AIAA 2001-4266

OPTIMIZATION AND COMPARISON OF ALTERNATIVE FLIGHT CONTROL SYSTEM DESIGN METHODS USING A COMMON SET OF HANDLING-QUALITIES CRITERIA

Mark B. Tischler*
Aeroflightdynamics Directorate (AMRDEC)
US Army Aviation and Missile Command
Ames Research Center
Moffett Field, CA
94035-1000

James A. Lee†
University of Southern California
Department of Electrical Engineering
Los Angeles, CA 90089

Jason D. Colbourne‡
Aeroflightdynamics Directorate (AMRDEC)
US Army Aviation and Missile Command
Ames Research Center
Moffett Field, CA
94035-1000

ABSTRACT

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, LQR, 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 RMS, while maintaining the 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 (1 or 2 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 perspective control designs against a comprehensive set of design requirements. CONDUIT® has proven to be an especially effective environment for this task.

INTRODUCTION

The design, integration, and flight test development of flight control systems factor significantly into the overall time and cost of aircraft development. By one estimate, over 25% of developmental flight testing hours for the UH-60 (BlackHawk) and RAH-66 (Comanche) helicopters were associated with flight control-related issues.[1] With costs that reach $50K per flight test hour for a modern flight test facility, there remains a considerable premium on control law design

*Flight Control Technology Group Leader, Associate Fellow AIAA
†Graduate Student
‡Research Scientist

and optimization methods that can streamline the development process.

Recently, the US Air Force completed an extensive review of flight control practices for avoiding pilot-induced-oscillations (PIO),[2] in response to the fact 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."[2] A second compilation of best flight control design practices[3] was prepared under the auspices of the NATO AGARD Flight Vehicle Integration Panel (now the RTO 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.

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 (e.g., bandwidth, rise time, phase lag, etc) and detailed supporting explanation are compiled in specification documents for fixed-wing handling-qualities (Military Standard 1797A), rotorcraft handling-qualities (Design Standard ADS-33E), general flight control stability requirements (MilSpec 9490), and much additional excellent design guidance.[3] 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.[4]

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, 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. Further, 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 usage. Since prospective alternative design methods are not optimized against a common set of dynamics response objectives, it is difficult to fairly compare their performance. Some useful explanations and comparisons of alternative design methods are presented in the AIAA Controls Design Challenge[5] and the GARTEUR Flight Control Design Challenge.[6]

The Control Designer's Unified Interface (CONDUIT®),[4] used in the current study, addresses the issues discussed above 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:

- Extensive pre-coded graphical libraries of key handling-qualities and dynamic response criteria design specifications (specs)
- Integrated environment to create, validate, and catalogue new user-defined specs
- Easy graphical selection, setup, and evaluation of specs
- Accommodation for any design method or architecture
- Automated tuning of user-selected design parameters using a vector optimization method which ensures that each dynamic response criteria is individually met – rather than only a weighted average being met
- Ultimate design optimization to minimize the selected performance objective that meet the design criteria with minimum overdesign
- Extensive supporting plots and integrated analysis tools

These features allow engineers to identify complex trade-offs, for example, in terms of performance and robustness of the proposed design methods, actuation techniques, and selection of sensors. 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 can be found in Ref. 4.

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 linearized dynamics. The design methods considered are: Classical, 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[7] and is shown in Fig. 1. The control system topology was selected for this case study and does not represent the actual implementation in the KC-135 aircraft.

The key elements of the block diagram are:

- 3 DOF state-space representation of KC-135 lateral/directional dynamics
- 2nd order actuator dynamics including rate and position limiting
- angular rate gyro filters
- washed-out yaw rate feedback (to enable turn coordination)
- roll rate and yaw command
- disturbance inputs

The feedback loop architecture shown in Fig.1 is appropriate to the *classical design.*

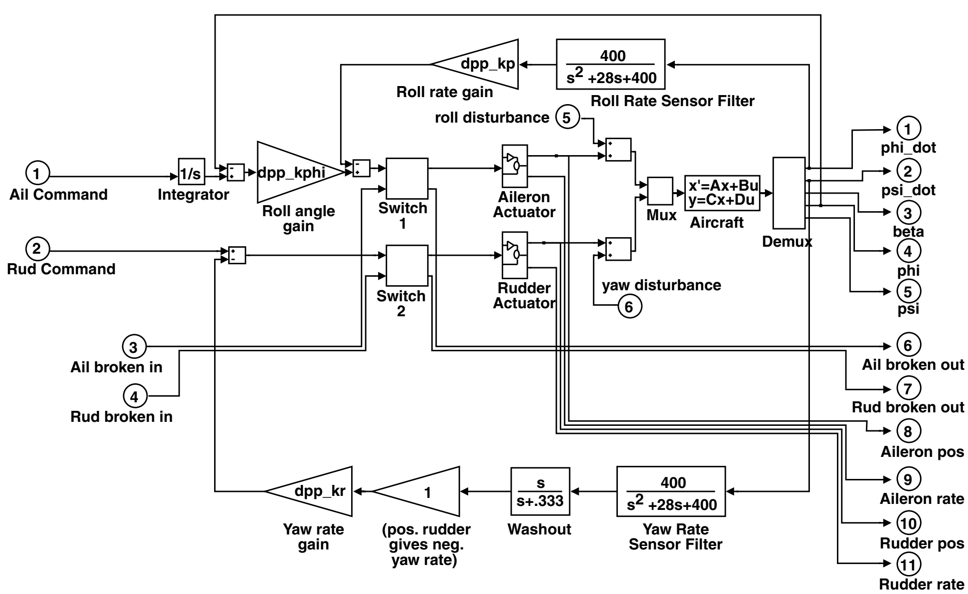

Fig. 1. Case Study Problem Based on KC-135 Lateral/Directional Dynamics.

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 spec is designated by the user as indicated by an "H", "S", or "J" indicated in the upper right hand corner of the spec. The role of the spec priority in the phases of CONDUIT® optimization process is summarized below, but is described more fully in Ref. 4.

In Phase I, the design parameters are tuned to attempt to meet the "hard specifications (H)" selected for this design study:

- All closed-loop eigenvalues must lie in the left-half plane (absolute stability) (EigLcG1)
- Gain/phase margin requirements (9490) to ensure satisfactory relative stability/robustness. (StbMgG1)

If Phase I specs are met, the solution enters Phase II, in which the design parameters are tuned to attempt to meet all of the "soft specifications (S)" and performance metrics included in the "summed objective (J)," with the earlier hard specs enforced as constraints in the optimization. The soft specs 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 soft specs are from MilStandard 1797A:

- Roll command response bandwidth (BnwRoD1)
- Equivalent system Dutch roll mode damping (DmpDrD3) and frequency (FrqDrD4)
- Roll time response quickness (QikAtG1)

Additional soft specs included from the CONDUIT® libraries are:

- Actuator saturation limits (SatAcG1)
- Disturbance response (HldNmH1)

Finally, if Phase II specs are met, the solution enters Phase III. Here, the design parameters are tuned to attempt to *minimize* the "summed objective specification (J)," with the earlier hard and soft specs enforced as constraints in the optimization.

