JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 25, No. 3, May–June 2002

Comparison of Flight Control System Design Methods Using the CONDUIT® Design Tool

Mark B. Tischler
U.S. Army Aeroflightdynamics Directorate, Moffett Field, California 94035-1000
James A. Lee
University of Southern California, Los Angeles, California 90089
and
Jason D. Colbourne
U.S. Army Aeroflightdynamics Directorate, Moffett Field, California 94035-1000

Optimization and comparison of several alternative control system design methods against a common extensive set of dynamics response criteria is demonstrated using the Control Designer's Unified Interface (CONDUIT®). The alternative methods considered are classical, linear quadratic regulator, dynamic inverse, and *H*-infinity as applied to the design of lateral/directional control laws for a transport aircraft. From poor initial guesses for the design parameters of each alternative method, CONDUIT first achieved a feasible design space that satisfied the stability and handling qualities to the best (level 1) criteria. Final controller tuning was accomplished to minimize the performance metrics of crossover frequency and actuator activity, while maintaining level 1 design criteria. An important finding of this research is that the alternative design methods optimized against a common set of design requirements yield controllers whose performance and stability robustness characteristics are quite similar to one another. A stronger discriminator than design method is the controller architecture (one or two degree of freedom), which plays an important role in determining the achievable design space. This research demonstrates the feasibility and emphasizes the need to analyze and optimize prospective control designs against a comprehensive set of design requirements. CONDUIT has proven to be an especially effective environment for this task.

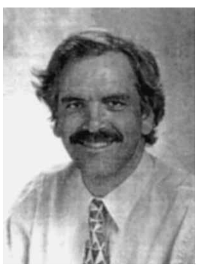

Mark B. Tischler, received B.S. and M.S. degrees in Aerospace Engineering from the University of Maryland, and his Ph.D. from the Department of Aeronautics and Astronautics at Stanford University. Tischler joined Ames Research Center in 1983, where he currently serves as Flight Control Technology Group Leader in the Army/NASA Rotorcraft Division, Aeroflightdynamics Directorate, leading research activities in rotorcraft manned and unmanned flight control, handling-qualities, system identification, and simulation. Tischler has authored or coauthored over 75 technical papers and reports and was the organizing editor and contributing author of the book *Advances in Aircraft Flight Control* (AIAA and Taylor & Francis). He is an AIAA Associate Fellow.

James A. Lee received a medical degree (M.D.) from the University of Texas, Houston in 1979. He earned an M.S. in Aerospace Engineering at the University of Colorado, Boulder in 1994. He attended Stanford University in 1996 for postmaster's work in astronautics. He is working on a second master's degree in Electrical Engineering at the University of Southern California. It was during this time he collaborated with Ames Research Center U.S. Army Aeroflightdynamics Directorate, on research involving the CONDUIT® program. Lee currently works for a research firm (VectorAero, a subsidiary of Network Telephone Services) as a Systems Engineer. He is an AIAA Senior Member.

Jason D. Colbourne obtained a B.S. degree in Aerospace Engineering from San Diego State University in 1993, and an M.S. degree in Aeronautical Engineering from California Polytechnic State University in 1997. Colbourne has worked as a Research Scientist in the Army/NASA Rotorcraft Division, Aeroflightdynamics Directorate, at Ames Research Center since 1995, with an emphasis on flight dynamics and control of rotorcraft unmanned air vehicles.

Received 9 July 2001; revision received 10 November 2001; accepted for publication 10 November 2001.

Introduction

THE U.S. Air Force has recently completed an extensive review of flight control practices for avoiding pilot-induced oscillations (PIO),[1] in response to that these problems have occurred during the development process for almost every new military aircraft. The role of checking and designing to meet handling-qualities criteria throughout the flight vehicle development life cycle figures prominently in the "Ten Steps to Reducing the Risk of PIO" reviewed in Sec. 1 of Ref. 1. A second compilation of best flight control design practices[2] was prepared under the auspices of the NATO AGARD Flight Vehicle Integration Panel (now the NATO Research and Technology Organization System Concepts and Integration Panel). In summarizing the lessons learned from much of the fly-by-wire aircraft experience, this excellent document also emphasizes the key role of handling qualities and robust stability criteria from the start and throughout the development and flight-test process. Automated optimization methods offer the potential of considerable time and cost savings in engineering over the manual evaluation and optimization of control laws to meet the many and competing design requirments that ensure good handling qualities and that limit the potential for PIO. Optimization methods are epsecailly valuable as an efficient means to update repeatedly control system parameters as accurate aircraft dynamic response data become available during flight testing.

Design requirements for satisfactory aircraft handling qualities and robust closed-loop stability have been developed and refined based on extensive flight-test, piloted-simulation, and analytical studies. Recommended values of dynamic response metrics, for example, bandwidth, rise time, phase lag, etc., and detailed supporting explanation are compiled in specification documents for fixed-wing handling qualities (MIL-STD 1797A), rotorcraft handling qualities (Design Standard ADS-33E), general flight-control stability requirements (MIL-SPEC 9490), and much additional excellent design guidance.[2] A complete set of design criteria may include as many as 50–100 individual specifications, including metrics based on time response, stability margin, and disturbance response.[3]

Many innovative flight-control design methods have been and continue to be proposed and advocated in a desire to improve flight-control system performance/robustness and to automate and thereby reduce the cost of the flight-control development process. Each flight-control design method invariably has a set of tuning parameters, for example, Q and R matrices for LQR design, weighting functions for H-infinity design, or target eigenspace locations for eigenvector assignment design. It is common for these parameters to be selected based on only one or two key design requirements, such as control system bandwidth or rise time. A complete evaluation against the full set of dynamic response criteria is then conducted, if at all, as a final check of the completed design.

In most of the proposed methods, there is little transparent connectivity between the tuning parameters and the ultimate dynamic response requirements. Furthermore, there is rarely an attempt to optimize the tuning parameters of a proposed method to achieve the complete set of dynamic requirements without overdesign and the resultant excessive actuator activity. Because prospective alternative design methods are not optimized against a common set of dynamics response objectives, it is difficult to compare their performance fairly. Some useful explanations and comparisons of alternative design methods are presented in Refs. 4 and 5.

The Control Designer's Unified Interface (CONDUIT®),[3] used in the current study, addresses the issues just discussed by facilitating the automated evaluation and optimization of any control law design method or architecture against a common comprehensive set of flight-control design requirements. As improved modeling data or design requirements become available during the development process, control laws can be updated rapidly and problems that arise during flight tests can be resolved quickly. Key features of CONDUIT include the following: 1) extensive precoded graphical libraries of key handling qualities and dynamic response criteria design specifications; 2) integrated environment to create, validate, and catalog new user-defined specifications; 3) easy graphical selection, setup, and evaluation of specifications; 4) accommodation for any design method or architecture; 5) automated tuning of user-selected design parameters using a vector optimization method that ensures that each dynamic response criteria is individually met, rather than only a weighted average being met; 6) ultimate design optimization to minimize the selected performance objective that meets the design criteria with minimum overdesign; and 7) extensive supporting plots and integrated analysis tools. Furthermore, the optimization results enable the proponent of a new design method to reverse engineer the selection of tuning parameters that will yield a design that best meets the criteria and thereby check the proposed rules of thumb. Details of the CONDUIT design environment may be found in Ref. 3.

This paper demonstrates the optimization and comparison of several alternative control design methods against a common set of dynamic response criteria. The case study performed is for the design of lateral/directional control laws for a transport aircraft based on KC-135 aircraft linearized dynamics. The key nonlinearities associated with the actuator position and rate limiting are included in the model. The design methods considered are classical, linear quadratic regulator (LQR), dynamic inverse, and H-infinity. The results show that each method can be tuned using CONDUIT to meet the dynamic requirements necessary to ensure good handling qualities and stability, while minimizing overuse of the actuators. An important result is that the alternative design methods optimized against a common set of requirements yield controllers with comparable performance and robustness characteristics.

Lateral/Directional Aircraft Flight-Control Design Problem

The case study problem is based on the lateral/directional dynamics of the KC-135 aircraft as given by Blakelock[6] and is shown in Fig. 1. The control system topology was selected for this case study is notional and does not represent the actual implementation in the KC-135 aircraft. A comprehensive design model would also inlude many additional elements, such as digital sampling effects, structural dynamics and related notch filters, sensor dynamics, mechanical hysteresis, and stick shaping.

The key elements of the block diagram are 1) three-degree-of-freedom (DOF) state-space representation of KC-135 lateral/directional dynamics, 2) second-order actuator dynamics including rate and position limiting, 3) angular rate gyro filters, 4) washed-out yaw rate feedback (to enable turn coordination), 5) roll rate and yaw command, and 6) disturbance inputs. The feedback loop architecture shown in Fig. 1 is (one of several) that would be appropriate to a classical design method.

The design objectives shown in Fig. 2 are selected from the CONDUIT libraries as appropriate to fixed-wing transport aircraft (1797A, 9490). The relative priority of each specification is designated by the user as hard specification (H), soft specification (S), or summed objective (J) indicated in the upper right-hand corner of the specification. The role of the specification priority in the phases of CONDUIT optimization process is summarized later, but is described more fully in Ref. 3.

In phase 1, the design parameters are tuned to attempt to meet the H specifications selected for this design study: 1) All closed-loop eigenvalues must lie in the left-half plane (absolute stability) (EigLcG1). 2) There are gain/phase margin requirements (MIL-SPEC 9490) to ensure satisfactory relative stability/robustness (StbMgG1).

