or, 34% of the commanded pitch.

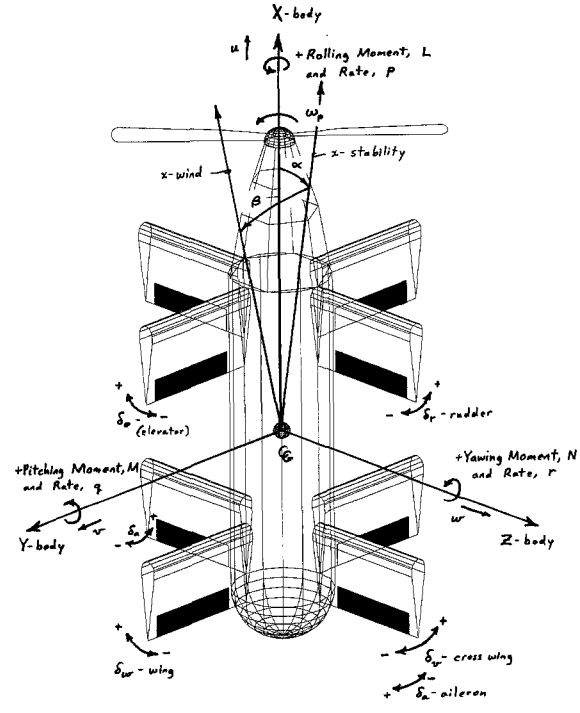

Fig. 2. Axis and control surface definitions

2.3.2. Gust Response. It is safe to assume that control surface responses to wind gust upsets will be similar with either Method (A) or (B). However, since the entire vehicle must rotate to compensate with Method (A), an additional lag will be seen in the response. The lag will be related to the maximum pitch and yaw accelerations that are available. The proposed standard actuators are believed to have a bandwidth of about 6 rad/s, and reasonable pitch and yaw bandwidths for the vehicle might be about 5 rad/s at best. So, Method (B) constant-attitude wind gust response can be expected to be about 4 times faster than Method (A).

2.3.3. Conclusion. Method (B) constant-attitude flight is preferred provided the specifications outlined in Section 2 can be met with reasonably sized flying surfaces and actuators. Further study will be needed to determine if this is possible. If a pure constant-attitude flight mode is not possible, then a mixed-mode should be considered with constant-attitude for low-speed translations and pitch-over for high-speed translations. While this approach will complicate the control system, it should generally simplify the most demanding task of holding position and attitude while locating and locking-on to a disk.

2.4. Guideline Specifications

It is likely that the competition will not enjoy dead calm wind conditions, so additional performance capability should be provided. Assume a worst case situation with a 15 ft/s (approximately 10 mph) steady wind blowing the length of the arena. Further, assume the vehicle is capable of ramping horizontal velocity at an average of 4 ft/s^2 to a maximum safe cruise airspeed of 20 ft/s using Method (B) constant-attitude mode of flight. This is a gross oversimplification, but the results are still valuable for first approximations of performance. The upwind traverse time will be about 14.8 seconds and the downwind time about 8.9 seconds. Thus, the round-trip time will be about 49 seconds including 25 seconds on station to acquire a disk in windy conditions. Even with the more realistic assumption the acceleration is reduced to, say, 1 ft/s^2 average in the 15 to 20 ft/s airspeed range, the upwind time will be about 21 seconds, resulting in a total round-trip time of about 56 seconds; that is, the entire mission can be completed in just under the 6 minute limit (Take-off and initial traverse time is assumed to at least be partially offset by not having to return to the pick-up area following the last transfer.)

As noted earlier, completion of the mission in 4 minutes versus 6 minutes offers very little advantage when a large number of disks are transferred. However, designing to such a specification affords an overall factor of safety of about 1.5 thereby providing some margin for lower than expected system performance and/or unexpected or difficult conditions (e.g., gusty winds).

Based upon the above analysis, the vehicle has been designed according to the following minimum guidelines:

• Maximum Safe Steady Wind	15 ft/s
• z-axis cruise airspeed	20 ft/s
• $\dot{w}_{max}$ (at w = 0 ft/s)	4 to 6 ft/s^2
• $\dot{w}_{min}$ (at w = 20 ft/s)	1 to 2 ft/s^2
• y-axis cruise airspeed	10 ft/s
• $\dot{v}_{max}$ (at v = 0 ft/s)	2 to 3 ft/s^2
• $\dot{v}_{min}$ (at v = 10 ft/s)	1 to 2 ft/s^2
• Disk acquisition time	8 seconds
• Minimum Excess Power	5% at max cruise airspeed

These specifications dictate that if the maximum allowable wind is cross course, then the capability of establishing a roll attitude such that the z-axis (or negative z-axis) is pointed upwind will be needed. The resulting complexity of mixed-mode control could be undesirable, and a more symmetric vehicle may be indicated.

3. Aircraft Design

The UTA AAV (Autonomous Aerial Vehicle) concept is derived from a prototype VTOL tail-sitter remotely piloted airplane developed by Sky-Technology. The original concept has the capability to transition into horizontal flight. The students redesigned the vehicle for the hover

mode only, and built the new design from scratch. The airfoil and fuselage have been redesigned to accommodate the competition flight profile. The vehicle is a symmetrical cruciform shape with regular control surfaces and canard surfaces immersed in the slipstream of the prop (Figures 1 and 2). Body axes are defined as a standard aircraft in vertical flight. The x-body axis is oriented out the vehicle nose, the y-body axis is out the right side, and the z-body axis follows standard right hand orientation.

3.1. Fuselage

The fuselage for the AAV is a streamlined aerodynamic body of semi-monocoque construction designed to provide stowage for the electronic components used in the mission. This is most efficiently done by using a simple square cross section with the nose of the aircraft transitioning to a circular cross section to conform to the spinner and provide a more aerodynamic shape. The base of the fuselage is tapered to facilitate the downward looking camera of the close range vision system. The major components to be mounted in the fuselage include the power plant, fuel tank, throttle servo, control computer, frame grabber, vision computer, rate gyros, radio receiver, CCD camera and electrical power system.

3.2. Wing Design

The UTA AAV wing design is a National Advisory Committee for Aeronautics (NACA) symmetrical airfoil section utilizing vertical end-plates at the tips. The spans of the wings and upper canard surfaces are driven by the estimated diameter of the propeller down-wash. The vertical end-plates serve to reduce losses due to low aspect ratios and therefore increase maneuverability and control response. Two of the lower wings are currently being tested with symmetric spoiler mechanisms to create excess drag, thereby providing a more responsive means to control altitude than available by adjusting the engine rpm. Control surface servos are installed in the wings to help free space in the fuselage for electronics and other equipment.

3.3. Power-plant

The power plant is an O.S. Max BGX-1 model aircraft engine. The manufacturer rates this two cycle engine at 4.1 maximum horsepower and a maximum speed of 10,000 rpm. The fuel tank capacity is 22 ounces and is estimated to give a run time of 9 minutes at full power. In order to determine the best propeller match for a given engine configuration, a test stand was designed and constructed. The test stand gives torque and thrust measurements at a given engine rpm through the use of electronic strain measurements. Strain gauges are mounted on two aluminum load cells which undergo pure bending during a test. Several different engine-propeller combinations were tested revealing that a maximum thrust of 28 pounds could be obtained with a $22'' \times 6$ pitch propeller. At this operating point the engine produces 3.0 ft-lb reaction torque at 6700 rpm which implies 3.82 BHP.

If this data is used to calculate the propeller efficiency, it is found to be about 73%. Efficiency could perhaps be improved with a propeller designed for zero speed forward flight, and this is currently being investigated. With a more efficient, smaller diameter propeller, the engine reaction torque could be redued for a given payload.

The engine reaction torque is, of course, a major issue in the planned vehicle because relatively short span flying surfaces must be used to cancel the torque. Also at issue is the prop-wash velocity over the flying surfaces and the propeller tip speed. Lower velocities imply larger surfaces for all control needs, while higher velocities imply increased horsepower demand. Tip speed must generally be held below Mach numbers of about 0.75 to 0.80 to avoid a radical drop-off in propeller efficiency and increased dynamic loading due to sonic shock.

4. Sensor Package

Inertial, absolute, and derived sensing schemes have been incorporated to obtain the location, orientation, velocities, and accelerations needed to control the vehicle; see Figure 3. The acoust-

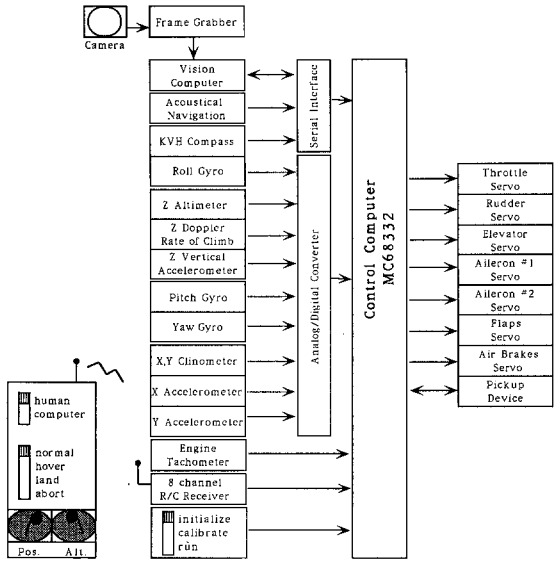

Fig. 3. Control computer block diagram

ical position system, and acoustical altimeter we have developed are based on a ranging module used in Polaroid cameras. The Doppler rate of climb sensor is based on transducers found in garage door openers. Accelerometers are integrated between position updates to provide accurate velocity and position data at a higher data rate than the sonar systems can provide. The scheme we employ is similar to that currently employed in aircraft and ships using a global positioning satellite system receiver. A similar scheme has been developed to derive heading from integrating a low-cost high drift rate gyro while updating with a flux gate compass which provides a heading sample every 100 ms. In the same way, a clinometer is used to update derived attitude from pitch and yaw rate gyros. An onboard vision system has also been developed to track disks and direct the craft once the inertial and position systems have placed the craft within the pickup area.