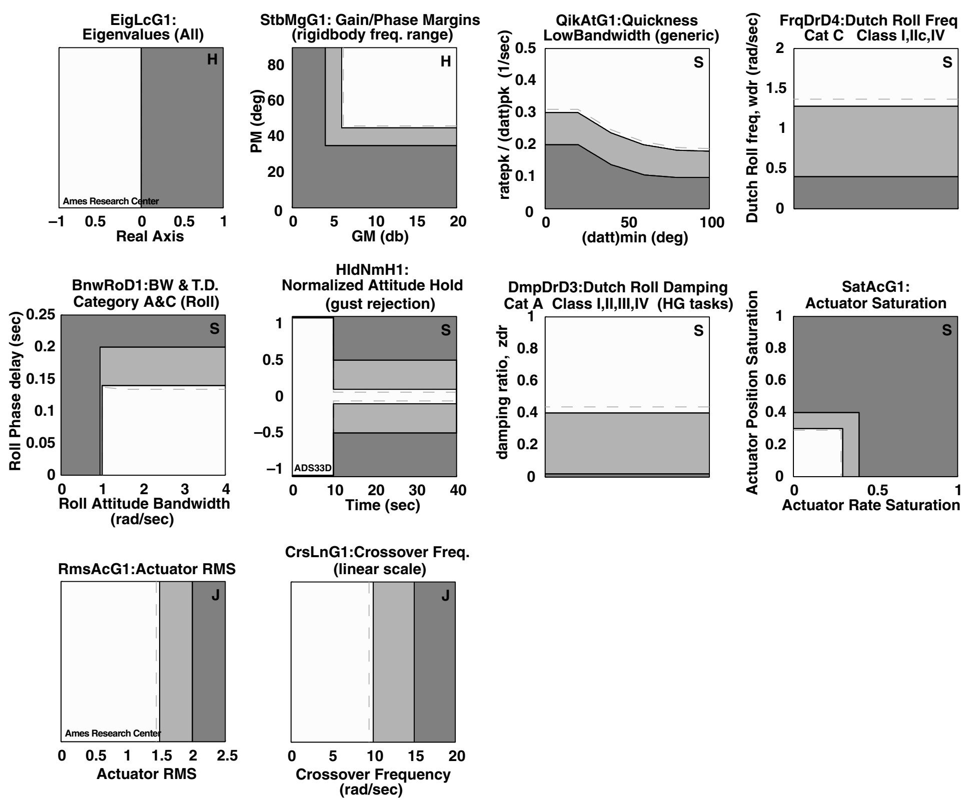

Fig. 2. Case Study Design Specifications.

Extensive practical design experience using CONDUIT® has indicated that good performance metrics to be included in the summed objective (J) are:

- Crossover frequencies for the individual broken loops (CrsLnG1)
- Actuator position response RMS (RmsAcG1)

The RMS values are determined from a power spectrum calculation in the frequency-domain, and are normalized to maximum control position, so that a value of unity reflects a command of full control surface authority (100% saturation).

By selecting these performance metrics, we ensure that the final design will meet all the requirement specs 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 III, in contrast to the min/max vector optimization of Phase I and II, 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 the loops to be tuned to their ultimate performance. Detailed explanation of this strategy is given in Ref. 4.

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 ("Milspec") documents (e.g., Ref. 3) 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 spec point lies a safe distance within the Level 1 region and not right on the Level 1/2 boundary. As shown in Ref. 4, 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 spec in Fig. 2.

ALTERNATIVE DESIGN METHODS

This paper presents and compares KC-135 case study results for four alternative design methods: (1) Classical, (2) LQR, (3) Dynamic Inverse, (4) H-Infinity.

The alternative methods can be characterized by the number of design parameters and the number of controller degrees-of-freedom. The design parameters are any quantities (e.g., 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 3 design parameters (feedback gains), while the remaining methods were implemented 4 design parameters.

The design methods can additionally be characterized by the number of controller degrees-of-freedom (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 a response feedback DOF and a command model (or forward loop) compensator (CM). 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 degree-of-freedom (RF), while the Dynamic Inverse method is inherently a two degree-of-freedom (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 above. 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 (Kphi) 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 three feedback gains (Kp, Kphi, Kr) are designated as CONDUIT® tuning parameters via the "dp_" or "dpp_" in prefix seen in the block diagram of Fig. 1, and comprise a 1DOF (RF) architecture as indicated in Table 1.

The initial values for the Classical design parameters were selected as:

$$\left.\begin{aligned} Kp &= 0.5 \text{ deg / deg / sec} \\ Kphi &= 0.25 \text{ deg / deg} \\ Kr &= 0.25 \text{ deg / deg / sec} \end{aligned}\right\} \qquad (1)$$

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. 3) and highlight the robustness of the CONDUIT® optimization to poor initial design guesses. As seen in Fig. 3, the initial design is in the Level 2 ("deficiencies warrant improvement") region for many of the design requirements.

The iteration history plot (Fig. 4) shows that after six iterations CONDUIT® reached Phase III, which is a "feasible solution" where all specs are in the Level 1 region. The fully converged result is shown in Fig. 5. The optimized solution achieved after 12 iterations meets all hard (H) and soft (S) requirements while minimizing the performance metrics (J). Convergence of the design in Phase III toward the final optimized solution is quite smooth as seen in Fig. 4, and is given by:

$$\left.\begin{aligned} Kp &= 0.12 \text{ deg / deg / sec} \\ Kphi &= 0.24 \text{ deg / deg} \\ Kr &= 2.82 \text{ deg / deg / sec} \end{aligned}\right\} \qquad (2)$$

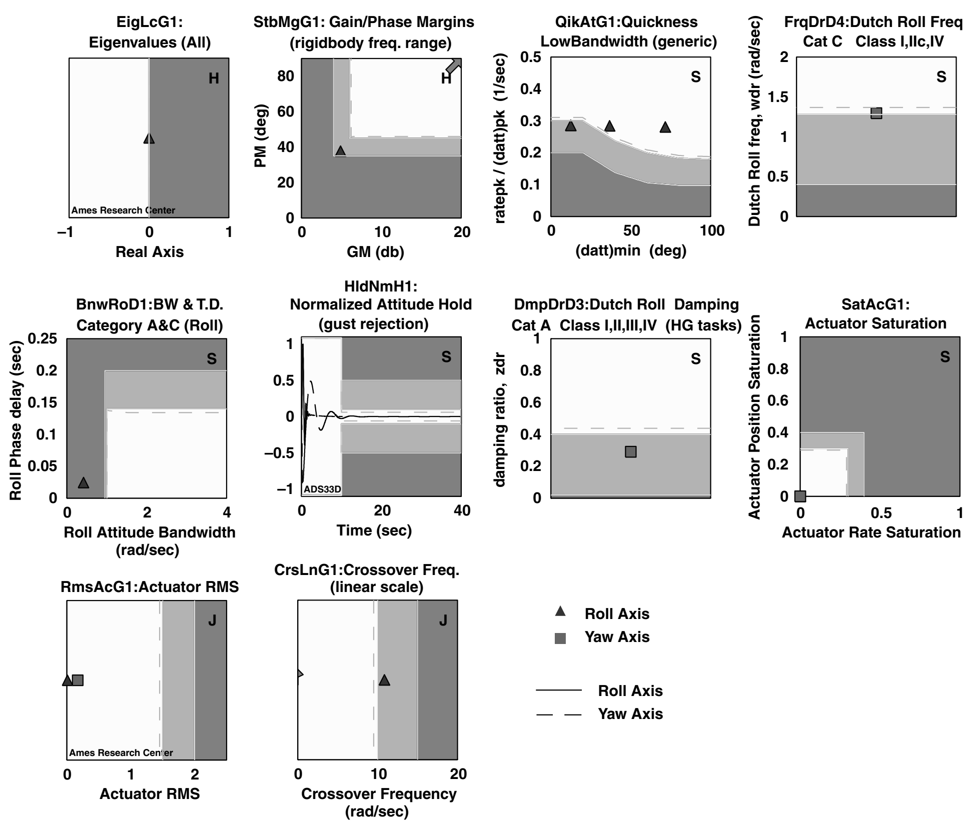

Fig. 3. Initial Design for Classical Method.