If phase 1 specifications are met, the solution enters phase 2, in which the design parameters are tuned to attempt to meet all of the S specifications and performance metrics included in the J objective. The S specifications are the handling-qualities metrics and can comprise as many as 50–100 individual requirements for a full-scale design problem. In the current study, the S specifications are from MIL-STD 1797A: 1) roll command response bandwidth (BnwRoD1), 2) equivalent system Dutch roll mode damping (DmpDrD3) and frequency (FrqDrD4), and 3) roll time response quickness (QikAtG1). Additional soft specs included from the CONDUIT libraries are 1) actuator saturation limits (SatAcG1) and 2) disturbance response (HldNmH1).

Finally, if phase 2 specifications are met, the solution enters phase 3. Here, the design parameters are tuned to attempt to

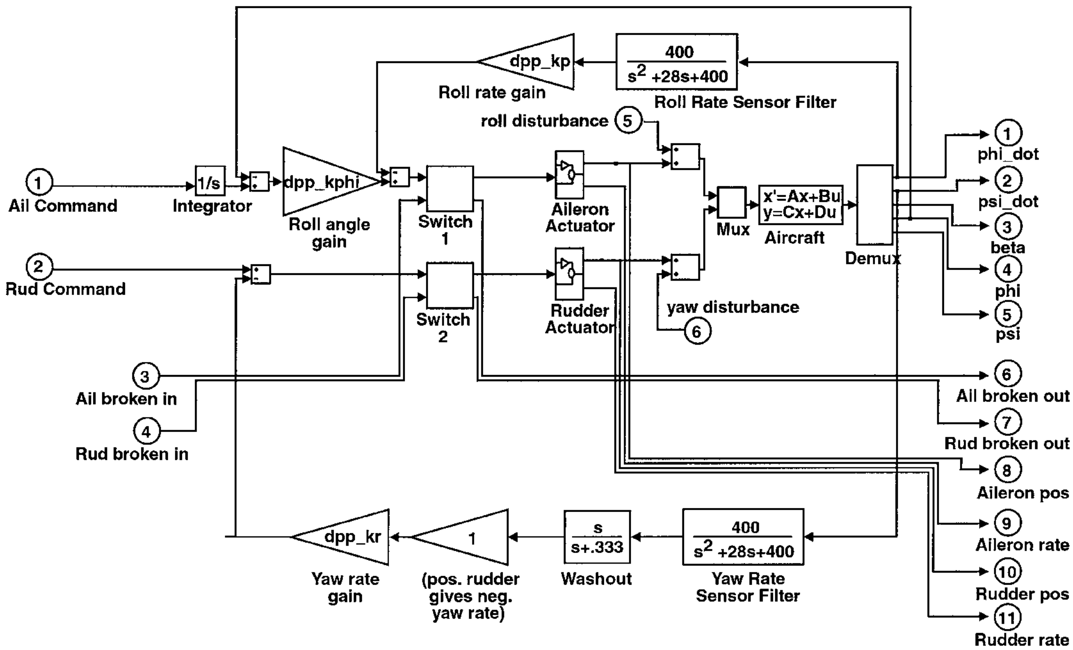

Fig. 1 Case study problem based on KC-135 lateral/directional dynamics.

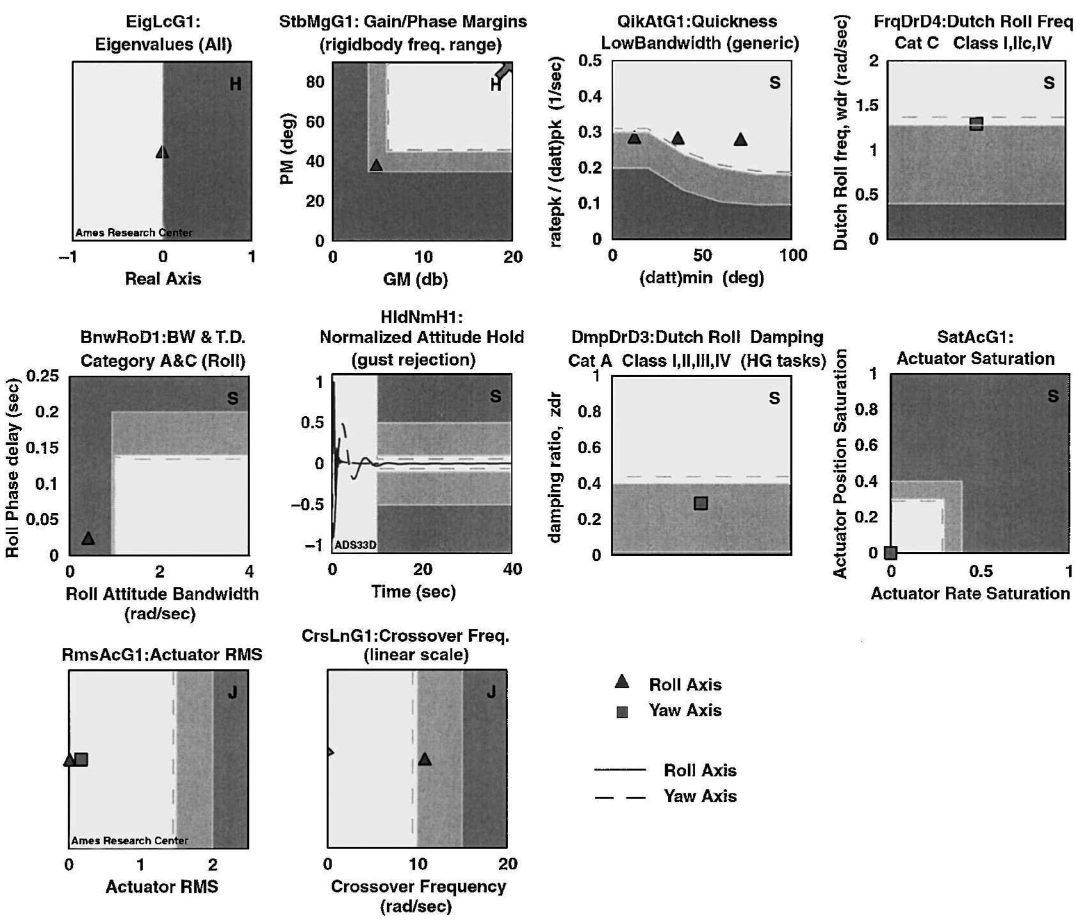

Fig. 2 Initial design for classical method.

minimize the J summed objective specification, with the earlier H and S specifications enforced as constraints in the optimization. Extensive practical design experience using CONDUIT has indicated that good performance metrics to be included in the J summed objective are 1) crossover frequencies for the individual broken loops (CrsLnG1) and 2) actuator position response as characterized by the root-mean square (rms) (RmsAcG1).

The rms value for each actuator is determined from a power spectrum calculation in the frequency domain, based on the transfer function from pilot stick input to actuator position response. The rms value is normalized to maximum control position, so that a value of rms = 1.0 indicates that full stick input will command full control surface authority (100% saturation).

By selecting these performance metrics, we ensure that the final design will meet all of the requirement specifications with minimum overdesign. This approach minimizes 1) sensitivity to sensor noise and unmodeled high frequency dynamics, 2) structural fatigue, and 3) actuator limiting.

The use of a summed objective for phase 3, in contrast to the min/max vector optimization of phases 1 and 2, allows the optimization to explore a broad space of possibilities that improve the ultimate performance of the individual loops while still maintaining compliance with the level 1 requirements. For example, the optimization will continue to improve the roll axis metrics (roll loop crossover frequency and aileron rms), even after the yaw axis metrics have reached optimum (but higher values) of these metrics. This allows all of the loops to be tuned to their ultimate performance. Detailed explanation of this strategy is given in Ref. 3.

The complete set of design specs are shown in Fig. 2 with the associated H, S, and J priority designations in the upper right of each specification. The lightest shade region reflects level 1 handling-qualities ratings, corresponding to characteristics that are satisfactory without improvement. The medium shade reflects level 2 handling-qualities ratings and corresponds to characteristics with deficiencies that warrant improvement. Finally, the darkest shade reflects level 3 handling-qualities ratings, corresponding to characteristics with deficiencies that require improvement. These boundaries are given by the various design standard (MIL-SPEC) documents (e.g., Ref. 2) based on extensive piloted handling-qualities data.

A desired amount of overdesign is selected (design margin) to provide uncertainty robustness, by ensuring that the desired specification point lies a safe distance within the level 1 region and not right on the level 1/level 2 boundary. As shown in Ref. 3, this parameter provides a direct mechanism for evaluating the tradeoff between improved performance and increased control usage. For the current study, the overdesign parameter was selected as 10% and is indicated by the additional dashed boundary line within the level 1 region on each specification in Fig. 2.

Alternative Design Methods

This paper presents and compares KC-135 case study results for four alternative design methods listed in Table 1: 1) classical, 2) LQR, 3) dynamic inverse, and 4) H-infinity. The alternative methods can be characterized by the number of design parameters and the number of controller DOF. The design parameters are any quantities, for example, gains, filter time constants, actuator limits, that CONDUIT is free to manipulate to arrive at an optimized solution to the control system design problem. In this case study, the classical method was implemented with three design parameters (feedback gains), whereas the remaining methods were implemented with four design parameters.

The design methods can additionally be characterized by the number of controller DOF, one or two. A one-DOF controller refers to a response feedback (RF) architecture in which the single compensator must be tuned to meet both performance and robustness requirements and, therefore, limits the design flexibility. A two-DOF controller architecture includes both an RF DOF and a command model (CM) (or forward-loop) compensator. Such an architecture allows a separate optimization of the regulator and the performance characteristics. In this study, the control law architectures for the classical, LQR, and H-infinity methods were configured with one DOF (RF), whereas the dynamic inverse method is inherently two DOF (RF and CM). It is important to recognize that a two-DOF architecture could be implemented in conjunction with any/all of the alternative design methods listed earlier. Furthermore, there are numerous combinations and permutations of the four methods that have been proposed in the literature. The objective of this case study is to show the feasibility and desirability of evaluating and systematically optimizing any selected method against a common wide-ranging set of design criteria.