4.1. Position (X, Y)

The competition requirements demand that new sensing schemes be developed to position the craft within the tolerances required. Radio Frequency time-of-flight transmissions like LORAN-C have been in use for many years, but they are only accurate to approximately 300 feet. Global positioning satellite systems have differential capability to within 3 feet in some areas, however, the update rates are relatively slow. New high-power solid state laser diodes and high speed electronics make laser ranging and Doppler velocity measurements possible, but the cost, size, and weight are prohibitive. Acoustics are well suited for this application. Sound travels at approximately 1100 ft/s and therefore accuracies to 1 inch and ranges of 150 feet are possible using low cost components.

An acoustical LORAN-like triangulation method is used for obtaining the X, Y location of the craft within the arena. Four ultrasonic transmitters located on the border of the arena transmit an ultrasonic pulse every 125 ms. A synchronization pulse sent over a coaxial cable connected to each of the four transmitters ensures that they all transmit at the same time. Each of the transmitters transmits at a different ultrasonic frequency between 40 and 80 kHz. An array of ultrasonic receivers on-board the craft receives the transmitted pulses. Four phased locked loops with center frequencies matching each of the four transmitted frequencies output a signal when these frequencies are received. A Motorola 68HC11 receives these signals and based on the relative arrival time of the pulses determines the craft position to within one foot in both the X and Y directions.

4.2. Acceleration (X, Y)

The acoustical navigation system provides position information eight times per second. The information provided by the location system is delayed 30 ms due to computation time of the triangulation algorithm. Therefore, velocities are derived from an integration of accelerometers in the X and Y axes (normally aligned with the aircraft y-, z-body axis system). A second integration is performed to provide position information in between the navigation samples.

The vehicle vibration due to the two cycle engine rotating at approximately 7200 rpm introduces a significant noise component at 120 Hz. This is eliminated using a combination of me-

chanical damping (foam mounts) and an analog low-pass filter. Another benefit of the low pass filter is the elimination of aliasing frequencies due to sampling rates close to that of the engine noise.

4.3. Altitude

Altitude is sensed through two Polaroid polapulse 50 kHz ultrasonic transducers. The control computer triggers these sensors to transmit a sound burst and then times the reflected pulse. Sound travels at approximately 1 ft/ms, therefore a 10 ms pulse corresponds to an altitude of approximately 5 feet. These sensors have an absolute accuracy of 0.1 foot. The maximum altitude that we intend to attain during flight testing and competition is 14 feet. Therefore we sample every 28 ms. This provides an altitude update rate of 30 Hz. A digital low pass filter provides filtering of high frequency noise spikes which occur due to missed echoes. Two sensors, forward and rear, are used to compensate for incorrect measurements which occur when crossing the central barrier.

4.4. Rate of Climb

In the 1991 craft, rate of climb was determined by differentiating altimeter data. This proved unsatisfactory due to the low sample rates of the acoustical altimeter system and the amplification of noise inherent in digital differentiation. An alternative method has been developed which uses the property of Doppler frequency shift. Doppler shift occurs when a waveform of a particular frequency is reflected off of a body moving at some velocity. The shift in frequency is proportional to the velocity of the object and inversely proportional to the wavelength of the transmitted frequency. An on-board ultrasonic transmitter commonly found in garage door openers sends a continuous 40 kHz tone. An on-board receiver turned for 40 kHz+/− 5 kHz receives the signal reflected from the ground. The changes in vertical velocity of the craft result a frequency modulation of the 40 kHz carrier signal. A phased locked loop is commonly used to demodulate frequency modulated signals. The output of a phase locked loop is proportional to the difference between a received frequency and a set frequency. In our case the phase locked loop is tuned to output 0 to 5v corresponding to vertical velocities of ± 1 ft/s. This corresponds to a Doppler shift of about 80 Hz.

4.5. Vertical Acceleration

The 1991 vehicle uses an accelerometer to sense vertical acceleration. An analog low pass filter is used to filter engine noise and aliasing frequencies. A skewed gravitational vector due to pitch, and yaw is another source of noise in the accelerometer data. However, since our approach is to maintain constant attitude (in so far as possible) we have assumed the gravity component to be constant.

4.6. Pitch and Yaw

A clinometer is used to measure the attitude (pitch and yaw) of the craft. A clinometer is like a carpenter's level with an analog output proportional to the pitch and yaw angle. One problem with the clinometer is the inaccuracies introduced from accelerations in the y, z plane. By measuring the accelerations in the y, z plane we can compensate for these inaccuracies. We are also investigating measuring the attitude of the vehicle sonar sensors similar to the sonar altimeter. By placing these sensors at the wing-tips, the differential can be used as an additional source of pitch and yaw information.

4.7. Pitch and Yaw Rates

Futaba rate gyros (FP-G154) are used to measure the pitch and yaw rates. These inexpensive rate gyros are sold primarily to hobbyists for radio-controlled applications. It would be useful to obtain attitude from the integration of the gyro output. An analysis of these gyros has shown significant hysteresis (4 deg/sec) and drift (3.4 deg/sec). Although this means that an inertial navigation system could not be based on these sensors alone, by incorporating them with other sensors, a reliable inertial navigation scheme can be realized.

4.8. Heading

A KVH flux-gate compass is used to obtain true heading. The compass provides 0.1 deg resolution every 100 ms. This axis is unique because of the "modulo" effect from the compass: the heading wraps around back to zero after 359.9 degrees. This complicates the error determination because heading can be obtained from two directions, counterclockwise and clockwise.

4.9. Roll Rate

A Futaba rate gyro is used for roll rate information. The characteristics of this gyro and implications of deriving heading from it are the same as those discussed in pitch and yaw rate section.

4.10. Vision

The vision system consists of an on-board CCD camera, frame buffer, and TMS320-C25 DSP. The frame buffer and DSP boards interface to the control computer through a serial connection. When the vehicle encounters the pick-up arena the vision system is turned on. When a disk is located, the vision system provides a position of the disk currently being tracked. The navigation system uses this information to direct the vehicle to hover directly above it.

Several ATD (Automated Target Detection) and tracking algorithms were evaluated [1,2]. Many conventional algorithms (i.e., edge detection) proved ineffective where glare or more than one disk were visible, trackers like AGC (Adaptive Gate Contrast) and Centroid have problems in the presence of ambient light [3]. The algorithm that we have developed applies a sum of absolute differences operator on every possible target location to the image. Once a match is found, a target state estimator calculates the relative velocity and position of the disk over the current and previous frame to predict the disk's next relative position. A search area is defined using this predicted position. This area is used as the starting point on the next frame. This reduces the number of calculations needed on each subsequent track. Typically, once a disk has been tracked for several frames, the track rate increases to thirty position updates per second with an accuracy of one quarter of an inch at an altitude of 4 feet.

5. Actuators

5.1. Servos

Futaba R/C servos are used to deflect aerodynamic surfaces, control the throttle, and lower and raise the disk pick-up mechanism. These servos require pulse width modulated signals with a period of 20 ms. We are currently investigating operating these servos at an update period of 25 ms. This will enable us to synchronize the output to the servo with the primary control loop sample rate of 40 Hz.

5.2. Pick-Up Mechanism

Five disk pickup mechanisms have been prototyped. Four of these were developed in a design class by engineering students who were competing for their design to be incorporated into the 1992 craft. Each of these utilizes a different concept for acquiring the disk. One prototype is designed to be lowered from the craft, extend a semi-flexible retrieval arm, and sweep a radius of up to two feet. The arm is configured so that only one disk can be retained. The primary advantage of this device is that accurate positioning is not required. A second prototype lowers an array of electromagnets and acquires all disks in a one foot diameter area. As the array is retrieved, magnets are switched off until only one disk remains. The third prototype extends a telescoping arm with an electromagnet attached at the end. The fourth prototype launches a permanent rare-earth magnet cluster at the disk. The cluster is attached to a bag partially filled with sand in order to dampen the impact and keep it from bouncing off. Backlash is avoided by spooling a lightweight fishing line from a mechanism similar to that of a fishing spinner reel. This also affords a rapid retrieval rate. This prototype is most similar to the 1991 design which uses an electric hoist to lower a permanent magnet to pick up the disk. An array of Hall-effect sensors detects

when a disk is acquired. At the drop-off area the hoist pulls the disk up against a stop to break the disk loose.