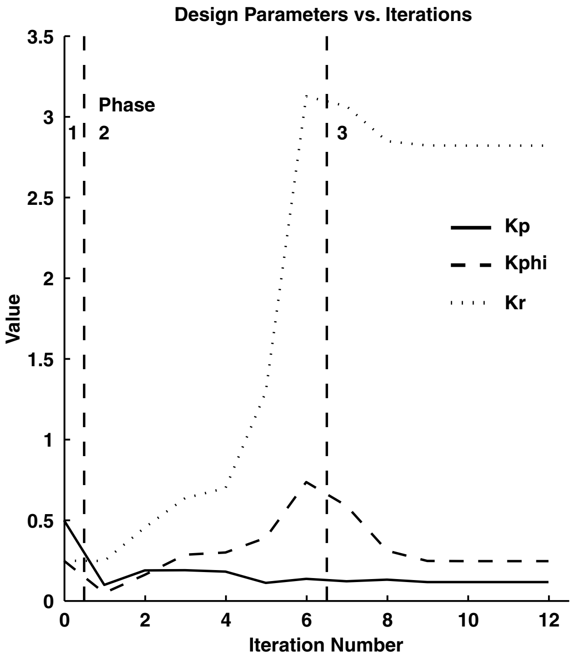

Fig. 4. Classical Design Optimization.

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

A key characteristic of the converged Phase III design in CONDUIT® as seen in Fig. 5 is that one (or more) of the soft or hard requirements lies on the design margin (dotted) boundary. This indicates that the Level 1 design requirements have been met with minimum overdesign, since 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:

- Control law requires measurements of p, ϕ, and r
- All Level 1 design requirements achieved (with 10% design margin)
- No control saturation for the largest expected command inputs
- Reasonable crossover frequencies (3-4 rad/sec)
- Overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements (e.g. quickness and bandwidth) for the 1DOF architecture.

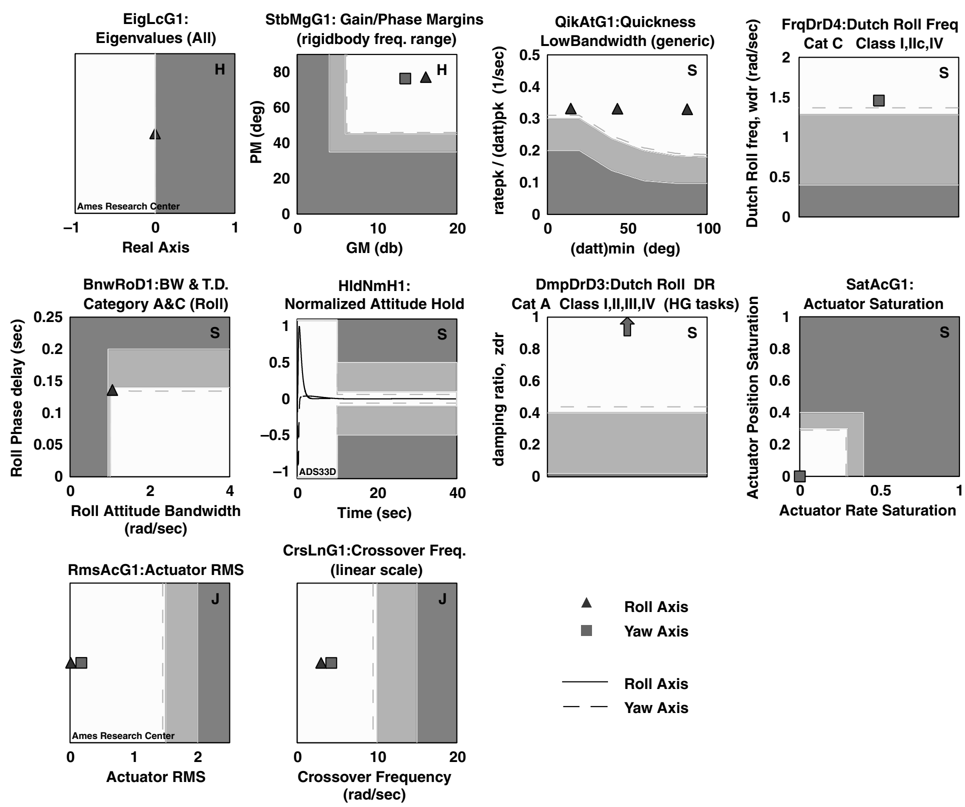

Fig. 5. Optimal Classical Design Characteristics

Each spec in CONDUIT® has an associated set of "supporting plots" which illustrates the various related analyses and calculations. For example, the supporting plots for the stability margin spec (StbMgG1 in Fig. 5) show the broken-loop response plots for the aileron (Fig. 6a) and rudder (Fig. 6b) loops that are used in the gain and phase margin calculations. Reference to Fig. 6a shows that the roll crossover frequency could be reduced considerably, while still maintaining adequate phase margin. However, the 1DOF architecture requires the higher gains to achieve the remaining handling-qualities requirements.

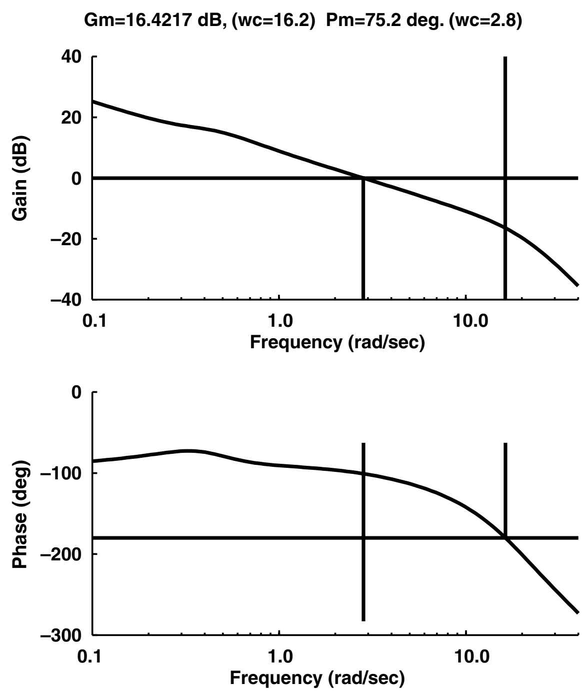

Fig. 6a. Aileron Broken-Loop Response (Optimized).

An interesting aspect of the results is the comparison of the optimized design parameters with "rules of thumb" for classical control. 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. 8. Referring to the block diagram of Fig. 1, the actuator and gyro filter dynamics contribute a total effective delay of $\tau_{eff} = 0.092$ sec.

The maximum achievable roll crossover frequency is then determined:

$$(\omega_c)_{max} = 0.370 / \tau_{eff} = 4.02 \text{ rad / sec} \qquad (3)$$

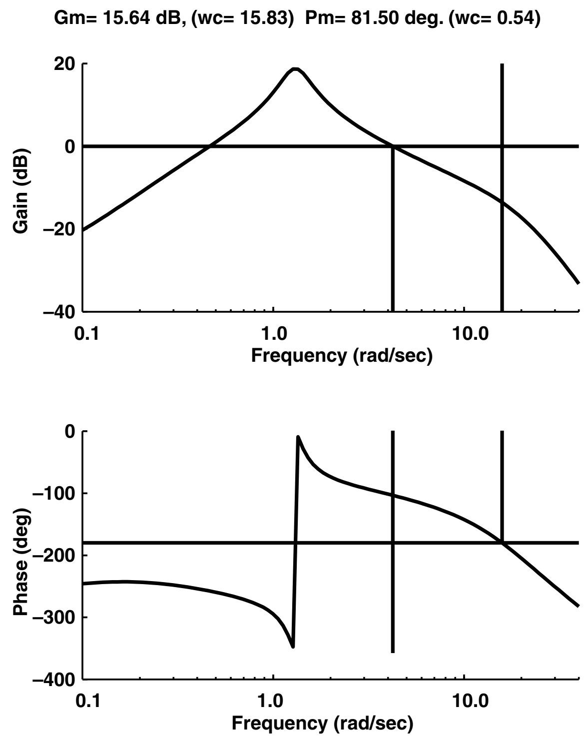

Fig. 6b. Rudder Broken-Loop Response (Optimized).