Classical Design

The classical control laws for the lateral/directional stabilization and command response are shown in Fig. 1. The classical architecture design is employed with feedback of roll rate *Kp* and roll angle *K*phi to aileron to achieve necessary stability margins and closed-loop control bandwidth. The washed-out yaw rate feedback to rudder *Kr* achieves the specified minimum Dutch roll damping and frequency. The feedbacks are also needed to reject atmospheric disturbances. The choice of washed-out yaw rate feedback has been

Table 1 Alternative design methods

Method	DOF	Tuning parameters (DP)	Number of DPs
Classical	1 (RF)	Feedback gains	3
LQR	1 (RF)	Q matrix	4
H infinity	1 (RF)	Loop shaping filters, $K/(s+a)$	4
Dynamic inverse	2 (CM, RF)	CM frequencies RF frequencies	4

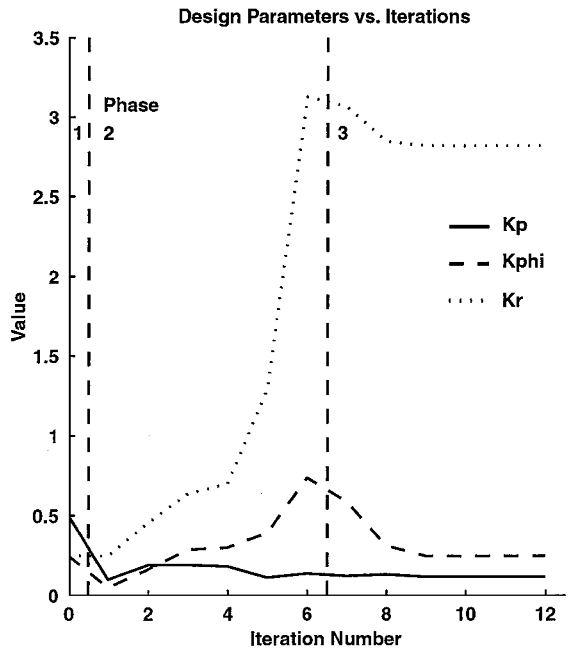

Fig. 3 Classical design optimization.

adopted herein to maintain consistency with Blakelock,[6] rather the the more currently used pseudo-$\dot{\beta}$ feedback in modern aircraft. The three feedback gains, *Kp*, *K*phi, and *Kr* are designated as CONDUIT tuning parameters via the positive definite design parameter (dpp_) in prefix seen in the block diagram of Fig. 1, and comprise a one-DOF (RF) architecture as indicated in Table 1.

The initial values for the classical design parameters were selected as

$$Kp = 0.5 \text{ deg/deg/s}, \qquad K\text{phi} = 0.25 \text{ deg/deg}$$

$$Kr = 0.25 \text{ deg/deg/s} \tag{1}$$

These values (and the initial values for the other design methods in this study) were intentionally selected to yield poor dynamic behavior (as shown in Fig. 2) and highlight the robustness of the CONDUIT optimization to poor initial design guesses. As seen in Fig. 2, the initial design is in the level 2 (deficiencies warrant improvement) region for many of the design requirements.

The iteration history plot (Fig. 3) shows that after six iterations CONDUIT reached phase 3, which is a feasible solution where all specifications are in the level 1 region. The fully converged result is shown in Fig. 4. The optimized solution achieved after 12 iterations meets all H and S requirements while minimizing the J performance metrics. Convergence of the design in phase 3 toward the final optimized solution is quite smooth, as seen in Fig. 3, and is given by

$$Kp = 0.12 \text{ deg/deg/s}, \qquad K\text{phi} = 0.24 \text{ deg/deg}$$

$$Kr = 2.82 \text{ deg/deg/s} \tag{2}$$

The individual performance metrics (crossover frequency and rms for both loops) that comprise the J summed objective are listed in Table 2.

A key characteristic of the converged phase 3 design in CONDUIT, as seen in Fig. 4, is that one (or more) of the S or H requirements lies on the design margin (dotted) boundary. This indicates that the level 1 design requirements have been met with minimum overdesign because further reduction in control usage (lower crossover frequency and/or lower control rms) will result in one (or more) requirements penetrating into the level 2 region. Important aspects of the classical design achieved with CONDUIT are 1) control law requires measurements of p, ϕ, and r; 2) all level 1 design requirements achieved (with 10% design margin); 3) no control saturation for the largest expected command inputs; 4) reasonable crossover frequencies (3–4 rad/s); and 5) overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements, for example, quickness and bandwidth, for the one-DOF architecture.

Each specification in CONDUIT has an associated set of supporting plots that illustrates the various related analyses and calculations. For example, the supporting plots for the stability margin specification StbMgG1 (Fig. 4) show the broken-loop response plots for the aileron (Fig. 5a) and rudder (Fig. 5b) loops that are used in the gain and phase margin calculations. Reference to Fig. 5a shows that the roll crossover frequency could be reduced considerably, while still maintaining adequate phase margin. However, the one-DOF architecture requires the higher gains to achieve the remaining handling-qualities requirements.

An interesting aspect of the results is the comparison of the optimized design parameters with rules of thumb for classical control.

Table 2 Comparison of performance metrics

Objective metric	Classical	LQR	Dynamic inverse	H infinity
ω_c, aileron, rad/s	2.89	3.10	0.344	4.12
ω_c, rudder, rad/s	4.24	1.92	1.772	3.96
rms, aileron, nondimensional	0.0118	0.0137	0.0276	0.0131
rms, rudder, nondimensional	0.1759	0.1702	0.1813	0.291

Design rules to achieving maximum crossover frequency and adequate stability margins for a given level of equivalent high-order loop delay are developed in Ref. 7. When the block diagram of Fig. 1 is referred to, the actuator and gyro filter dynamics contribute a total effective delay of $\tau_{\mathrm{eff}} = 0.092$ s.

The maximum achievable roll crossover frequency is then determined,

$$(\omega_c)_{\max} = 0.370/\tau_{\mathrm{eff}} = 4.02 \text{ rad/s} \tag{3}$$

which sets the roll angle feedback gain based on the roll control effectiveness

$$K_\phi = \frac{\omega_c^2}{2.48\left(L_{\delta_{\mathrm{ail}}}\right)} = 0.29 \text{ deg/deg} \tag{4}$$

The inverse lead time constant is also determined from the crossover frequency:

$$1/T_p = K_\phi/K_p = 0.442\omega_c = 1.78 \text{ rad/s} \tag{5}$$

which finally determines the roll rate gain,

$$K_p = 0.16 \text{ deg/deg/s} \tag{6}$$

Analogously, the yaw rate gain is easily determined from the maximum achievable yaw crossover frequency as

$$K_r = \omega_c \big/ N_{\delta_{\mathrm{rud}}} = 2.94 \text{ deg/deg/s} \tag{7}$$

Comparing these results with the optimized gains from CONDUIT [Eq. (2)] shows that the classical rules of thumb give good first guesses to the optimized solution.

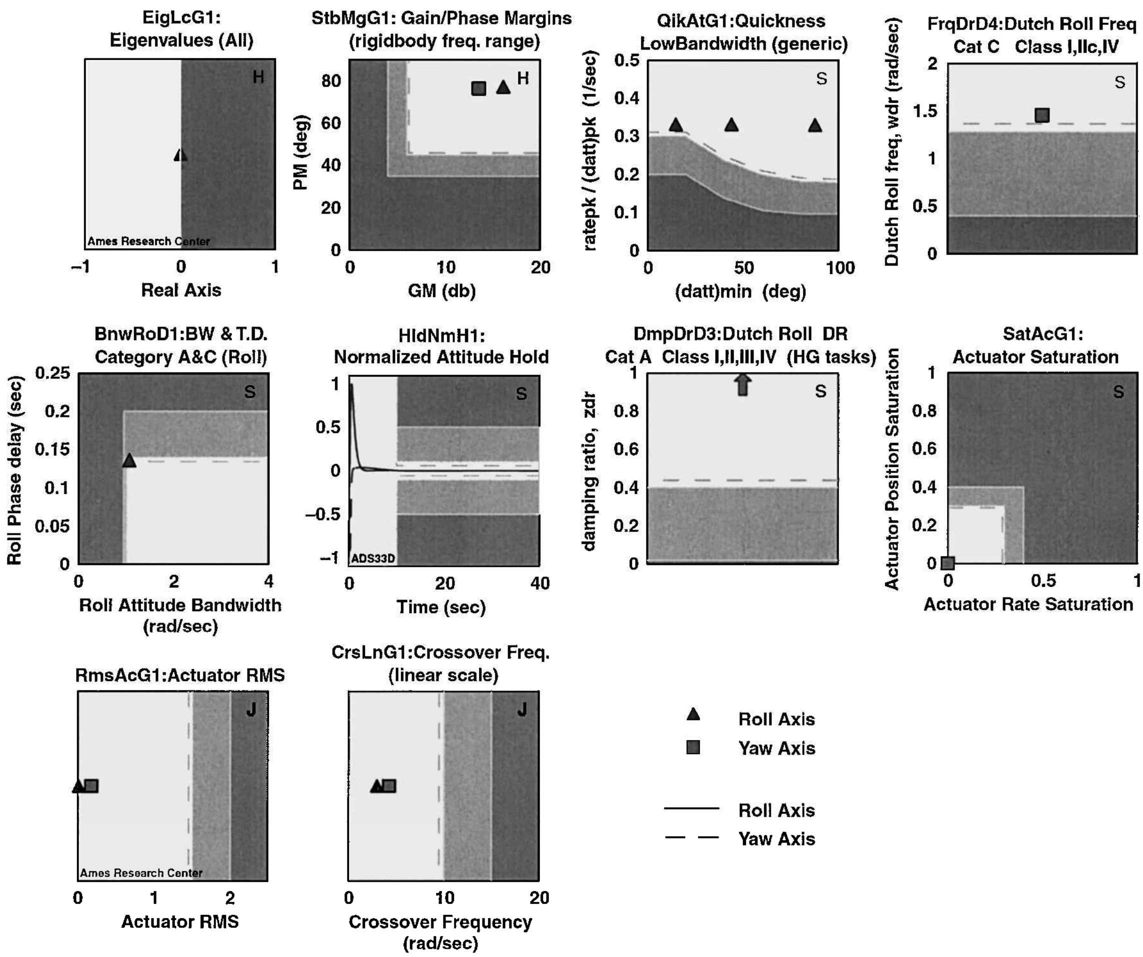

Fig. 4 Optimized classical design characteristics.