6. Control System

The vehicle control system consists of three hierarchical control systems each interfacing to the next through variables representing the current state of the craft. Supervisory control is performed in the outermost level. This system has been implemented on the control computer as a state machine. The next level is the navigation control system. This system is responsible for providing set-points to the stability augmentation system thereby either repositioning the craft or maintaining a current global coordinate position. Craft stabilization occurs in the innermost control level. All three of these functions are performed on a 16.7 MHz Motorola 68332 32-bit microprocessor and 68882 floating point coprocessor. As discussed in earlier sections, these control systems use sensor data from accelerometers, rate gyros, clinometer, acoustical altimeter, Doppler acoustic rate-of-climb, flux-rate compass, acoustical navigation system, disk detector, and on-board vision system.

6.1. Supervisory Control

The AUVS competition is well defined and occurs in a structured environment. Based on the scoring analysis in Section 2.1, the flight plan can be decomposed into a small number of tasks that the vehicle must perform. Various internal and external events will trigger the transition between tasks (see Figure 4). Notice that although these commands may seem to occur in a simple sequence it is the responsibility of the supervisory control computer to determine if a disk has been dropped, fuel is running low, time is running out, etc., and to make appropriate changes in the flight plan. For example, if a disk is dropped while transporting it to the drop-off area, the vehicle should stop moving toward the drop-off area and retrieve another disk.

Predefined strategies have been developed and programmed into the supervisory control system. The strategy for a particular flight will

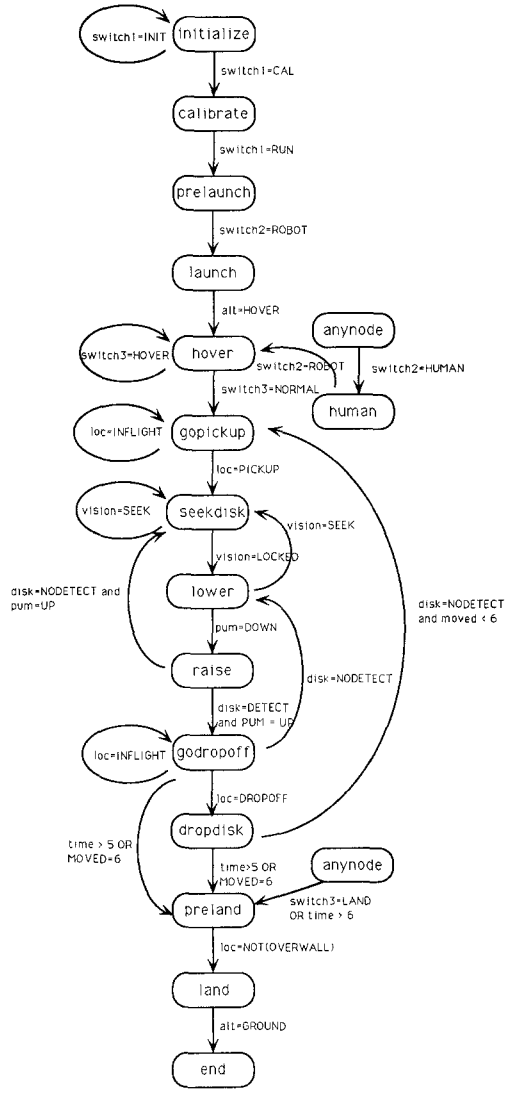

Fig. 4. High-level state diagram

be chosen by the team captain prior to takeoff. The choice will be based on a variety of factors including the weather, our performance to date, state of disrepair of the craft, current point totals in the competition, evaluation of opponents' potential performance, etc. The supervisory control system does not have the capability to change the strategy in mid-flight. For example one simple strategy may be based on the assumption that the pickup mechanism has failed and therefore acquisition of a disk is impossible. In this case to maximize point accumulation we would take-off, fly for 30 seconds, and land. An-

other example with more complex decision making would be when there is a choice between landing and acquiring or delivering another disk. If only 30 additional points are required to win, risking the acquisition and delivery of another disk would not be the best strategy.

External sensory inputs trigger transitions between states in the state diagram. The behavior which is exhibited by this simple mechanism may appear to be complex, even though its implementation in software is straightforward. When power is applied, the AAV is in the "initialize" state. An external switch moved to the calibrate position transitions the craft to the "calibrate" state. Once the vehicle has finished the calibration task and the engine has been started, the external switch is changed to the RUN position. This commands the vehicle to transition to the "prelaunch" state which waits until the remote control switch is placed into the COMPUTER position whereby the vehicle will "launch." The sequence of events which now occurs is dependent upon previous states and sensory input. If the remote control unit is switched to HUMAN, the computer will immediately transition into the "human" state from which the pilot commands the vehicle. This transition will override all other switch settings.

If the second switch on the remote control is placed in the HOVER position, the craft will transition immediately "hover" and remain there until the switch is changed. In the LAND position the computer will immediately transition to the "preland" state and land when the vehicle has determined that it has moved away from the 3 foot high barrier. If the switch is placed in the ABORT position the computer will immediately transition to the LAND state thereby commanding the craft to land regardless of its current location. While in the NORMAL switch position, the state machine will transition between "hover," "gopickup," "seekdisk," "lower" (pickup mechanism), "raise," "godropoff," and "dropdisk" as the strategy, current state, and sensory inputs demand.

6.2. *Navigation*

The command generator uses high level commands passed from the robot pilot state machine and preprocessed sensor data to set reference values for the stabilization module (Figure 5). This module executes a minimum of 10 times each second. If the vehicle is under partial or full human control, the command generator receives high level commands from the remote control. The command generator sends its signals through prefilters designed to help improve the response of the closed-loop system. The filtered signals are then used as reference inputs to the stability augmentation and attitude controller.

6.3. *Stability Augmentation and Attitude Control for 1992*

6.3.1. Motivation and Background. In this section, the automatic controller design techniques being employed for the second year (1992) UTA effort will be presented. Maintaining adequate control was a problem with the 1991 UTA vehicle. Classical controller design techniques were used, and a considerable amount of ad hoc tuning was required when the system was implemented on the vehicle. Even then, the best control achieved was not very precise, and an undershoot of minimum altitude contributed to the crash at the competition. The additional time and resources available for design of the 1992 entry will be used to develop and implement a much

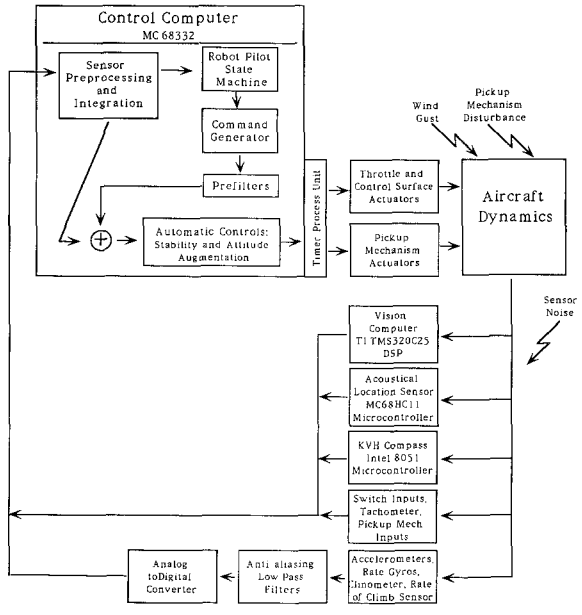

Fig. 5. Control system block diagram

more robust stability augmentation system. This section will report on the design of the block labeled "Automatic Controls" in Figure 5, and in particular on the design of the altitude/heading control loop. During this effort, we have made some extensions and modifications to a recently published technique for mixed $\mathcal{H}_2/\mathcal{H}_\infty$-optimal control design. These results are reported.

There are many uncertainties in the model that may be attributed to nonlinearities (e.g., aileron induced drag effect upon altitude hold), modeling errors, ignored dynamics, discretization errors, disturbances, sensor noise, sensor bias/hysteresis, and more. Therefore, the controller design should be maximally robust in terms of both stability and performance. It must also satisfy the physical constraints imposed by available sensor quality and bandwidth, digital sampling rate, and finite power control effectors.

The controls to be implemented have been designed to a mixed $\mathcal{H}_2/\mathcal{H}_\infty$ optimality criteria [4]. The design technique employed is an extension of an $\mathcal{H}_2$-optimization method commonly known as LQG/LTR (Linear Quadratic Gaussian with Loop Transfer Recovery); see [5, 6]. While it has been shown that optimal LQG designs can exhibit arbitrarily poor stability margins [7], it has also been demonstrated that $\mathcal{H}_2$-optimal controllers may generally achieve a good combination of nominal performance and stability robustness properties [8]. On the other hand, $\mathcal{H}_\infty$-optimal designs having wide stability margins are usually possible, but good performance with these controllers may be more difficult to realize. In fact, one usually opts for a sub-optimal design in order to retain better performance. For the case when the order of the controller is fixed at less than or equal to the order of the plant model, it has also been shown that the design equations yielding an $\mathcal{H}_\infty$ controller are a generalization of the classic estimator and regulator Riccati equations used in LQG designs [9, 10]. So it is natural to consider a combined optimization problem.