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

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

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

$$\frac{1}{T_p} = \frac{K_\phi}{K_p} = 0.442\,\omega_c = 1.78 \text{ rad / sec} \qquad (5)$$

which finally determines the roll rate gain:

$$K_p = 0.16 \text{ deg / deg / sec} \qquad (6)$$

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

$$K_r = \frac{\omega_c}{N_{\delta_{rud}}} = 2.94 \text{ deg/ deg/ sec} \qquad (7)$$

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.

LQR DESIGN

The control law architecture for the LQR design is shown in Fig. 7.

The LQR (linear quadratic regulator) method calculates the optimal gain matrix K for that the (full) state feedback law that minimizes the quadratic cost function comprised of 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 the elements represent weights for the corresponding states. In practice, R is usually set to the identity matrix (Ref. 9). 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 (i.e., 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 (number of states = 4: p, r, β, ϕ), in the Riccati solution. This ensures that we will only require these aircraft states as feedback measurements, and not all the internal states of the control system. However, this approach also degrades the "guaranteed stability" of the LQR design, since the omitted elements (e.g., 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) 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 degree-of-freedom. Therefore, as indicated in Table 1, the LQR implementation is 1 DOF with 4 design parameters.

Fig. 8 shows that the handling-qualities specifications for the initial design parameter guesses (Q=I) are far from ideal. We see that the initial design parameters result in Level 3 behavior for the eigenvalue and stability margin requirements and Level 2 behavior for several other requirements, thus presenting a poor initial design with which to challenge CONDUIT®.

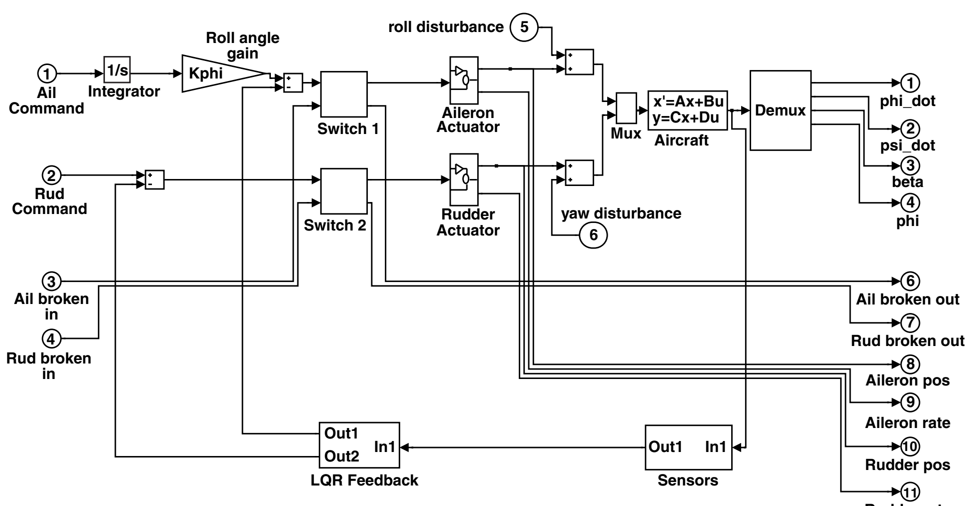

Fig. 7. LQR Control Law Architecture.

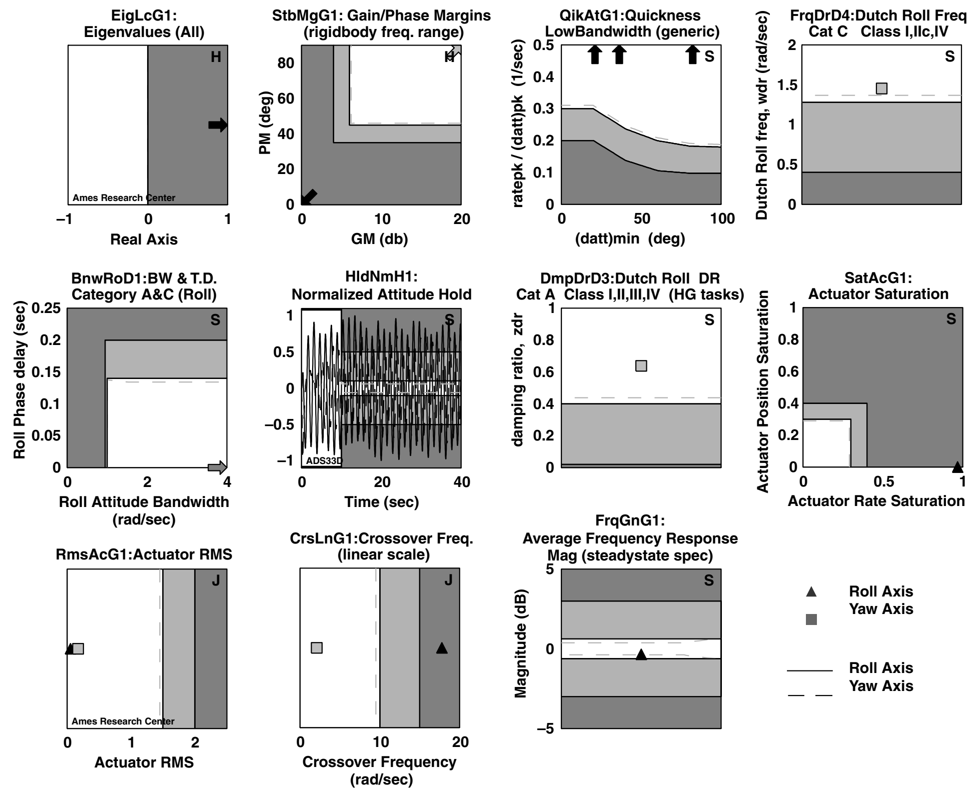

Fig. 8. Initial Design for LQR Method

CONDUIT® tuned the Q matrix to achieve an LQR design that satisfied all the handling qualities, even when the initial guess for the design parameters was poor. The design parameters are fine tuned in Phase III to a converged solution after 12 iterations (Fig. 9) that meets all Level 1 design requirements (Fig. 10) 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:

$$\left.\begin{aligned} Q_{pp} &= 0.0079 \\ Q_{rr} &= 0.81 \\ Q_{\beta\beta} &= 0.95 \\ Q_{\phi\phi} &= 0.071 \end{aligned}\right\} \tag{9}$$

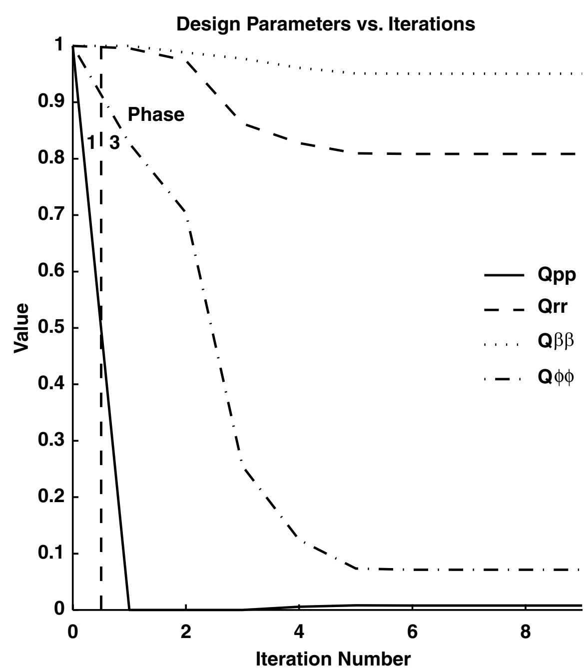

Fig. 9. LQR Design Optimization.