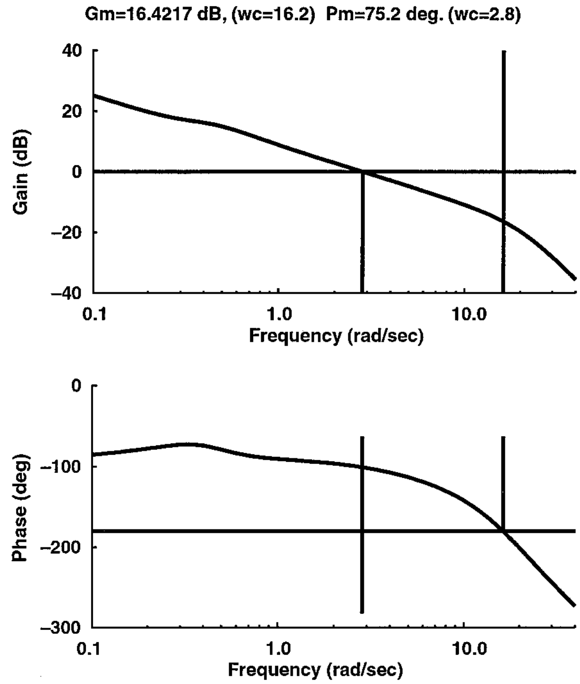

Fig. 5a Aileron broken-loop response (optimized).

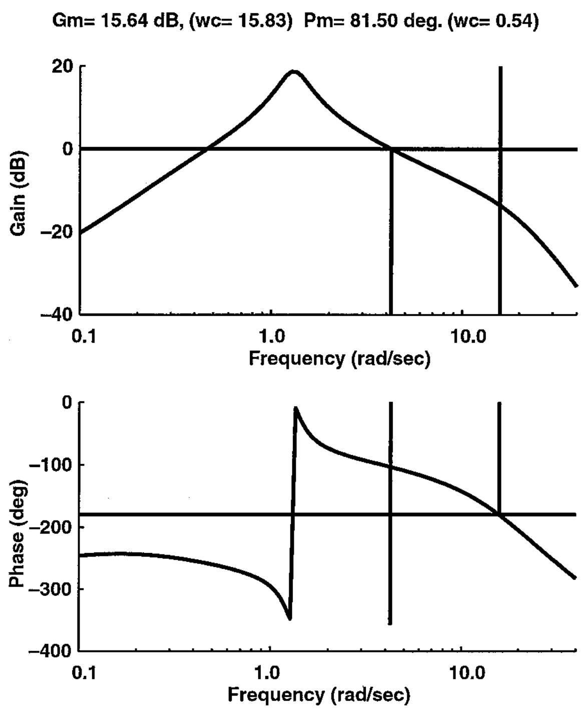

Fig. 5b Rudder broken-loop response (optimized).

LQR Design

The control law architecture for the LQR design is analogous to the classical design architecture of Fig. 1, but now with an LQR full-state feedback block replacing the individual feedback gains. The LQR method calculates the optimal gain matrix K for the (full) state feedback law that minimizes the quadratic cost function comprising state errors weighted by the Q matrix and control usage weighted by the R matrix.

The matrix Q is assumed to be a diagonal, so that the elements represent weights for the corresponding states. In practice, R is usually set to the identity matrix. When the elements of Q are selected to be small with respect to one, the penalty for the state errors is small with respect to the penalty for control usage, thus producing lower feedback gain, that is, lower crossover frequency and lower bandwidth, corresponding to a sluggish response. In the opposite case, when the elements of Q are selected as large relative to one, the controller will utilize a large amount of control to minimize state error, and thus, the response will be more brisk corresponding to a higher crossover frequency.

We adopt a common implementation of the LQR design method by including only the bare airframe dynamics (four states, p, r, β, and ϕ), in the Riccati solution. This ensures that we will only require these aircraft states as feedback measurements, and not all of the internal states of the control system. However, this approach also degrades the guaranteed stability of the LQR design because the omitted elements, for example, actuators, filters, etc., will contribute phase lag to the overall system.

CONDUIT is used to optimize the elements of Q as design parameters to achieve the level 1 requirements with a minimum overdriving of the control:

$$Q = \begin{bmatrix} Q_{pp} & & & \\ & Q_{rr} & & \\ & & Q_{\beta\beta} & \\ & & & Q_{\phi\phi} \end{bmatrix} \tag{8}$$

As before, we intentionally chose poor initial values for the design parameters ($Q = I$, the unity matrix) to show that CONDUIT is robust to initial design choices and that a good controller is quickly achieved.

A roll angle stick command gain and associated specification FrqGnG1 is included to ensure constant unity steady-state sensitivity in the roll axis, but this does not constitute an additional controller DOF. Therefore, as indicated in Table 1, the LQR implementation is one-DOF with four design parameters.

The handling-qualities specifications for the initial design parameter guesses ($Q = I$) are far from ideal, with level 3 behavior indicated for the eigenvalue and stability margin requirements and level 2 behavior indicated for several other requirements, thus presenting a poor initial design with which to challenge CONDUIT.

CONDUIT tuned the Q matrix to achieve an LQR design that satisfied all of the handling qualities, even when the initial guess for the design parameters was poor. The design parameters are fine tuned in phase 3 to a converged solution after 12 iterations that meets all level 1 design requirements (Fig. 6) at minimum crossover frequencies and actuator rms (Table 2). The optimized design parameters, corresponding to the diagonal elements of the Q matrix [Eq. (8)], are

$$Q_{pp} = 0.0079, \qquad Q_{rr} = 0.81$$
$$Q_{\beta\beta} = 0.95, \qquad Q_{\phi\phi} = 0.071 \tag{9}$$

The corresponding feedback gain matrix is

$$K = \begin{matrix} & p & r & \beta & \phi & \\ & \left[\begin{matrix} 0.1141 \\ 0.0088 \end{matrix}\right. & \begin{matrix} -0.2024 \\ -0.8456 \end{matrix} & \begin{matrix} 0.2262 \\ 0.2455 \end{matrix} & \left.\begin{matrix} 0.2823 \\ 0.0061 \end{matrix}\right] & \begin{matrix} \delta_{\text{ail}} \\ \delta_{\text{rud}} \end{matrix} \end{matrix} \tag{10}$$

The key characteristics of the optimized LQR solution are 1) control law requires measurements of p, r, β, and ϕ; 2) all level 1 design requirements are achieved (with 10% design margin); 3) no control saturation for the largest expected command inputs; 4) reasonable crossover frequencies (2–3 rad/s), and 5) overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements, for example, quickness and bandwidth, for the one-DOF architecture.

It is interesting to compare the optimized LQR design parameters with the rule of thumb guidance for Q matrix selection given by Bryson as described in Refs. 8 and 9. Bryson recommends setting the diagonal elements of Q to correspond to the square of the ratio of maximum control authority to maximum expected response, that is,

$$Q_{ii} = [(u_i)_{\max}/(x_i)_{\max}]^2 \tag{11}$$

When we assume a maximum control input of 5 deg (aileron and rudder) and maximum expected state responses [50 deg/s, 5 deg/s,