A straightforward design methodology to perform frequency domain loop shaping with a mixed optimality criterion has recently been presented by Yeh, et al. [11]. A mixed $\mathcal{H}_2/\mathcal{H}_\infty$ version of the Doyle and Stein condition [5] which forms the basis of the loop transfer recovery procedure was derived. The presentation herein will closely follow the development in [11] with three notable modifications. First, the dual problem for output loop shaping rather than input loop shaping will be developed. Recovery to the plant output is generally more appropriate when good command following is required, and it is essential whenever the plant model has more inputs than outputs. While the dual nature of the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ optimal control problem is already well established [12, 13], it is hoped that this presentation will help to clarify the design process as it relates to the theory and notation presented in [10]. Second, a significant restriction on the choice of a design parameter will be removed. We have found that the removal of this restriction affords the designer an additional freedom that allows for better performance to be maintained while improving the stability margins. Finally, we will present a modified loop design and recovery procedure which can be used to further refine the optimum tradeoff between performance and stability robustness.

6.3.2. Essential $\mathcal{H}_2/\mathcal{H}_\infty$-Optimization Theory. A complete discussion of the combined $\mathcal{H}_2/\mathcal{H}_\infty$-optimal control problem formulation is presented in Appendix A.

6.3.3. $\mathcal{H}_2/\mathcal{H}_\infty$-Optimal Control Design Applied to the AAV Altitude/Heading Control.

6.3.3.1. INTRODUCTION. The proposed new vehicle is a tail-sitter similar in concept to the 1991 aircraft. It is modeled as an airplane in vertical flight with the roll or x-axis nominally perpendicular to the earth (Figure 2). Essential aircraft flight dynamics may usually be captured in 8th order models which can be decoupled into a pair of 4th order models representing so-called longitudinal and lateral modes. However, the unusual characteristics of this tail-sitter operating close to hover at all times naturally led to the separations presented herein.

The current configuration (it is still being refined) calls for primary altitude control via a throttle, δ_T, and roll attitude control via a set of ailerons, δ_a. Both past experience and the model below demonstrate significant coupling between altitude and roll control through the control connection matrix. The engine relatively slowly compared to other control actions, so it is a pri-

mary limiting factor on the final bandwidth of the system. Also, the relatively large ratio of engine and propeller gyroscopic inertia to vehicle inertia, together with prop-wash and direct reaction torque effects, means that throttle setting changes induce large rolling moments.

6.3.3.2. THE ALTITUDE/HEADING MODEL. With reference to axis and control surface definitions in Figure 2, the aircraft altitude/heading nominal model may be represented in state-space by

$$\dot{x} = Ax(t) + Bu(t), \tag{6.1a}$$
$$y(t) = Cx(t), \tag{6.1b}$$

where

$$A = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & X_u & 0 & 0 & X_{\omega_P} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & L'_u & 0 & L'_p & L'_{\omega_P} \\ 0 & 0 & 0 & 0 & -1 \end{bmatrix};$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix};$$

$$B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ L'_{\delta_T} & L'_{\delta_a} \\ k_{\delta_T} & 0 \end{bmatrix};$$

The states $\mathbf{x} = [\, x\ u\ \phi\ p\ \omega_P\,]^T$ are vertical position (perturbed altitude), vertical velocity, roll angle (heading), roll rate, and propeller speed. The controls $\boldsymbol{u} = [\delta_T\ \delta_a]^T$ are the normalized throttle and aileron inputs. Note that $|\delta_a|$ is treated as a deterministic disturbance input to the vertical direction states. Simulations to follow will show its effect.

The coefficients of the model (so-called stability derivatives), to which the control system presented herein is designed, were obtained almost exclusively from purely theoretical approximations. For example,

$$X_{\omega_P} = \frac{1}{m}\left[\frac{\partial F_T}{\partial \omega_P} - \frac{\partial \mathcal{D}}{\omega_P}\right] \tag{6.2}$$

where m is the vehicle mass, F_T is the thrust force, and $\mathcal{D}$ is the drag force. Our airframe group provided estimates of final gross weight and total drag. Then, theoretical and empirical relationships between propeller thrust, size, and shaft horsepower along with an estimate of propeller efficiency were used to develop relationships for $F_T = f(\omega_P)$. Drag force was similarly treated, and the necessary stability derivative calculated. The resulting model parameters are

$$X_u = -0.06987\,\frac{1}{s}, \quad X_{\omega_P} = 0.0366\,\frac{ft}{s},$$
$$L'_u = -0.1904\,\frac{1}{ft \cdot s}, \quad L'_p = -9.521\,\frac{1}{s},$$
$$L'_{\omega_P} = 0.018773\,\frac{1}{s}, \quad L'_{\delta_T} = -18.733\,\frac{1}{s^2},$$
$$L'_{\delta_a} = 59.82\,\frac{1}{s^2}, \quad k_{\delta_T} = 59.82\,\frac{1}{s^2}.$$

The two assumed measurements are derived from data taken by the on-board sensor package. From an analysis of the available (affordable) sensors, we expect these measurements to begin to exhibit increasingly significant error above about 4 rad/s.

6.3.3.3. AAV CONTROLLER DESIGN AND RESULTS. The loop transfer recovery procedure presented in Appendix A.5 was used to design a combined $\mathcal{H}_2/\mathcal{H}_\infty$-optimal controller. A desirable loop shape to recover with an open-loop plant is one that is integrating (for good tracking), satisfies low frequency minimum magnitude constraints (for good disturbance rejection), has a 0 dB cross-over slope of about -20 dB/decade (for good nominal stability margins), and has a reasonably limited bandwidth (for good noise rejection). Often a desirable loop shape can not be achieved with *any* choice of D_1, and augmentation of the plant with additional dynamics is required. However, since the AAV and its model are designed with this optimization procedure in mind, the model is already square and integrating as required.

Besides loop shapes, another factor in the controller design is its maximum frequency. Since the sample rate is planned at 40 Hz, we decided to place a limit of 4 Hz (25 rad/s) on the controller. This value should allow us to implement the controller with acceptably small discretization error. It will also help in avoiding control servo saturations.

Figure 6 shows the open-loop plant, target feedback loop, and Kalman filter loop transfer functions. Also illustrated are the constraints placed upon the loop shape per the above discussion. The design parameter choices are given in Appendix B. These choices result in an $\mathcal{H}_2$ controller with a maximum frequency just under 25 rad/s. Note that per the theoretical development in Appendix A, we have selected $E_{1\infty} \neq C$. We were never able to obtain results as good as described below when the equality was enforced. Of course, this is not to say that better results can not be obtained, but we do assert that the extra freedom of adjusting $E_{1\infty}$ makes reaching a given level of performance and robustness "easier" to accomplish.

Figure 7 shows the final design loop shapes (solid lines) and the sensitivities. The maximum filter loop singular value is plotted (dashed line) to aid in comparing plots. At the final design point, D_1 has been increased by a factor of two over the values given in Appendix B. For the sensitivity and complementary sensitivity plots, the dashed lines illustrate the $\mathcal{H}_2$ controller for comparison. The maximum sensitivity has improved from 5.29 to 2.55 dB, and the maximum complementary sensitivity from 3.44 to 1.33 dB. The final adjustment made to D_1 accounted for about 0.12 dB of the improvement. It also raised the controller maximum frequency from about 23 to 25 rad/s.

Figures 8 and 9 give the step altitude and step roll command responses. A low-pass filter with a break frequency of 0.7 rad/s is used to shape the altitude command. The solid lines are for the final $\mathcal{H}_2/\mathcal{H}_\infty$ controller, the dashed are before increasing D_1, and the dash-dot are for the $\mathcal{H}_2$ controller. Figure 9 shows increase of D_1 resulted in recovery of some of the $\mathcal{H}_2$ controller nominal performance.