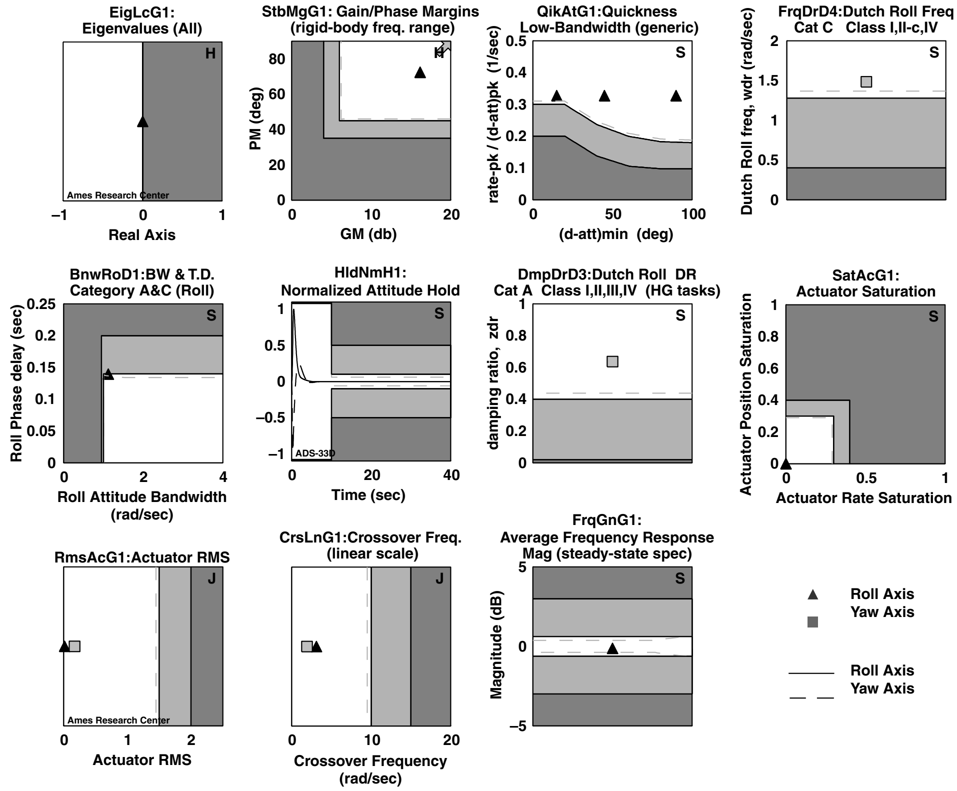

Fig. 10. Optimized LQR Control System Characteristics.

The corresponding feedback gain matrix is:

$$
K = \begin{array}{c} \begin{array}{cccc} p & r & \beta & \phi \end{array} \\ \begin{bmatrix} .1141 & -.2024 & .2262 & .2823 \\ .0088 & -.8456 & .2455 & .0061 \end{bmatrix} \end{array} \begin{array}{c} \delta_{ail} \\ \delta_{rud} \end{array} \qquad (10)
$$

The key characteristics of the optimized LQR solution are:

- Control law requires measurements of p, r, β, φ
- All Level 1 design requirements are achieved (with 10% design margin)
- No control saturation for the largest expected command inputs
- Reasonable crossover frequencies (2-3 rad/sec)
- Overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements (e.g. quickness and bandwidth) for the 1DOF architecture.

It is interesting to compare the optimized LQR design parameters with the rule of thumb guidance for Q matrix selection given by Bryson in Ref. 10. Professor Bryson recommends setting the diagonal elements of Q to correspond to the square of the ratio of maximum control authority to maximum expected response, i.e.:

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

When we assume a maximum control input of 5 deg (aileron and rudder) and maximum expected state responses [50 deg/sec, 5 deg/sec, 5 deg, 50 deg], corresponding to roll rate, yaw rate, sideslip, and roll angle, respectively, then we obtain:

$$
\left.\begin{aligned} Q_{pp} &= .01 \\ Q_{rr} &= 1.0 \\ Q_{\beta\beta} &= 1.0 \\ Q_{\phi\phi} &= .01 \end{aligned}\right\} \qquad (12)
$$

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

DYNAMIC INVERSE DESIGN

The dynamic inverse control method (Ref. 11) 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 2 DOF architecture (Table 1). However, the method requires carrying a look-up table of accurate stability and control derivatives in the on-board software for the linear implementation, or carrying a complete nonlinear simulation on-board in the nonlinear implementation. The linear implementation is adopted herein (Fig. 11) following closely the approach outlined in Franklin (Ref. 12). 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 acceleration-like dynamics are easily tuned to produce the desired regulator bandwidth and stability margins. Finally, the command models (CM) 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. 11.

Four design parameters were selected for CONDUIT® optimization: inverse time constant (i.e., bandwidth) for the first-order roll command model (Tp), roll regulator natural frequency (omphi), Dutch Roll command model 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:

$$\left.\begin{aligned} \text{Tp} &= 1 \text{ sec} \\ \text{comdr} &= 1 \text{ rad / sec} \\ \text{ombet} &= 1 \text{ rad / sec} \\ \text{omphi} &= 1 \text{ rad / sec} \end{aligned}\right\} \qquad (13)$$

resulting in the control system performance far from ideal (Fig. 12). The roll bandwidth is in the Level 3 region, and the Dutch Roll frequency is in the Level 2 region.

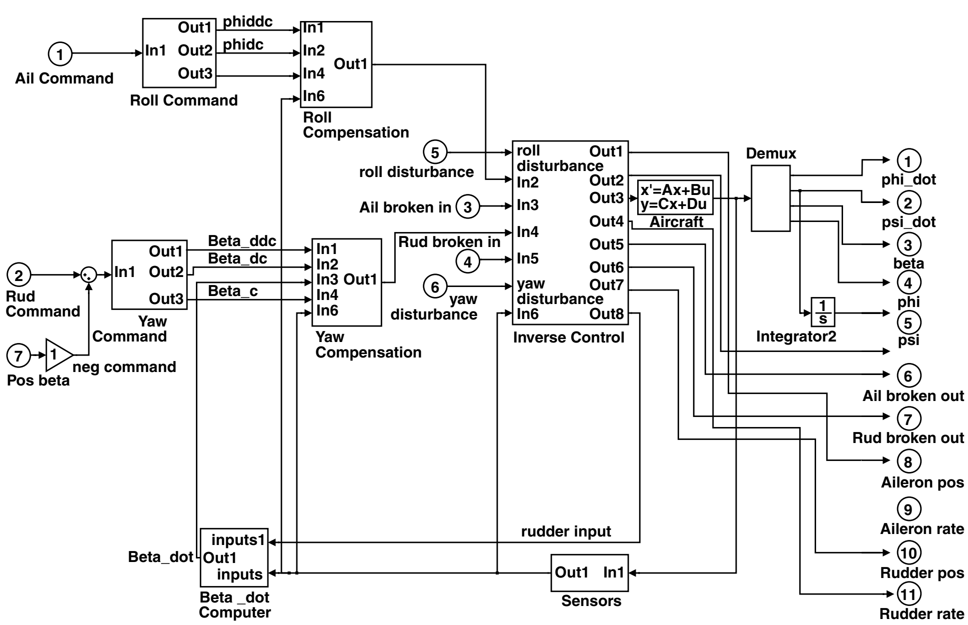

Fig. 11. Dynamic Inverse Control Architecture.