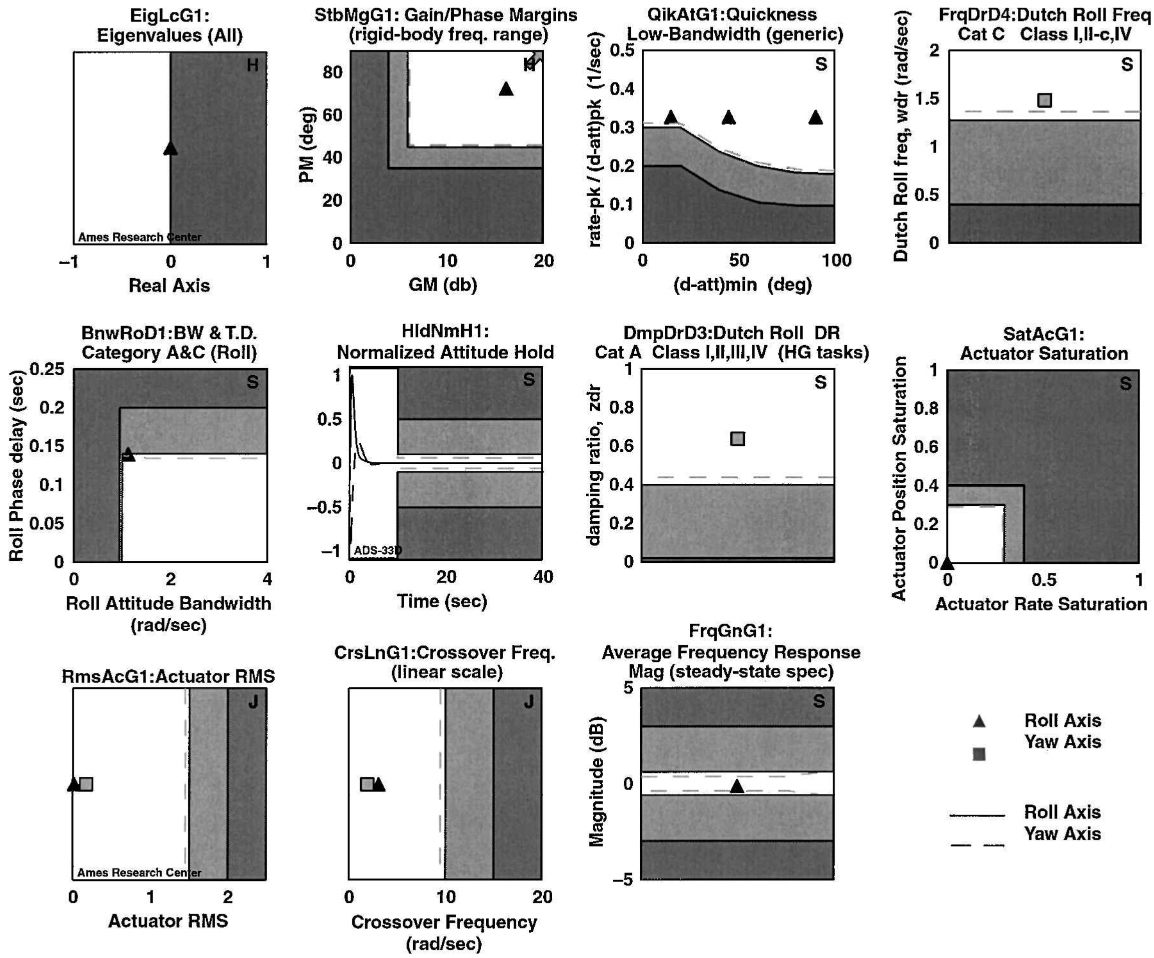

Fig. 6 Optimized LQR control system characteristics.

5 deg, 50 deg], corresponding to roll rate, yaw rate, sideslip, and roll angle, respectively, then we obtain

$$Q_{pp} = 0.01,\ Q_{rr} = 1.0,\ Q_{\beta\beta} = 1.0,\ Q_{\phi\phi} = 0.01 \qquad (12)$$

These initial guesses are reasonably close to the CONDUIT optimized values and, thus, support Bryson's rule of thumb.

Dynamic Inverse Design

The dynamic inverse control method[10] provides a very intuitive way of designing a control system by allowing a direct and independent selection of regulation loop bandwidth and command response bandwidth in a two-DOF architecture (Table 1). However, the method requires carrying a look-up table of accurate stability and control derivatives in the onboard software for the linear implementation or carrying a complete nonlinear simulation onboard in the nonlinear implementation. The linear implementation is adopted herein (Fig. 7) following closely the approach outlined by Franklin.[11] The control law architecture can conceptually be divided into three key elements: command model, regulator, and aircraft inversion.

The inverse control acts to cancel the basic aircraft dynamics to yield a simple ($1/s^2$) system. Then the regulator loops (RF) closed around these accelerationlike dynamics are easily tuned to produce the desired regulator bandwidth and stability margins. Finally, the CMs, contain first- or second-order transfer-function characterizations of the desired closed-loop aircraft response, tuned to meet the handling-qualities requirements. The control law implementation contains a dynamic algorithm that requires the measurement of rudder position δ_{rud} or lateral acceleration a_y to reconstruct the needed sideslip rate $\dot{\beta}$, as shown in Fig. 7.

Four design parameters were selected for CONDUIT optimization: inverse time constant, that is, bandwidth, for the first-order roll rate command model Tp, roll regulator natural frequency omphi, Dutch roll CM frequency comdr, yaw regulator natural frequency ombet. The roll and yaw regulation loop damping ratios were fixed at zeta = 0.7.

An arbitrary initial choice of design parameters was selected as

$$Tp = 1\ \text{s}, \qquad \text{comdr} = 1\ \text{rad/s}$$
$$\text{ombet} = 1\ \text{rad/s}, \qquad \text{omphi} = 1\ \text{rad/s} \qquad (13)$$

resulting in the initial control system performance that again is far from ideal. The roll bandwidth is in the level 3 region, and the Dutch roll frequency is in the level 2 region.

CONDUIT converges in 11 iterations to an optimized design that satisfies all level 1 requirements (Fig. 8) while minimizing overdesign (Table 2). The optimized design parameters are

$$Tp = 0.837\ \text{s}, \qquad \text{comdr} = 1.21\ \text{rad/s}$$
$$\text{ombet} = 0.76\ \text{rad/s}, \qquad \text{omphi} = 0.835\ \text{rad/s} \qquad (14)$$

The generation of high actuator rates and/or displacements has been reported in some previous applications of the dynamic inverse control method.[12] Although the aileron rms is substantially higher (a factor of two) than for the other methods (see Table 2), there is no rate or position saturation even for the largest commanded inputs. An important distinguishing characteristic of this design is the low roll crossover frequency ($\omega_{c_\phi} = 0.344$ rad/s), and reduced (but still level 1) phase margins for both loops (about 45 deg) as compared to the other methods. These features arise as a result of the two-DOF controller architecture, which allows independent tuning of the feedback response, that is, crossover characteristics, from the command response.

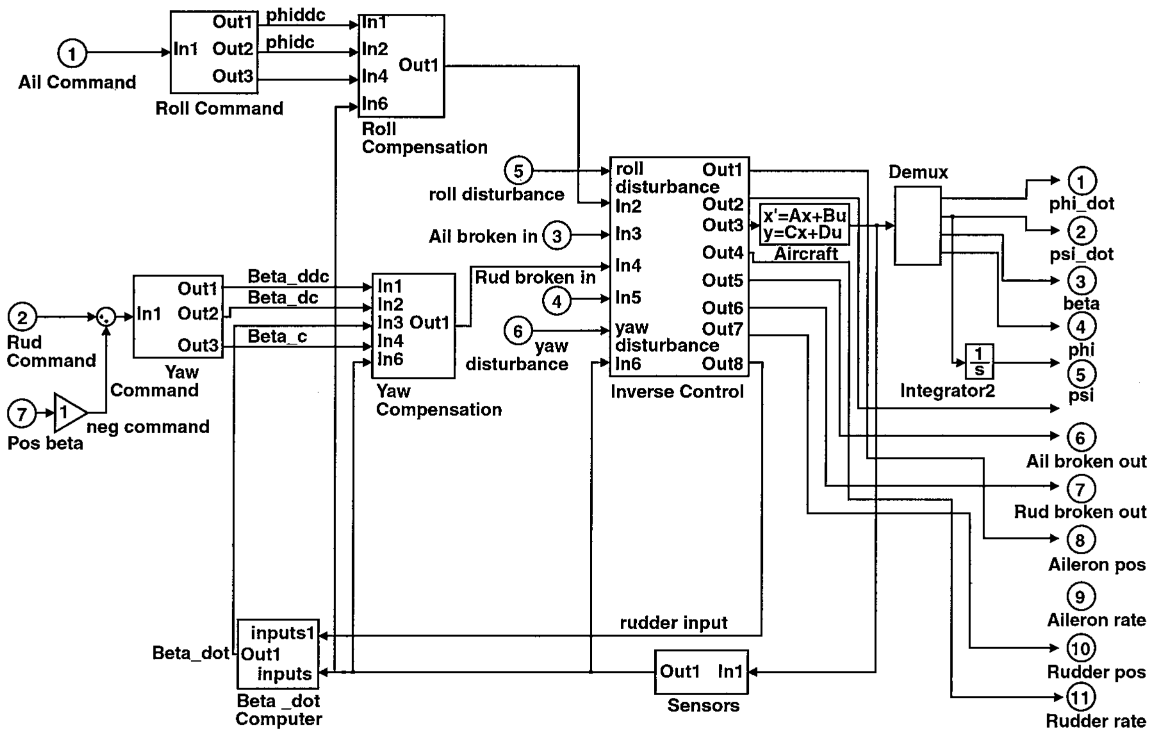

Fig. 7 Dynamic inverse control architecture.