From tests that are currently being conducted we have determined that some of our parameter estimates were off by as much as a factor of three, and the controller coefficients will therefore need to be recalculated. However, we decided to try the controller designed above on the

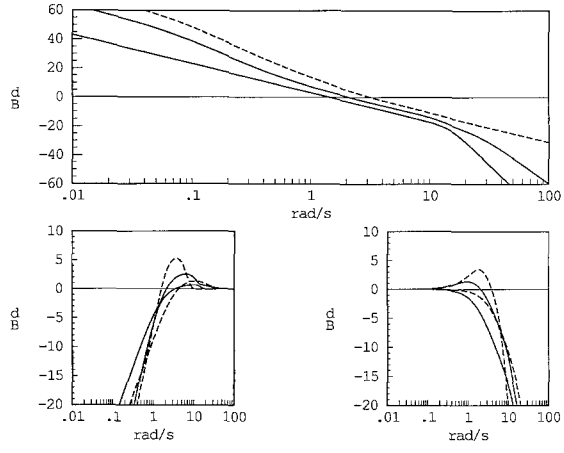

Fig. 7. Final recovered loop and sensitivities

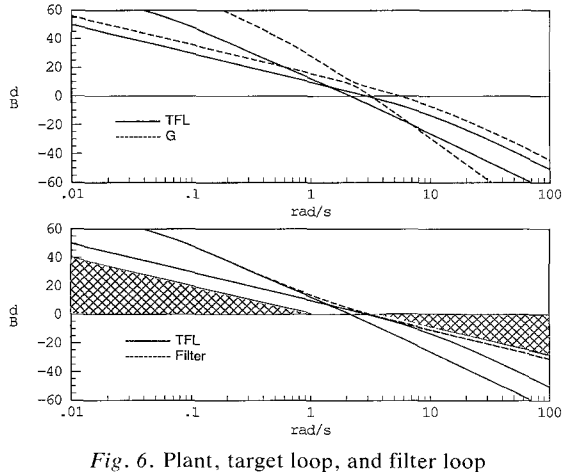

Fig. 6. Plant, target loop, and filter loop

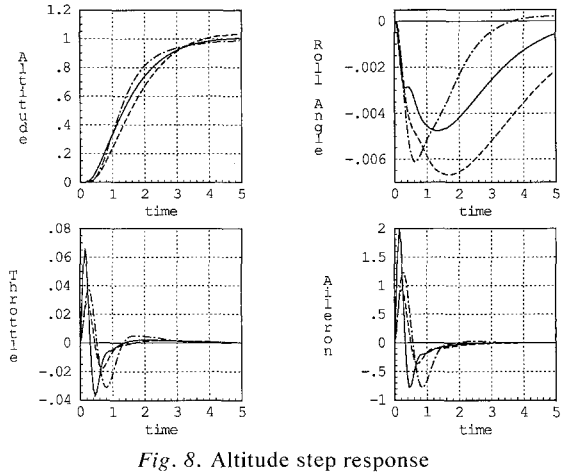

Fig. 8. Altitude step response

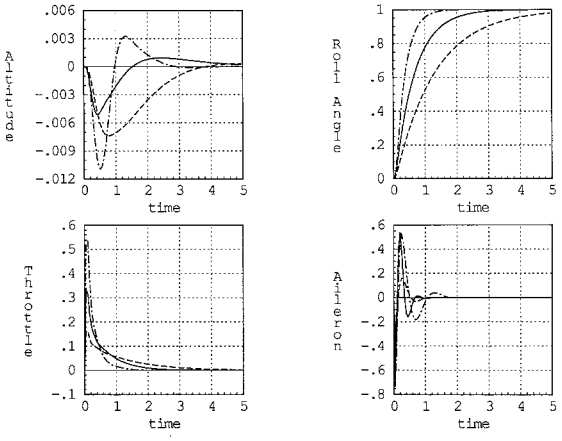

Fig. 9. Roll step response

revised plant model. We found very little change in the responses. Indeed, the controller design has proven to be robust to every worse case scenario that has been tried so far.

Figure 10 illustrates the responses obtained using the final controller design in a much more complete and realistic nonlinear model. This model includes actuator dynamics, saturation limits, and most significantly the nonlinear $|\delta_a|$ coupling to the vertical acceleration equation. The solid lines for comparison are the same as given in Figure 9. The dashed lines are for the nominal plant given above, and the dash-dot lines are for the revised plant.

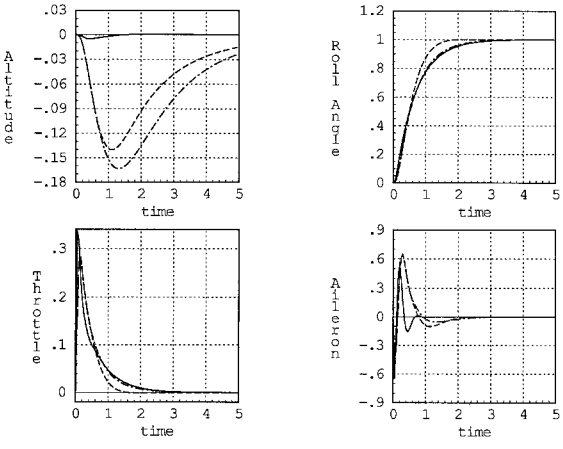

Fig. 10. Nonlinear and perturbed plant step response

6.4. *Human-Computer Control*

One problem when developing an autonomous control system is the transition period between human and computer pilot. The UTA AAV team has developed a unique solution to this problem. The control computer reads four proportional channels (throttle, pitch, roll, and yaw) and three nonproportional channels (on/off) from the receiver on-board the craft. One non- proportional channel is used to transfer control between the on-board computer and the human pilot. Two more non-proportional channels are used to set the mode of the vehicle. This switch is used to command the vehicle to continue performing the task, or it can be used to force the vehicle to go into hover, or to land immediately.

While the vehicle is under human control the computer passes three PWM signals directly to the outputs without modification. In this mode the human pitch roll and yaw commands will fly the craft. When the human switches to select computer control the computer then generates the output. This allows human control of pitch and yaw while working on a simple control scheme for the roll and altitude channels.

Unused receiver channels can also be used as reference inputs. For example a channel can be dedicated to the altitude reference signal and a step input to this reference signal can be changed from the ground while the vehicle step response is examined. This allows a reference signal or other variable in the control software such as control gains or time constants to be changed during flight.

The 1992 vehicle will incorporate a control scheme which enables a less experienced test pilot to fly the craft. New pilots have a difficult time learning the complex transformations required to place themselves in the cockpit on the craft. This is complicated even further in a tail-sitter. For example pitch and yaw are reversed when the craft has rolled 180 degrees. The stabilization of the craft will be performed directly by the control computer. The pitch, roll, and yaw commands in human control mode will control a commanded navigation position in a 3D coordinate system relative to the human pilot. Since the craft will be maintaining a constant heading, zero pitch, and zero yaw during the entire flight plan,

the pilot will not have control over the pitch, roll, or yaw axes.

7. Conclusions and Future Research

Participation in this challenging event has afforded students a valuable opportunity to apply their classroom studies and to broaden their skills. The work has led to a better understanding and appreciation by all involved for the kinds of technical and nontechnical problems and that must be identified and overcome in practice. The work has also spurred our research in the area of applied optimal controls. Our effort to design a combined $\mathcal{H}_2/\mathcal{H}_\infty$-optimal controller for the aircraft has resulted in an extension to the loop transfer recovery theory. A future research goal is to further extend and clarify the meaning of these results. We also will emphasize the reduced-order problem. Another goal for the future is to investigate the use of neural networks and/or adaptive techniques to try to overcome some of the shortcomings of low-cost sensors.

Acknowledgments

This project would not have been possible without the support of the following corporate and individual sponsors: Sky Technology, Microtech Research, Inc., Polaroid Corporation, Crossland Research, Motorola, Hughes Simulation Systems, Roy's Hobby Shop, UTA College of Engineering, IEEE-Fort Worth Section, Bell Helicopter Textron, Collmer Semiconductor, E-Systems, Futaba Corporation of America, KVH Corporation, Bits-n-Bytes, LTV Missiles and Electronics, General Dynamics, Apple Computer, Tandy Corporation, Texas Instruments, and National Semiconductor.

Kliffton M. Black is a Laboratory Graduate Fellow sponsored by the Air Force Office of Scientific Research, Bolling AFB, DC. His research is sponsored by the AFOSR AFSC, USAF, under contract FH6920-86-C-0127.

In addition to the authors, the following individuals have made significant contributions toward the technical success of this project: Jay Bowen, Randy Lawrence, Matt Wilson, William Page, Stephen Maduzia, Mike Walti, Shane Scott, Steve Spence (Pilot), and Dr. D. Tuckness, J.D. Sanders, Rick Roberts, Derek Barnhill, Randy Siemens, Paul Wanamaker, Kyle Barker, and Rodney Fannin.

Appendix A

A.1. The Combined $\mathcal{H}_2/\mathcal{H}_\infty$-Optimal Control Problem Formulation

We begin by introducing the $\mathcal{H}_\infty$-constrained LQG control problem. This problem may be posed in the following manner using the notation and theory developed in [10]. Consider the state-space system description

$$\dot{x} = Ax(t) + Bu(t) + D_1 w(t), \tag{A.1a}$$

$$y(t) = Cx(t) + D_2 w(t), \tag{A.1b}$$

$$z(t) = E_1 x(t) + E_2 w(t), \tag{A.1c}$$

$$z_\infty(t) = E_{1\infty} x(t) + E_{2\infty} w(t) \tag{A.1d}$$

having one exogenous input vector and two observed output vectors where $A \in \mathcal{R}^{n \times n}$, $B \in \mathcal{R}^{n \times m}$, $C \in \mathcal{R}^{l \times n}$, $D_1 \in \mathcal{R}^{n \times p}$, $D_2 \in \mathcal{R}^{l \times p}$, $E_1 \in \mathcal{R}^{q \times n}$, $E_2 \in \mathcal{R}^{q \times m}$, $E_{1\infty} \in \mathcal{R}^{q_\infty \times n}$, $E_{2\infty} \in \mathcal{R}^{q_\infty \times m}$. The triple (A, B, C) represents the stabilizable and detectable minimum phase nominal system. All other matrices in are taken to be "free" design parameters subject the following standard assumptions and definitions [12]:

i) (A, D_1) stabilizable and $(E_{1\infty}, (A)$ detectable;

ii) $E_1^T E_2 = 0; \quad E_2^T E_2 = R_2 > 0;$

iii) $E_{1\infty}^T E_{2\infty} = 0; \quad E_{1\infty}^T E_{1\infty} = R_{1\infty}; E_{2\infty}^T E_{2\infty} = R_{2\infty};$

iv) $D_1 D_2^T = 0; \quad D_1 D_1^T = V_1;$

$D_2 D_2^T = V_2 > 0.$

The design objective is to design a stabilizing linear feedback controller

$$K(s) = C_c(sI - A_c)^{-1} B_c = \begin{bmatrix} A_c & B_c \\ C_c & 0 \end{bmatrix} \tag{A.2}$$

such that

$$\| H(s) \|_\infty \geq \gamma \tag{A.3}$$

and the cost functional

$$\mathfrak{g} = \mathrm{tr}\{Q\tilde{R}\} \tag{A.4}$$

is minimized.