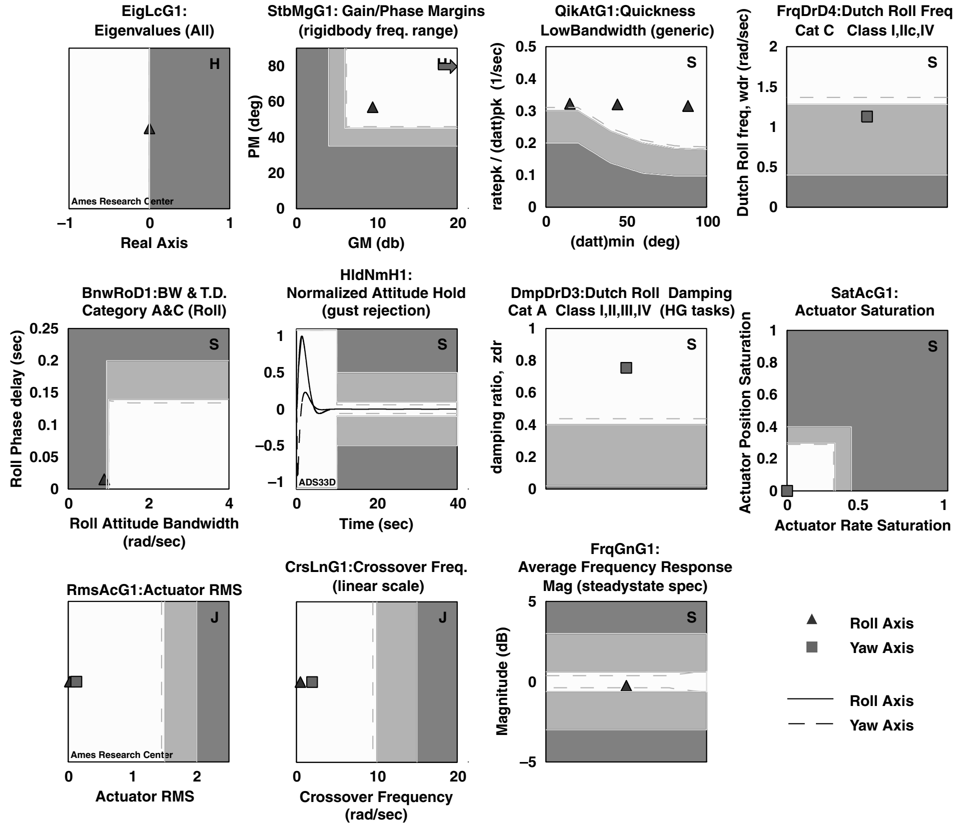

Fig. 12. Initial Design for Dynamic Inverse

As seen in Fig. 13, CONDUIT® converges in 11 iterations to an optimized design that satisfies all Level 1 requirements (Fig. 14) while minimizing overdesign (Table 2). The optimized design parameters are:

$$\left.\begin{aligned} \text{Tp} &= 0.837 \text{ sec} \\ \text{comdr} &= 1.21 \text{ rad / sec} \\ \text{ombet} &= 0.76 \text{ rad / sec} \\ \text{omphi} &= 0.835 \text{ rad / sec} \end{aligned}\right\} \qquad (14)$$

The generation of high actuator rates and /or displacements has been reported in some previous applications of the dynamic inverse control method (Ref. 13). While 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 ($\omega_{c_\phi} = 0.344 \text{ rad / sec}$), 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 2DOF controller architecture, which allows independent tuning of the feedback response (i.e. crossover characteristics) from the command response.

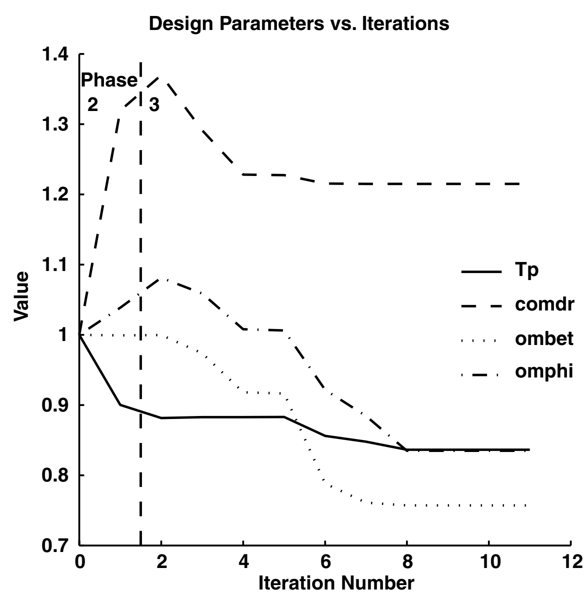

Fig. 13. Dynamic Inverse Method Optimization.

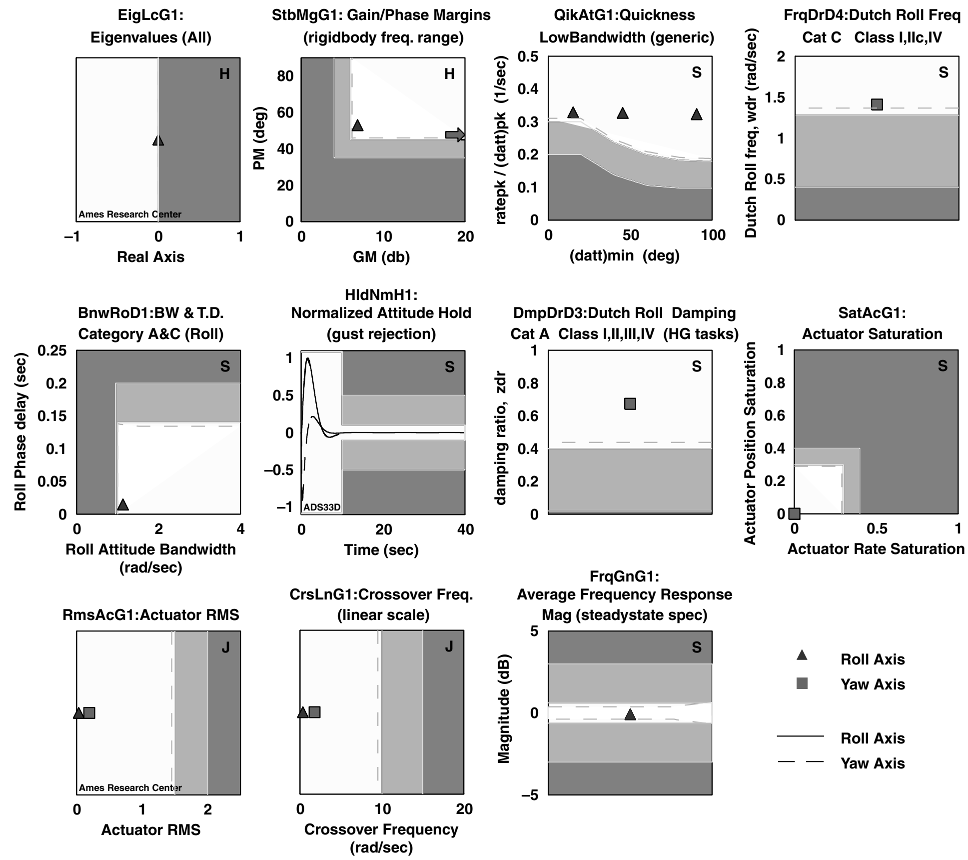

Fig. 14. Optimized Dynamic Inverse Control System Characteristics.

In the roll channel, satisfactory stability is achievable at low crossover frequencies as can be seen in Fig. 6a. Therefore for this (2DOF) design, CONDUIT® reduces the crossover frequency as an element of the summed objective until the attitude hold spec (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. It is important to emphasize that these same characteristics would be achieved with any of the other methods with the introduction of a second controller DOF.

Summarizing the results of the Dynamic Inverse design method:

- Control law requires measurements of p, r, β, ϕ, δ_{rud} or a_y
- Control law requires dynamic estimate of $\dot{\beta}$
- Level 1 design requirements achieved (with 10% design margin)
- Increased aileron control RMS
- No control saturation for the largest expected command inputs
- Low roll loop crossover frequency
- No overdesign in stability margins or disturbance rejection

H-INFINITY DESIGN

The H-infinity control method implemented in this case study used the 1DOF loop-shaping approach described by Postlethwaite (Ref. 14). 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 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) is comprised of first-order filters:

$$\begin{aligned} W1_{ail} &= k_{ail} / (s + a_{ail}) \quad \text{, for the aileron loop} \\ W1_{rud} &= k_{rud} / (s + a_{rud}) \text{ , for the rudder loop} \end{aligned} \qquad (15)$$

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. 15, and is listed as 1DOF with four dp’s 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 (Fig. 16):

$$\left.\begin{aligned} k_{ail} &= 6 \\ a_{ail} &= 50 \\ k_{rud} &= 160 \\ a_{rud} &= 50 \end{aligned}\right\} \qquad (16)$$

CONDUIT® optimized the design parameters to satisfy Level 1 requirements (Figs. 17, 18), with some overdesign in several specs due to the 1DOF nature of the controller:

$$\left.\begin{aligned} k_{ail} &= 17.11 \\ a_{ail} &= 48.63 \\ k_{rud} &= 159.8 \\ a_{rud} &= 50.58 \end{aligned}\right\} \qquad (17)$$

which shows a significant change only the first design parameter. As might be expected, the optimized roll-off in the loop transmission corresponds to just below the actuator bandwidth (62.8 r/s), to keep from overdriving the control inputs.