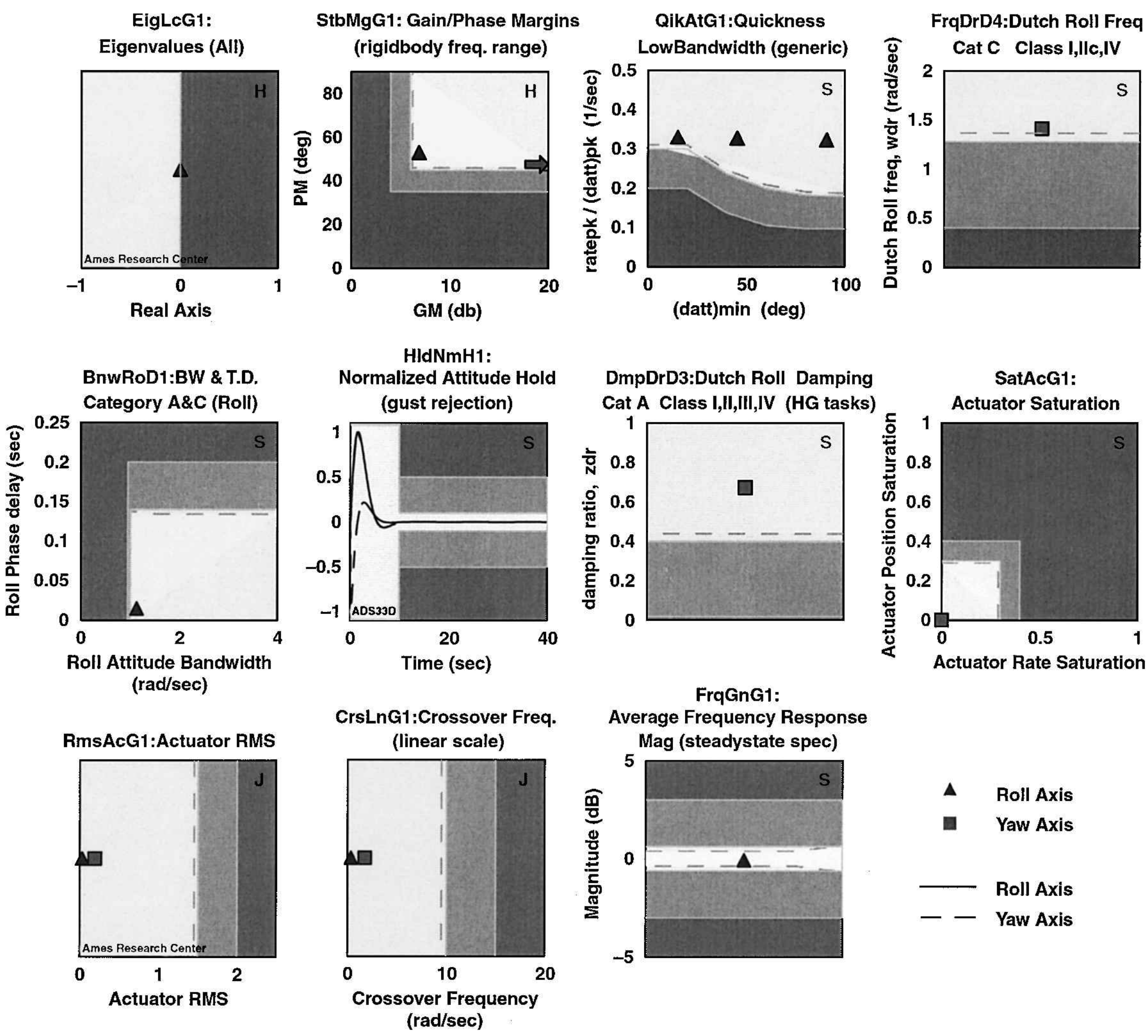

Fig. 8 Optimized dynamic inverse control system characteristics.

In the roll channel, satisfactory stability is achievable at low crossover frequencies, as can be seen in Fig. 5a. Therefore for this (two-DOF) design, CONDUIT reduces the crossover frequency as an element of the summed objective until the attitude hold specification HldNmH1 is just met. Tighter attitude hold or gust rejection requirements would drive the roll crossover frequency higher. As long as all of the level 1 design characteristics are preserved, the reduced crossover frequency decreases sensitivity to measurement noise, relaxes actuator hardware requirements, reduces actuator saturation, reduces structural fatigue, and reduces the possible excitation of unmodeled high-frequency modes. Note that these same characteristics would be achieved with any of the other methods with the introduction of a second controller DOF.

The results of the dynamic inverse design method are 1) control law requires measurements of p, r, β, ϕ, and δ_{rud} or a_y; 2) control law requires dynamic estimate of $\dot{\beta}$; 3) level 1 design requirements achieved (with 10% design margin); 4) increased aileron control rms; 5) no control saturation for the largest expected command inputs; 6) low roll loop crossover frequency; and 7) no overdesign in stability margins or disturbance rejection.

H-Infinity Design

The H-infinity control method implemented in this case study used the one-DOF loop-shaping approach described by Postlethwaite.[13] In this method, two diagonal weighting matrices $W1$ and $W2$ are introduced to shape the bare-airframe response. Then, a high-order controller is synthesized to minimize the H-infinity norm of the shaped aircraft dynamics system. Finally, the controller order is reduced to match that of the aircraft dynamic system. In an approach analogous to the R matrix in the LQR method, the $W2$ weighting matrix is set to identity for simplicity. The remaining weighting function $W1$ comprises first-order filters for the aileron loop

$$W1_{\mathrm{ail}} = k_{\mathrm{ail}}/(s + a_{\mathrm{ail}}) \tag{15a}$$

and for the rudder loop

$$W1_{\mathrm{rud}} = k_{\mathrm{rud}}/(s + a_{\mathrm{rud}}) \tag{15b}$$

where the filter parameters k_{ail}, a_{ail}, k_{rud}, and a_{rud} are used as design parameters in CONDUIT.

The complete H-infinity control system architecture is shown in Fig. 9 and is listed as one DOF with four design parameters (DPs) as indicated in Table 1. As before, an additional (but dependent) parameter was included to maintain unity steady-state stick sensitivity in the roll axis.

The initial choice of H-infinity design parameters yields a range of level 1, 2, and 3 characteristics:

$$k_{\mathrm{ail}} = 6, \qquad a_{\mathrm{ail}} = 50, \qquad k_{\mathrm{rud}} = 160, \qquad a_{\mathrm{rud}} = 50 \tag{16}$$

CONDUIT optimized the design parameters to satisfy level 1 requirements (Fig. 10), with some overdesign in several specs due to the one-DOF nature of the controller:

$$k_{\mathrm{ail}} = 17.11, \qquad a_{\mathrm{ail}} = 48.63$$
$$k_{\mathrm{rud}} = 159.8, \qquad a_{\mathrm{rud}} = 50.58 \tag{17}$$

which shows a significant change only the first DP. As might be expected, the optimized rolloff in the loop transmission corresponds to just below the actuator bandwidth (62.8 rad/s), to keep from overdriving the control inputs.

The ability to achieved an optimized solution for the H-infinity design was quite sensitive to the initial values for the design parameters. If the design was started too far from a feasible solution, the optimization tended to wander and often did not reach a satisfactory result. This problem was not encountered for the other design methods, which achieved a unique converged solution independent of the initial guesses for the design parameters. The design parameters for the chosen H-infinity implementation (first-order filter coefficients in $W1$) did not appear to map well into the design requirements, which in turn caused problems for the optimization. It is possible that a more general implementation of the H-infinity design method would resolve this problem. The performance metrics for this design are listed in Table 2.

The results of the H-infinity design method are 1) control law requires full state measurements of p, r, β, and ϕ; 2) sixth-order dynamic compensator (fourth-order controller plus $W1$ filters); 3) all level 1 design requirements achieved (with 10% design margin); 4) increased rudder control rms; 5) no control saturation for the largest expected command inputs; 6) reasonable crossover frequencies (4 rad/s); 7) overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements, for example, quickness and bandwidth, for the one-DOF architecture; and 8) optimization quite sensitive to initial choices of the design parameters.

Comparison of Stability Robustness for Alternative Design Methods

An analysis of the stability robustness to uncertainties in the aircraft parameters was conducted for the alternative design methods. The aircraft state-space model matrices (A and B) were perturbed element by element, singly and in combination, by plus or minus 20%, with the optimized DPs fixed as determined earlier [Eqs. (2),

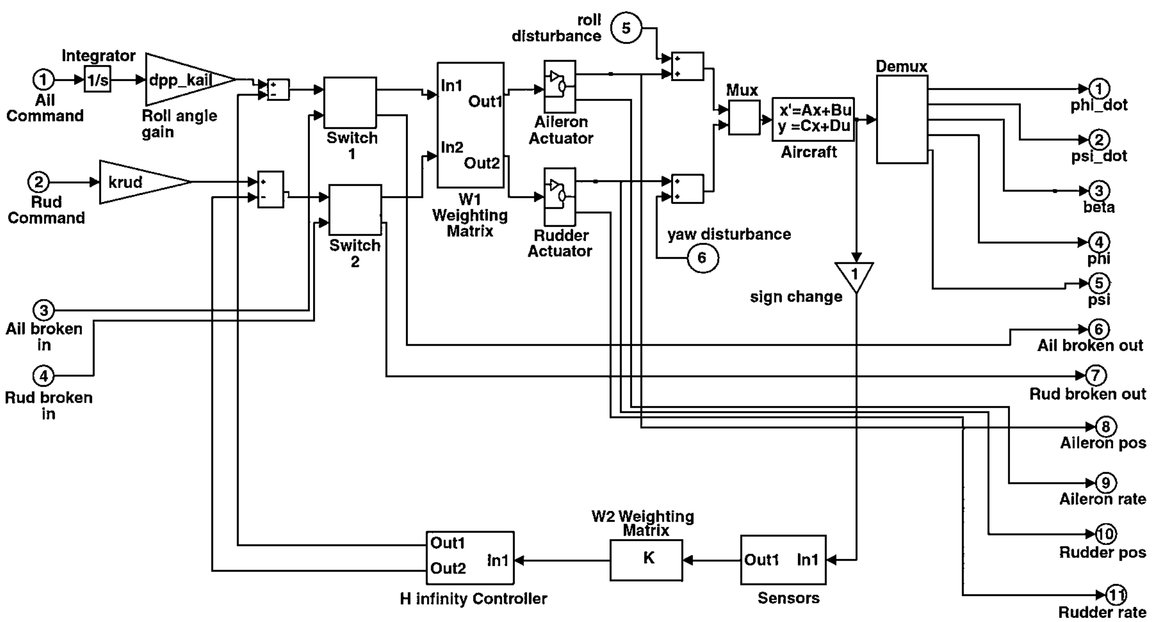

Fig. 9 *H*-infinity control system architecture.