The closed-loop transfer function whose $\mathcal{H}_\infty$-norm is to be bounded by some given constant γ is given by

$$H(s) = \tilde{E}_\infty(sI - \tilde{A})^{-1}\tilde{D} = \begin{bmatrix} A & BC_c & D_1 \\ B_cC & A_c & B_cD_2 \\ E_{1\infty} & 0 & 0 \end{bmatrix} = \begin{bmatrix} \tilde{A} & \tilde{D} \\ \tilde{E}_{1\infty} & 0 \end{bmatrix}. \quad (A.5)$$

It can be shown that the cost functional $\mathfrak{g}$ is dependent upon

$$\mathfrak{Q}: \quad 0 = \tilde{A}\,\mathfrak{Q} + \mathfrak{Q}\tilde{A}^{\mathrm{T}} + \gamma^{-2}\mathfrak{Q}\tilde{R}_\infty\,\mathfrak{Q} + \tilde{V} \quad (A.6)$$

where

$$\tilde{R}_\infty = \tilde{E}_\infty^{\mathrm{T}}\tilde{E}_\infty$$

$$\tilde{V} = \tilde{D}\tilde{D}^{\mathrm{T}}\,\tilde{R} = \begin{bmatrix} R_1 & 0 \\ 0 & C_c^{\mathrm{T}}R_2C_c \end{bmatrix}$$

It can also be shown that when w is standard Gaussian white noise, $\mathfrak{g}$ provides an optimal overbound for the usual LQG cost functional

$$J = \lim_{t\to\infty} \int_0^t (z^{\mathrm{T}}z)dt. \quad (A.7)$$

That is, as outlined in [8], the results of [14] demonstrate that LQG problem can be converted into an equivalent $\mathcal{H}_2$-problem via Parseval's theorem. Thus, a weighed tradeoff between system transfer functions can be accomplished with a loop transfer recovery procedure to the plant output breakpoint.

Making the simplification $E_{2\infty} = \beta E_2$ and taking $\beta = 0$, then the following lemma due to Bernstein and Haddad [10] gives sufficient conditions for closed-loop stability with mixed $\mathcal{H}_2/\mathcal{H}_\infty$ optimality (see [12] for relationships to the dual development in [15]).

Lemma A.1 – Suppose there exists positive semidefinite solutions Q, P, and $\hat{Q}$ to

$$0 = AQ + QA^{\mathrm{T}} + V_1 + \gamma^{-2}QR_{1\infty}Q - QC^{\mathrm{T}}V_2^{-1}CQ, \quad (A.8)$$

$$0 = (A + \gamma^{-2}(Q + \hat{Q})R_{1\infty})^{\mathrm{T}}P + P(A + \gamma^{-2}(Q + \hat{Q})R_{1\infty}) + R_1 - PBR_2^{-1}B^{\mathrm{T}}P, \quad (A.9)$$

$$0 = (A - C^{\mathrm{T}}V_2^{-1}CP + \gamma^{-2}QR_{1\infty})\hat{Q} + \hat{Q}(A - C^{\mathrm{T}}V_2^{-1}CP + \gamma^{-2}QR_{1\infty})^{\mathrm{T}} + \gamma^{-2}\hat{Q}R_{1\infty}\hat{Q} - QC^{\mathrm{T}}V_2^{-1}CQ \quad (A.10)$$

such that

$$A_c = A - B_cC + BC_c + \gamma^{-2}QR_{1\infty}, \quad (A.11)$$

$$B_c = QC^{\mathrm{T}}V_2^{-1}, \quad (A.12)$$

$$C_c = -R_2^{-1}B^{\mathrm{T}}PS \quad (A.13)$$

stabilizes $\tilde{A}$ given in (A.5). Then $\| H(s) \|_\infty \leq \gamma$ and $J \leq \mathfrak{g}$.

It should be noted that the Lemma A.1 can be generalized to include nonzero $E_{2\infty}$ as well as controller order constraints (i.e., the theory provides for the direct design of reduced-order controllers).

A.2. Discussion

The first part of the design objective is to minimize the steady state power of the output z when the input w is taken to be Gaussian white noise. That is, minimize the $\mathcal{H}_2$-norm of the transfer matrix from w to z. This may generally be thought of as a time-domain performance constraint, and when coupled with LTR it also yields guaranteed minimimum stability margins. The remainder of the design objective is to minimize the magnitude of the output z_∞ when w is taken to be a deterministic signal having unity power. That is, bound by some constant, γ, the $\mathcal{H}_\infty$-norm which is the peak over all frequencies of the maximum singular value of the disturbance transfer matrix from w to z^∞. This objective serves to maximize robustness with respect to the worst case destablizing uncertainty that might be present.

In performing an $\mathcal{H}^2$-optimization, a formal tradeoff between transfer functions of the system is accomplished by freely using the usual LGQ design parameters as design variables. The approach is to manipulate one of these variables to achieve a desirable open-loop transfer function that may then be recovered by manipulating another variable.

In much the same manner as established for LQG/LTR design, the authors of [11] have detailed a procedure for implementing combined $\mathcal{H}^2/\mathcal{H}_\infty$-optimization via LTR. Our development will closely parallel theirs except that we will

deal with recovery to the plant output breakpoint instead of the input. As noted earlier, this is merely the dual of the output loop shaping problem, and it applies to plants that have no more outputs than inputs ($l \leq m$).

A.3. Removal of a Design Parameter Constraint

The following extends the results of [11] to allow for the case when $E_{1\infty} \neq C$. It is shown that a mixed $\mathcal{H}_2/\mathcal{H}_\infty$ version of the Doyle and Stein condition [5] can still be used to obtain a simple loop transfer function that can be asymptotically recovered. The structure of the mixed $\mathcal{H}_2/\mathcal{H}_\infty$ controller (Figure 11) is completely analogous to the standard LQG structure [16], except that we have added the $\mathcal{H}_\infty$ constraint block, $\Gamma = \gamma^{-2}QR_{1\infty}$. Let

$$\begin{aligned}\Phi_\gamma &= (sI - A - \Gamma)^{-1} \\ &= (\Phi - \Gamma)^{-1} \\ &= \Phi(I - \Gamma\Phi)^{-1}. \end{aligned} \tag{A.14}$$

Φ and Φ_γ are both functions of s. We will suppress the argument to make the equations more readable. If the loop is broken at point 1, then the return ratio y' is given by

$$\begin{aligned} y' = {} & C\Phi BC_c(I - \Phi_\gamma BC_c)^{-1}\Phi_\gamma B_c \mathrm{r} \\ & - \mathrm{C}(\mathrm{I} - \Phi_\gamma \mathrm{BC}_c)^{-1}\Phi_\gamma B_c r. \end{aligned} \tag{A.15}$$

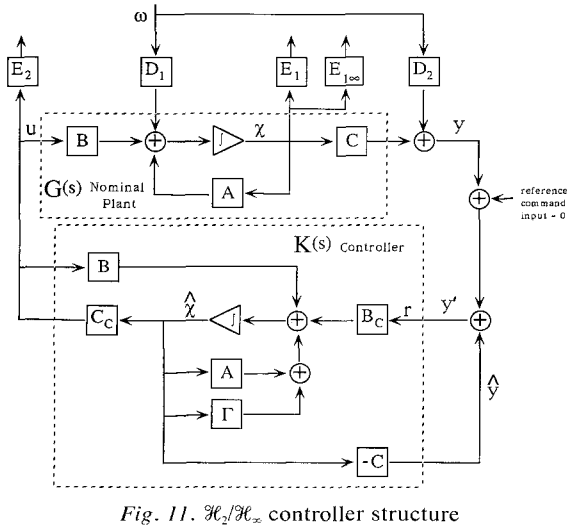

Fig. 11. $\mathcal{H}_2/\mathcal{H}_\infty$ controller structure

If $C\Phi B$ is not square, then the output matrix must be augmented. Let

$$\tilde{C} = \begin{bmatrix} C \\ C_a \end{bmatrix} \tag{A.16}$$

where C_a such that $(\tilde{C}\Phi B)$ and $(\tilde{C}\Phi B)$ are both minimum phase and invertible. Then we can write the output matrix as

$$C = [\, I\ 0\,]\tilde{C}. \tag{A.17}$$

Substituting into (A.15) and using some simple algebra, we obtain

$$\begin{aligned} y' = {} & [\, I\ 0]\Big\{\tilde{C}\Phi B(I - C_c\Phi_\gamma B)^{-1}C_c\Phi_\gamma B_c r - \tilde{C} \\ & [I + \Phi_\gamma B(I - C_c\Phi_\gamma B)^{-1}C_c]\Phi_\gamma B_c \mathrm{r}\Big\} \\ = {} & [\, I\ O\,]\tilde{C}\Phi B\Big\{(I - C_c\Phi_\gamma B)^{-1}C_c - \\ & [(\tilde{C}\Phi_\gamma B)^{-1}\tilde{C} + (I - C_c\Phi_\gamma B)^{-1}C_c]\Big\}\Phi_\gamma B_c r. \end{aligned} \tag{A.18}$$