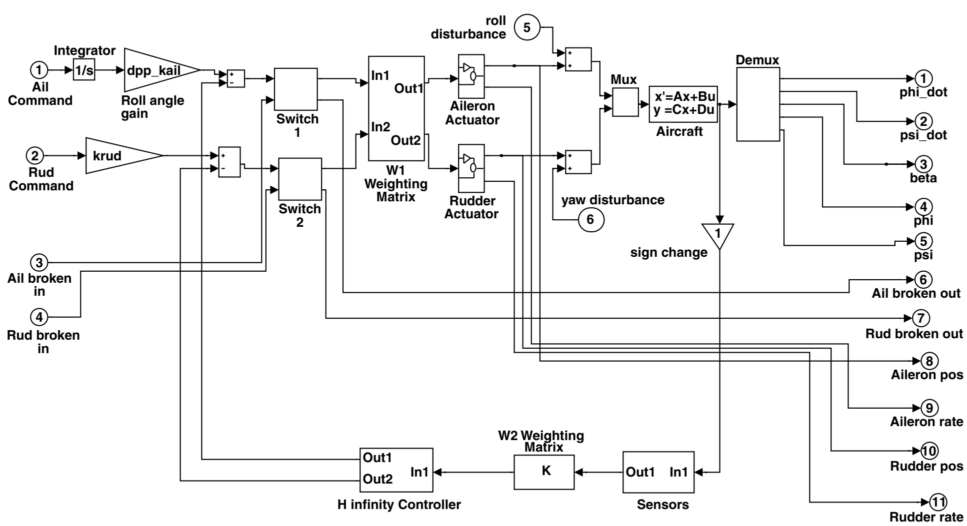

Fig. 15. H-infinity Control System Architecture.

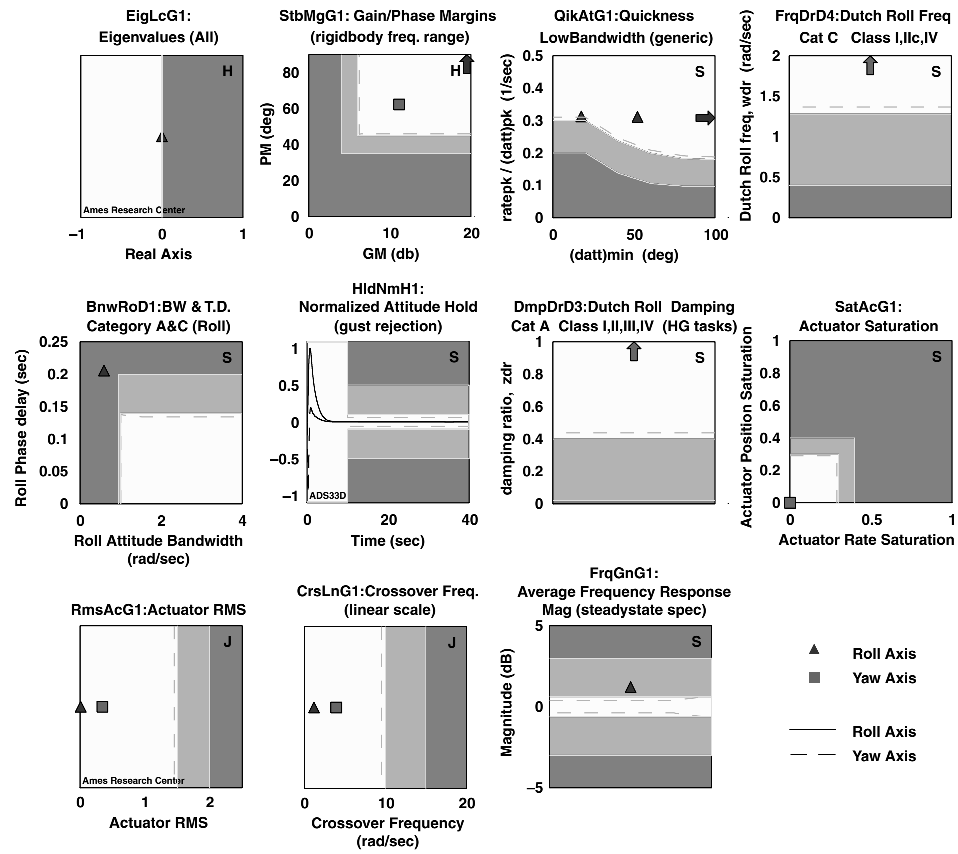

Fig. 16. Initial Design for H-infinity Control Law Design.

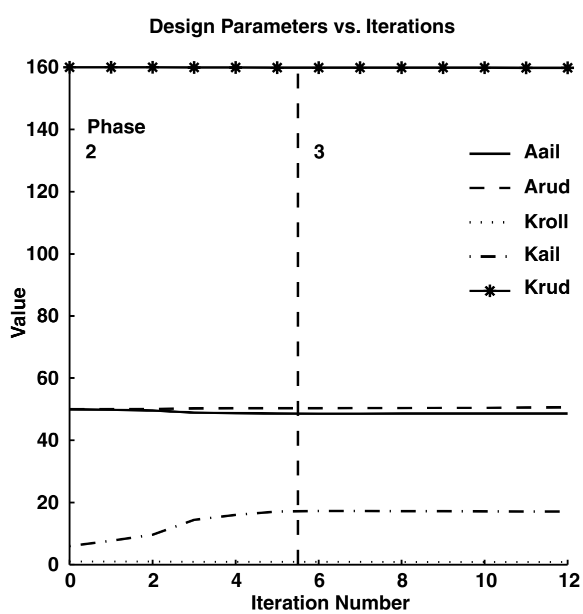

Fig. 17. H-infinity Design Method Optimization.

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 (1st 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.

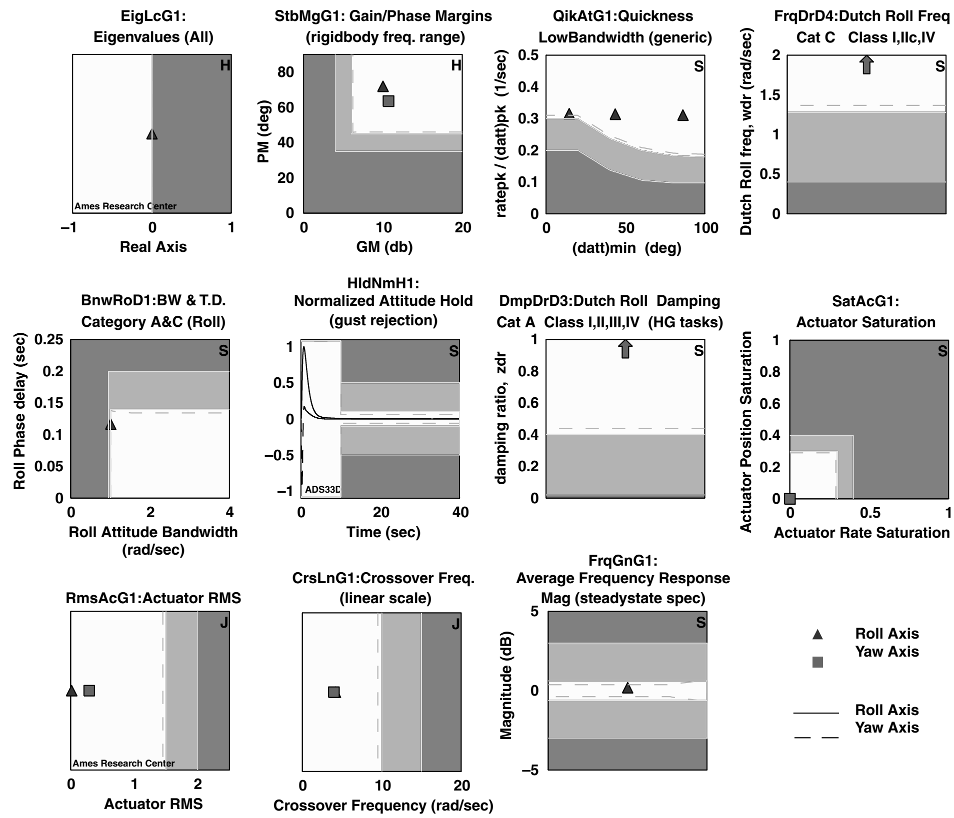

Fig. 18. Optimized H-infinity Control System Characteristics.