(9), (14), and (17)]. Figure 11 shows the robustness analysis for the four designs on the stability margin boundary.

The nominal results (without perturbation) are shown for reference and correspond to the earlier CONDUIT optimization results. The one-DOF design methods (classical, LQR, and H-infinity) all had large nominal values for stability margins, and therefore, perturbed cases do not penetrate the level 2 region. In the case of the two-DOF dynamic inverse design, the nominal margins are somewhat reduced, so that the perturbated cases penetrate into the level 2 region. This could be easily alleviated by tightening the level 1 stability margin boundaries used for the nominal design, or by increasing the overall design margin. The results of Fig. 11 also reveal that the classical and H-infinity designs display a very small scatter between all of the aileron and rudder margin cases, as compared to the LQR and dynamic inverse designs.

Table 3 compares the relative robustness for the four designs, as measured by the maximum deviation in gain and phase margin from the nominal point (data from Fig. 11). The classical and H-infinity methods display improved robustness (reduced scatter) for the rudder loop compared to LQR and dynamic inverse, but overall the differences seem much less pronounced than one might expect at the outset considering the large differences in design methods. In addition to checking the sensitivity to parametric variations around a single reference condition, a complete robustness evaluation should also include aircraft configuration and flight condition changes that are not explicitly included in a gain schedule implementation.

Table 3 Relative robustness for 20% variations in aircraft dynamics model (maximum deviation from nominal value shown)

Objective metric	Classical	LQR	Dynamic inverse	H infinity
Aileron gain margin, dB	−1, +2	−2, +2	−2, +3	−2, +3
Aileron phase margin, deg	−9, +9	−10, +9	−8, +9	−7, +8
Rudder gain margin, dB	−2, +2	−4, +2	−2, +3	−2, +3
Rudder phase margin, deg	−6, +3	−22, +19	−11, +14	−7, +9

Discussion

The key results of this comparative design study are summarized in Tables 1–3. The design methods were intentionally setup for a comparable number of design parameters (3–4) and a single-degree-of-controller freedom (except for the dynamic inverse). CONDUIT readily tuned all four design methods via the selected tuning parameters to meet the level 1 standards and minimize overdesign.

As seen in Table 2, the three methods that use a one-DOF architecture (classical, LQR, and H-infinity) achieve comparable values of crossover frequency as needed to meet both the disturbance response and command response specifications. The addition of a second DOF in the dynamic inverse method relaxes the constraint between the feedback and command response characteristics as imposed by the one-DOF architecture. This permits the significant reduction in aileron loop crossover frequency until the disturbance response just reaches the level 1/level 2 boundary (Fig. 8). These characteristics are illustrated well in Fig. 12, which is a comparison of the aileron broken-loop responses (aileron broken out/aileron broken in) for the optimized designs. The classical and LQR designs are seen to have an identical K/s loop shape as needed to achieve the common feedback requirements, for example, stability margin and disturbance. The H-infinity loop shape is essentially the same but with a slightly higher loop gain, resulting in the increased crossover frequency (Table 2). Also, the H-infinity shaping filters $W1$ cause the additional phase rolloff at higher frequencies (beyond 20 rad/s). Whereas the dynamic inverse loop shape is also K/s and parallel to the other designs, the magnitude curve is shifted down by −20 dB. This results

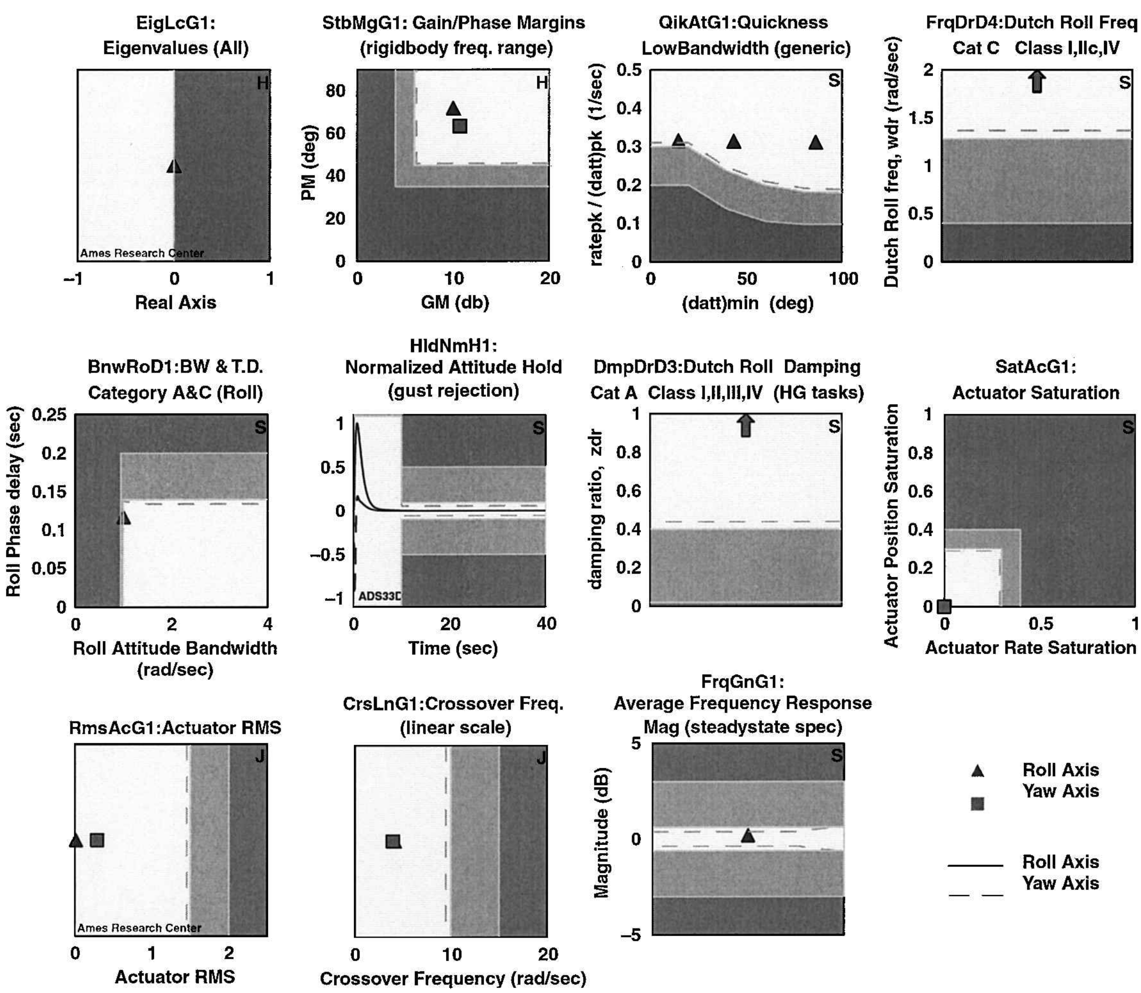

Fig. 10 Optimized H-infinity control system characteristics.

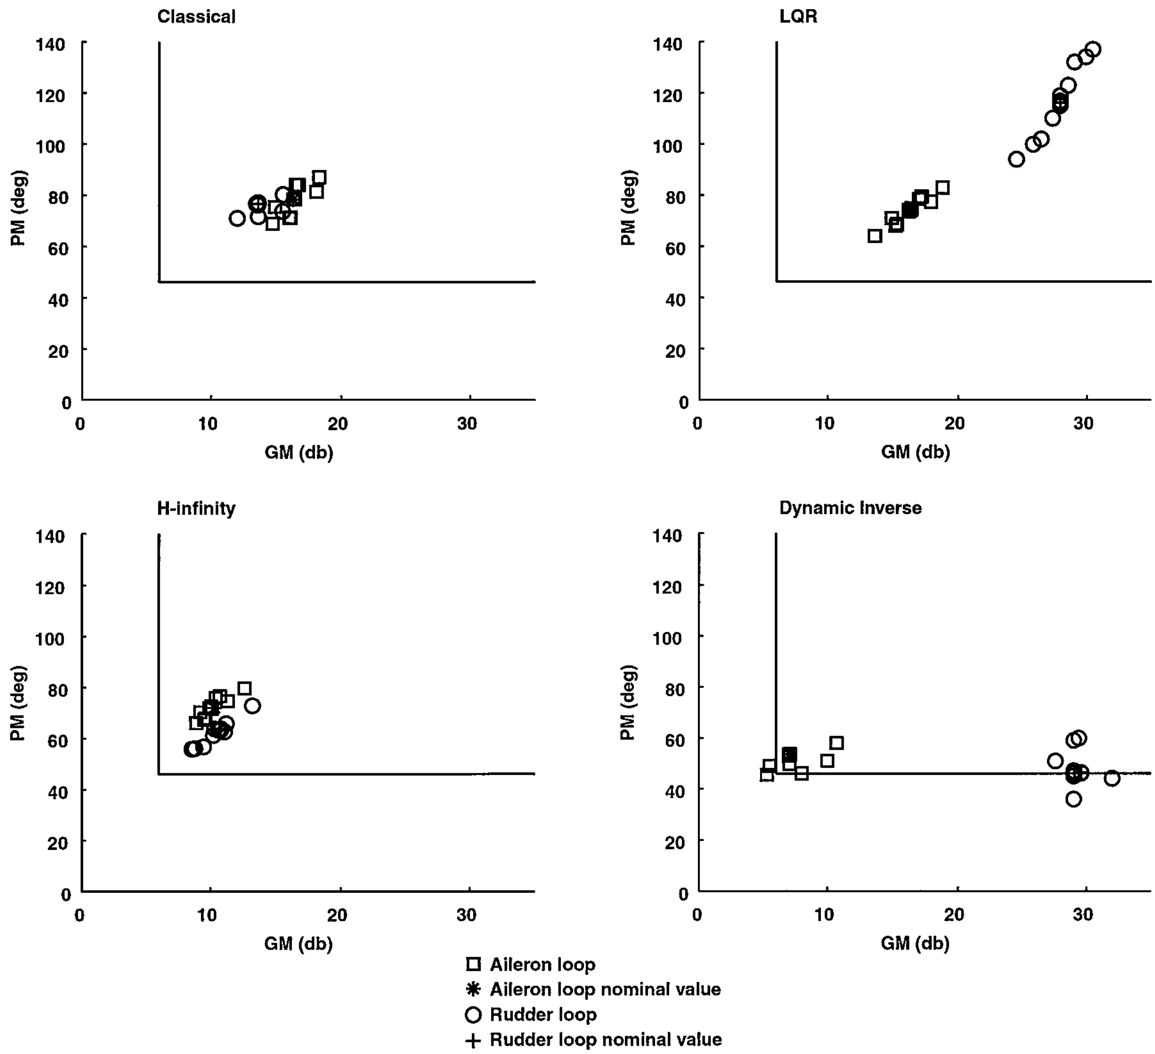

Fig. 11 Robustness comparison for 20% plant variations.