Clearly, the second term of (A.18) can be made to vanish if C_c can be chosen such that

$$\begin{aligned} \tilde{C}\Phi_\gamma B)^{-1}\tilde{C} + (I \\ - C_c\Phi_\gamma B)^{-1}C_c = 0. \end{aligned} \tag{A.19}$$

In that case,

$$\begin{aligned} y' &= [\, I\ 0\,]\tilde{C}\Phi B[- (\tilde{C}\Phi_\gamma B)^{-1}\tilde{\mathrm{C}}]_\gamma B_c \mathrm{r} \\ &= - C\Phi[B(\tilde{C}\Phi_\gamma B)^{-1}\tilde{\mathrm{C}}\Phi_\gamma]B_c r \\ &= - C\Phi B_c r. \end{aligned} \tag{A.20}$$

The last line of (A.20) is a consequence of [17, Theorem 1.3.2]. Equation (A.19) is the output breakpoint $\mathcal{H}_2/\mathcal{H}_\infty$ version of the Doyle and Stein condition, and it was obtained in essentially the same manner as the input version was obtained in [11]. However, our derivation in (A.20) allows for $E_{1\infty} \neq C$ (equivalent to $B_1 \neq B_2$ in the input breakpoint formulation).

We will now show that (A.19) can be satisfied for the choice of C_c as given in (A.13). Let

$$A_P = (A + \gamma^{-2}(Q + \hat{Q})R_{1\infty}), \tag{A.21}$$

$$E_1 = \begin{bmatrix} M\tilde{C} \\ 0_{m \times m} \end{bmatrix} \tag{A.22}$$

$$E_2 = \begin{bmatrix} PI_m \\ 0_{m \times m} \end{bmatrix}, \tag{A.23}$$

where $M \in \mathcal{R}^{m \times m}$, M^{-1} exists. The augmentation giving $\tilde{C}$ and the transformation M must be chosen such that $M\tilde{C}[sI - A_p]^{-1}B$ is minimum phase. Substituting into (A.9),

$$0 = A_p^T P + PA_p + \tilde{C}^T M^T M\tilde{C} - PB\left(\frac{I}{\rho^2}\right)B^T P. \quad \text{(A.24)}$$

It can now be shown that (see [18])

$$\lim_{\rho \to 0} P = 0, \quad \text{(A.25)}$$

$$\lim_{\rho \to 0} PB\left(\frac{1}{\rho^2}\right)B^T P = \tilde{C}^T M^T M\tilde{C}, \quad \text{(A.26)}$$

$$\lim_{\rho \to 0} \frac{1}{\rho} B^T P = UM\tilde{C}, \quad \text{(A.27)}$$

where $U \in \mathcal{R}^{m \times m}$, $\quad U^T U = I$.

Substituting (A.23) into (A.13) and then combining with (A.27) and rearranging gives

$$C_c = -\lim_{\rho \to 0} \frac{UM\tilde{C}}{\rho}. \quad \text{(A.28)}$$

It is now easy to show that C_c does indeed satisfy the mixed $\mathcal{H}_2$ and $\mathcal{H}_\infty$ Doyle and Stein condition given by (A.19).

A.4. Implications of $E_{1\infty} \neq C$

For completeness, we should now examine the Kalman equality. When the selection $E_{1\infty} = C$ is made, then decreasing γ has precisely the same effect upon the modified filter equation (A.8) as a simple scaling of D_2 would have. That is, it serves primarily to increase the magnitudes of the loop singular values resulting in a higher bandwidth for the filter loop. This has the adverse effect of increasing controller bandwidth and final open-loop response bandwidth as the $\mathcal{H}_\infty$ norm constraint is reduced. In deriving a mixed version of the Kalman equality, the authors of [11] determined that a sensible alternative was to scale D_1 so that at low frequencies the filter loop principal gains are unchanged.

The choice $E_{1\infty} = C$ also has the effect of placing a practical lower bound of 1 on γ. However, for other choices of $E_{1\infty}$, both the lower bound on gamma and the scaling that should be applied to D_1 are unclear. Based primarily upon numerical experience, we have opted to retain the scaling used in [11]. That is, let

$$\alpha^{-2} = 1 - \gamma^{-1} \quad \text{(A.29)}$$

and use $\frac{1}{\alpha}D_1$ in place of D_1 when solving the modified filter equation (A.8) for Q. We have also chosen to scale any choice of $E_{1\infty}$ so that solutions to (A.8) exist for $\gamma > 1$, but do not exist for $\gamma \leq 1$. This scaling must be done carefully since it is possible for $\gamma = 1$ to be a singular point of (A.8).

A.5. A Modified Loop Transfer Recovery Procedure

The procedure outlined is for output breakpoint recovery, and it is applicable only to minimum phase plants having a number of outputs less than or equal to the number of inputs. A dual of this procedure may be applied to the input breakpoint of minimum phase plants having a number of inputs less than or equal to the number of outputs. The first five steps are identical to the standard LQG/LTR procedure. Before beginning, a nonsquare plant should be augmented according to (A.17). For formal loop shaping, the plant should also generally be integrating. So, it may be necessary to append additional dynamics to the input before proceeding; see [19].

1. Manipulate D_1 to achieve a good Target Feedback Loop (TFL) shape. The TFL transfer function is given by

$$L_{TFL}(s) = C\Phi(s)D_1. \quad \text{(A.30)}$$

2. Choose $D_2 = I$ and solve the Kalman Filter equation ($\gamma \to \infty$ in (A.8)) for Q_{H2}. We will use the subscript "H2" to refer to the solutions and controllers obtained with no enforcement of the infinity-norm constraint. Use (A.12) to obtain B_c and plot the Filter loop shape where

$$L_{Filt}(s) = C\Phi(s)B_{cH2}. \quad \text{(A.31)}$$

3. The Filter loop shape should match closely with the TFL at low frequencies, and it will exhibit a 20dB/decade roll-off at high frequencies. If the TFL appears unacceptable, return to step 1 and readjust D_1.

4. Select E_1 and E_2 according to (A.22) and (A.23) and let $\rho = 1$. Solve (A.9) for P_{H2}, form the controller according to (A.11)-(A.13), and plot final open loop shape $G(s)K_{H2}(s)$. If the open loop shape is not sufficiently close to the Target loop shape (or if the controller bandwidth is too high), then decrease (increase) ρ in equation (A.23) and repeat this step. If an adequate tradeoff can not be achieved, it may be necessary to return to step 1.

5. Examine the time responses of the resulting $\mathcal{H}_2$-optimal system design. Our experience indicates that these should be reasonably accepted before proceeding. A close examination of the tradeoffs that are being effected by the combined optimization problem also supports this approach. It is also advisable to compute the infinity-norms of $H(s)$ as given in (A.5) and of the Sensitivity and Complementary Sensitivity functions (see [20] for a computational method). With experience, these results can be used as a gauge to estimate what final results might be possible when the infinity-norm constraint is imposed. For reference, the Sensitivity and Complementary Sensitivity functions are

$$S(s) = \left(I + G(s)\,K(s)\right)^{-1}, \tag{A.32}$$

$$T(s) = G(s)K(s)\left(I + G(s)K(s)\right)^{-1}. \tag{A.33}$$

6. Choose $E_{1\infty} = C$, select $\gamma > 1$, and solve (A.8)-(A.10) for P, Q, and $\hat{Q}$. Obtain the controller using (A.11)-(A.13) and examine performance and robustness.

7. Iterate over γ and manipulate $E_{1\infty}$ to obtain desired or best possible levels of combined performance and robustness. D_1 might also be adjusted as discussed below. One should keep in mind that changes in $E_{1\infty}$ and D_1 affect the minimum value of γ, so rescaling of $E_{1\infty}$ may be necessary.

A.6. Some Observations and Numerical Considerations

We have had success in the limited number of examples studied to date by using some of the same guidelines for adjusting $E_{1\infty}$ as might be employed for initially selecting D_1. That is, we examine the transfer function $(E_{1\infty}\Phi\, D_1)$ and try to improve upon what might be called a modified Target loop response. On examination of (A.8), this approach makes sense because as γ is decreased, the modified filter "sees" a modified output.