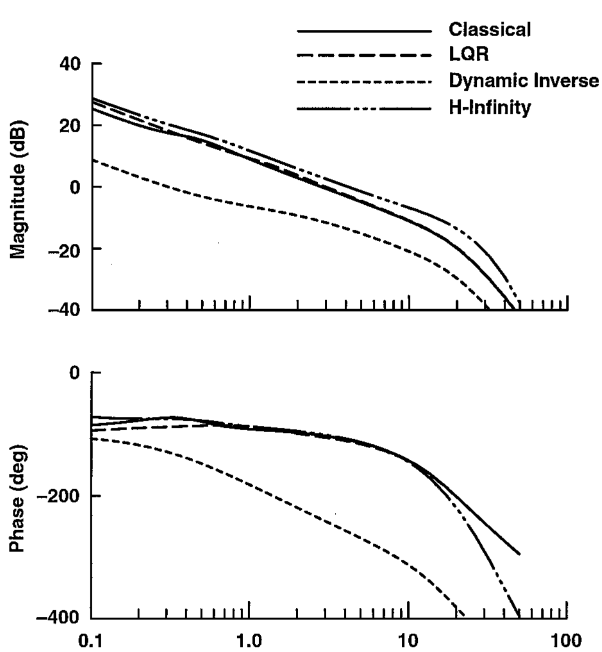

Fig. 12 Aileron broken-loop response comparison.

in a crossover frequency that is about one-tenth of the other designs, but still meets the disturbance response requirement with adequate stability margin. The stability robustness associated with the feedback response is somewhat better for the classical and H-infinity design methods, but still quite comparable across the four methods.

The responses to roll command (phi_dot/Ail command) for closed-loop roll rate are compared for the four optimized designs in Fig. 13. Once again, we see that the responses for the classical and LQR designs track very closely to meet the common command response requirements, for example, bandwidth and quickness, which explains the nearly identical values for aileron rms. The H-infinity design exhibits a slightly higher closed-loop bandwidth consistent with the increased crossover frequency, but displays a comparable gain rolloff at higher frequencies resulting in an identical value of aileron rms. Finally, the dynamic inverse response tracks the other designs at low frequencies (up to 2 rad/s) to achieve common closed-loop bandwidth and quickness requirements, but then follows the first-order roll-rate command model at higher frequencies. The increased response gain at higher frequencies accounts for the sizeable increase in aileron rms for the dynamic inverse design. The increased rudder rms (and increased aileron crossover frequency) for the H-infinity method reflects some overdesign that appears to results from the lack of direct mapping between the H-infinity design parameters ($W1$ shaping filters) and the handling-qualities requirements. This may also be evidenced in the sensitivity of the H-infinity convergence to initial design parameter values.

The overall results of this study suggest the careful definition of design requirements sets a unique broken-loop shape and CM response. CONDUIT will alternatively drive on the respective DP for

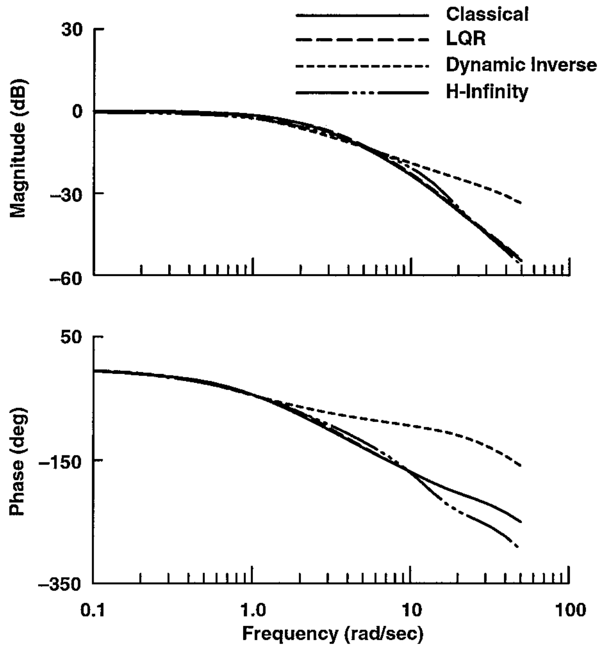

Fig. 13 Aileron closed-loop response comparison.

classical (feedback gains), LQR (Q matrix), etc., using the associated mathematical mechanisms (direct gains tuning, Riccati equations, etc.) to achieve the common required loop shape and, thus, the common set of handling-qualities requirements. Also, the optimized controller performance and robustness do not vary much between the different design methods. Perhaps achieving such similar results using different but all linear controllers might have been expected, but this is now made clearly apparent with the opportunity to optimize the methods on a level playing field using CONDUIT. Thus, the results here lend further support to the view that the mathematical mechanism to achieve the level 1 requirements is not nearly as important as the 1) proper selection of and optimization against a comprehensive set of specifications, 2) selection of a one- vs two-DOF controller architecture, 3) flight measurements required for each method, and 4) degree of transparency in the design method for efficiently resolving problems that may arise in flight test.

As stated at the outset, the objective of this research work was not to demonstrate the superiority of one method of control system design over another. Moreover, an expert in any one of the methods used may come up with a more effective implementation or a better design based on a combination of the methods. The intent herein was to demonstrate the feasibility of using CONDUIT to analyze and optimize a range of familiar control design methods against a broad-spectrum set of design requirements.

Conclusions

1) Proposed alternative design methods can and should be evaluated against a consistent set of design specifications that include all relevant handling-qualities and controller performance metrics.

2) Four alternative design methods (classical, LQR, dynamic inverse, and H-infinity) optimized against a common and comprehensive set of requirements achieved similar performance and robustness.

3) CONDUIT served as an efficient and effective tool for evaluating and optimizing alternative control system designs. Level 1 performance specifications were produced after a small number of iterations with poor initial guesses for the DP.

4) The final DP were reasonably close to predicted values based on standard rules of thumb for these design methods.

References

[1]Mitchell, D. G., and Hoh, R. H., "Development of Methods and Devices to Predict and Prevent Pilot-Induced Oscillations," AFRL-VA-WP-TR-2000-3046, U.S. Air Force Research Lab., Wright–Patterson AFB, OH, Jan. 2000.

[2]Moorhouse, D. J. (ed.), *Flight Control Design—Best Practices*, NATO Research and Technology Organization, RTO-TR-029, Task Group SCI-026, Dec. 2000.

[3]Tischler, M. B., Colbourne, J. D., Morel, M. R., Biezad, D. J., Cheung, K. K., Levine, W. S., and Moldoveanu, V., "A Multidisciplinary Flight Control Development Environment and Its Application to a Helicopter," *IEEE Control Systems Magazine*, Vol. 19, No. 4, 1999, pp. 22–33.

[4]Duke, L., "AIAA Controls Design Challenge," *Proceedings of the 1992 AIAA Guidance, Navigation and Control Conference*, AIAA, Washington, DC, 1992, pp. 1527–1580.

[5]Magni, J.-F., Bennani, S., and Terlouw, J. (eds.), *Flight Control: A Design Challenge, (GARTEUR)*, Springer-Verlag, Berlin, 1997.

[6]Blakelock, J. H., *Automatic Control of Aircraft and Missiles*, Wiley-Interscience, New York, 1991, pp. 122–124.

[7]Tischler, M. B., "Digital Control of Highly Augmented Combat Rotorcraft," NASA TM 88346, ARMY TR 87-A-5, May 1987.

[8]Franklin, G. F., and Powell, J. D., *Digital Control of Dynamic Systems*, Addison Wesley Longman, Reading, MA, 1998, pp. 458–459.

[9]Bryson, A. E., and Ho, Y. C., *Applied Optimal Control*: *Optimization, Estimation, and Control*, Hemisphere, New York, 1975, p. 149.

[10]Lane, S. H., and Stengel, R. F., "Flight Control Design Using Non-Linear Inverse Dynamics," *Automatica*, Vol. 24, No. 4, 1988, pp. 471–483.

[11]Franklin, J. A., "Attitude Control Augmentation and Tracking Display for XV15," NASA TM-209599, Sept. 2000.

[12]Schumacher, C., "Missile Autopilot Designs Using H-infinity Control and Dynamic Inversion," *Journal of Guidance, Control, and Dynamics*, Vol. 21, No. 2, 1998, pp. 234–243.

[13]Postlethwaite, I., *Multivariable Feedback Control*, Wiley, New York, 1996, pp. 384–388.