Generally, as ρ is decreased all of the peak sensitivities (infinity-norms of $H(s)$, $S(s)$, and $T(s)$) decrease. The tradeoff is that the eigenvalues of the controller also increase, so a faster and more powerful control system is required. Also, the equations become harder to solve.

Figure 12 was generated from the AAV control problem solved in Section 6.3.3.3. Notice that the L_2 cost peaks at around $\ln(\gamma) = 1.6$. This behavior is the result of scaling D_1 and of selecting $E_{1\infty} \neq C$.

As a final tuning parameter, we have sometimes found it advantageous to also adjust the magnitude of D_1 when nearing the final design. The $\frac{1}{\alpha}$ scaling of D_1 sometimes causes the L_2 costs and the eigenvalues of the controller to decrease with γ after it has been lowered below some critical level. When this has occurred, D_1 can sometimes be increased by a some scalar magnitude to recover performance but with little or no adverse effect on the stability margins. The design application presented herein exhibits this behavior.

As effective numerical algorithm for solving

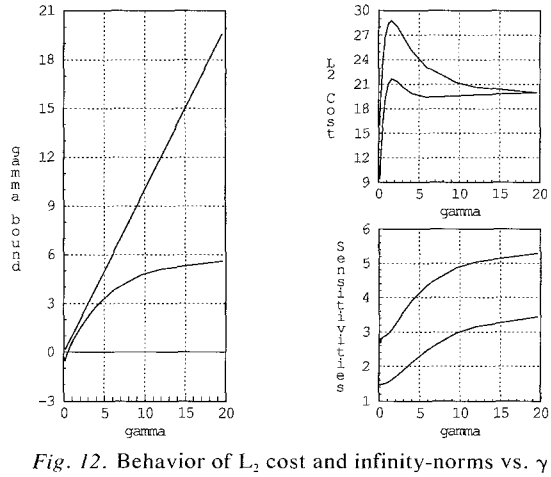

Fig. 12. Behavior of L_2 cost and infinity-norms vs. γ

(A.8)-(A.10) has been reported in [21] and [22]. It is based upon rewriting the $\hat{Q}$ equation in a Lyapunov form dependent upon the prior step solution and then iterating between the highly coupled P and $\hat{Q}$ equations. It exhibits very good numerical stability and can obtain solutions arbitrarily close to the minimum γ for a given problem.

Appendix B—Control Design Parameters

$$D1 = \begin{bmatrix} 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ -5.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ -0.9367 & 29.9100 & 0.0000 & 0.0000 \\ 49.9500 & 0.0000 & 0.0000 & 0.0000 \end{bmatrix}$$

$$D2 = \begin{bmatrix} 0. & 0. & 1. & 0. \\ 0. & 0. & 0. & 1. \end{bmatrix}$$

$$E1 = \begin{bmatrix} 1. & 0. & 0. & 0. & 0. \\ 0. & 0. & 1. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. \\ 0. & 0. & 0. & 0. & 0. \end{bmatrix}$$

$$E2 = \begin{bmatrix} 0.0000 & 0.0000 \\ 0.0000 & 0.0000 \\ 0.1400 & 0.0000 \\ 0.0000 & 0.1400 \end{bmatrix}$$

$$E1INF = \begin{bmatrix} 0.6250 & 0.3125 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.6250 & 0.1875 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{bmatrix}$$

References

1. K. Dzialo and R. Schalkoff, "Control implications in tracking moving objects using time-varying perspective projective imagery," *IEEE Transactions on Industrial Electronics*, 1986.
2. A. Gilbert, "A real-time video tracking system," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1980.
3. M. Baird, "Image segmentation technique for locating automotive parts on belt conveyors," in *IEEE Proceedings on Image Understanding II*, 1987.
4. J.C. Doyle, K. Glover, P.P. Khargonekar and B.A. Francis, "State-space solutions to standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems," *IEEE Transactions on Automatic Control*, vol. 34, pp. 831–847, 1989.
5. J.C. Doyle and G. Stein, "Multivariable feedback design: Concepts for a classical/modern synthesis," *IEEE Transactions on Automatic Control*, vol. AC-26, pp. 4–16, 1981.
6. J.M. Maciejowski, *Multivariable Feedback Design*, Addison-Wesley: New York, 1989.
7. J.C. Doyle, "Guaranteed margins for LQG regulators," *IEEE Transactions on Automatic Control*, vol. AC-23, pp. 756–757, 1978.
8. G. Stein and M. Athans, "The LQG/LTR procedure for multivariable feedback control design," *IEEE Transactions on Automatic Control*, vol. AC-32, pp. 105–114, 1987.
9. I.R. Petersen, "Disturbance attenuation and $\mathcal{H}_\infty$ optimization: A design method based on the algebraic Riccati equation," *IEEE Transactions on Automatic Control*, vol. AC-32, pp. 427–429, 1987.
10. D.S. Bernstein and W.M. Haddad, "LQG control with

an H_∞ performance bound: A Riccati equation approach," IEEE Transactions on Automatic Control, vol. 34, pp. 293–305, 1989.
11. H.-H. Yeh, S.S. Banda, and A.G. Sparks, "Loop shaping in mixed $\mathcal{H}_2$ and $\mathcal{H}_\infty$ optimal control," *Proceedings of the American Control Conference,* 1991.
12. H. Yeh, S.S. Banda and B. Chang, "Necessary and sufficient conditions for mixed $\mathcal{H}_2$ and $\mathcal{H}_\infty$ optimal control," to appear *IEEE Transactions on Automatic Control,* also in *Proceedings of the 29th IEEE Conference on Decision and Control,* 1990.
13. A.G. Sparks, S.S. Banda, and H.-H. Yeh, "A comparison of loop shaping techniques," preprint submitted to the *30th IEEE Conference on Decision and Control Control,* 1991.
14. M.G. Safonov, A.J. Laub and G.L. Hartmann, "Feedback properties of multivariable systems: The role and use of the return difference matrix," *IEEE Transactions on Automatic Control,* vol. AC-26, pp. 47–65, 1981.
15. J.C. Doyle, K. Zhou and B. Bodenheimer, "Optimal control with mixed H_2 and H_∞ performance objectives," in *Proceedings of the American Control Conference,* pp. 2065–2070, 1989.
16. B.D.O. Anderson and J.B. Moore, *Linear Optimal Control,* Prentice-Hall: NJ, 1971.
17. S.L. Campbell and C.D. Meyer, Jr., *Generalized Inverses of Linear Transformations,* Pitman, Ltd.: London, 1979.
18. W.M. Wonham, *Linear Multivariable Control: A Geometric Approach,"* 3rd edition, Springer-Verlag: New York, 1985.
19. D.B. Ridgely and S.S. Banda, "Introduction to robust multivariable control," Control Dynamics Branch, FDL Report AFWAL-TR-85-3102, AFSC, WPAFB, OH 45433, February 1986.
20. N.A. Bruinsma and M. Steinbuch, "A fast algorithm to compute the $\mathcal{H}_\infty$-norm of a transfer function matrix," *Systems and Control Letters,* vol. 14, pp. 287–293, 1990.
21. K.M. Black, "A new algorithm for solving a mixed $\mathcal{H}_2/\mathcal{H}_\infty$ optimization problem," in *Proceedings of the American Control Conference,* 1991.
22. K.M. Black, A. Sparks, S.S. Banda and H. Yeh, "Combined $\mathcal{H}_2/\mathcal{H}_\infty$ optimal control design applied to an advanced fighter aircraft," in *30th Aerospace Sciences Meeting,* Reno, NV, 1992.

Jeffrey O. Smith is currently an Assistant Instructor and Ph.D. candidate in the Computer Science Engineering Department at The University of Texas at Arlington. He also works as a Computer Specialist at the Super-Conducting Super-Collider. Prior to attending U.T.A. Jeff was a Sr. Analyst and Development Engineer at Motorola in Fort Worth, Texas. While at U.T.A. Jeff received The American Electronics Association Faculty Development Fellowship 1989–1991 and Outstanding Contribution Award in Computer Science 1989. He was project leader of the winning entry in the First International Aerial Robotics Competition 1991. Jeff is a member of Tau Beta Pi, Upsilon Pi Epsilon, IEEE, ACM, and AAAI.

Kliffton M. Black received B.S. (1987) and M.S. (1990) degrees in mechanical engineering from the University of Texas at Arlington where he is now completing work toward a Ph.D. with emphasis in control system theory and design.

In 1989, he was named as a recipient of a Laboratory Graduate Fellowship sponsored by the United States Air Force Office of Scientific Research, Bolling AFB, DC. He is a member of Tau Beta Pi and Pi Tau Sigma, and he is currently serving as an assistant editor for the ASME Dynamic Systems and Controls newsletter. His main research interests are applied optimal control of multivariable systems and stability robustness measures in the time domain.

Fahrad A. Kamangar Associate Professor in Computer Science Engineering Department (1983). B.S. University of Tehran, 1975; M.S. University of Texas at Arlington, 1977; Ph.D. 1980.

Areas of Research: Artificial Intelligence Computer Graphics, Systems Programming, Digital Signal Processing.

Jack Fitzer Professor in Electrical Engineering Department (1967). B.S. University of Missouri, 1951; M.S. Washington University, 1960; D.Sc. 1962.