The performance metrics for this design are listed in Table 2.

Summarizing the results of the H-infinity design method:

- Control law requires full state measurements of p, r, β, ϕ
- 6^{th} order dynamic compensator (4^{th} order controller plus W1 filters)
- All Level 1 design requirements achieved (with 10% design margin)
- Increased rudder control RMS
- No control saturation for the largest expected command inputs
- Reasonable crossover frequencies (4 rad/sec)
- Overdesign in stability margins and disturbance rejection to meet the handling-qualities requirements (e.g. quickness and bandwidth) for the 1DOF architecture
- Optimization was quite sensitive to initial choices of the design parameters

COMPARISON OF STABILITY ROBUSTNESS FOR ALTERNATIVE DESIGN METHODS

An analysis of the stability robustness for the alternative design methods was conducted. 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 design parameters fixed as determined previously (Eqs 2, 9, 14, 17). Fig. 19 displays 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 previous CONDUIT® optimization results. The 1DOF design methods (Classical, LQR, 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 2DOF Dynamic Inverse design, the nominal margins are somewhat reduced, so 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. 19 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. 19). 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.

DISCUSSION:

The key results of this comparative design study are summarized in Tables 1, 2, and 3. The design methods were intentionally set-up 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 3, the three methods that use a 1DOF architecture (Classical, LQR, H-infinity) achieve comparable performance as defined by the performance metrics: crossover frequency and control RMS. The addition of a second DOF in the Dynamic Inverse method relaxes the constraint between the feedback characteristics and the command response characteristics as imposed by the 1DOF architecture. This permits a significant reduction in aileron loop crossover frequency. The control RMS values is comparable for the four methods, although there is an increase in aileron RMS for the Dynamic Inverse design and an increase in rudder RMS for the H-infinity design. The stability robustness is somewhat better for the classical and H-infinity design methods, but still quite comparable across the four methods.

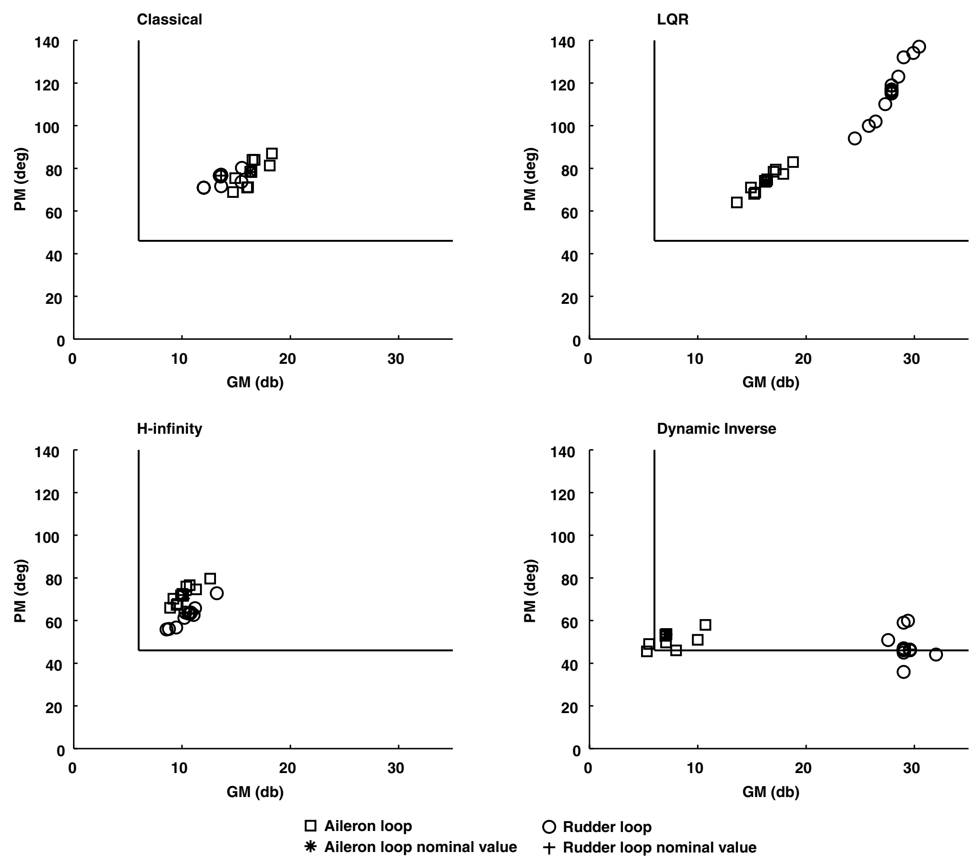

Fig. 19. Robustness Comparison for 20% Plant Variations.

The overall results of this study suggest the careful definition of design requirements sets a unique broken-loop shape and command model response. CONDUIT® will alternatively drive on the respective design parameters for 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 specs; (2) selection of a 1 vs. 2DOF 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 in this paper 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 which include all relevant handling-qualities and controller performance metrics.

2. Four alternative design methods (Classical, LQR, Dynamic Inverse, 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 design parameters.

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

REFERENCES:

1. Crawford, C., "Potential for Enhancing the Rotorcraft Development Qualification Process using System Identification," RTO SCI Symposium on System Identification for Integrated Aircraft Development and Flight Testing, Madrid, Spain, May 5-7, 1998.

2. 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, January, 2000.

3. Moorhouse, D. J. (Ed.), "Flight Control Design – Best Practices," RTO-TR-029, Task Group SCI-026, December 2000.

4. Tischler, M. B., et al, "A Multidisciplinary Flight Control Development Environment and Its Application to a Helicopter", *IEEE Control Systems Magazine*, Vol. 19, No. 4, August, 1999, pp. 22-33.

5. Duke, L. AIAA Controls Design Challenge, 1990.

6. Magni, J-F, Bennani, S., and Terlouw, J (Eds), *Flight Control: A Design Challenge (GARTEUR)*, Springer-Verlag, 1997.

7. Blakelock, J. H., *Automatic Control of Aircraft and Missiles*, Wiley-Interscience, 1991.

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

9. Franklin, G. F, and Powell, J. D, *Digital Control of Dynamic Systems*, Addison-Wesley, 1998.

10. Bryson, A. E., *Control of Spacecraft and Aircraft*, Princeton Press, Princeton 1994.

11. Lane, S. H. and Stengel, R. F., "Flight Control Design Using Non-Linear Inverse Dynamics," *Automatica*, Vol 24, No. 4, 1988.

12. Franklin, J. A., “Attitude Control Augmentation and Tracking Display for XV15,” NASA TM-209599, Sept. 2000.

13. Schumacher, C., "Missile Autopilot Designs Using H-infinity Control and Dynamic Inversion," *Journal of Guidance, Control, and Dynamics*, Vol. 21, No. 2, March-April, 1998.

14. Postlethwaite, I., *Multivariable Feedback Control*, Wiley, 1996, pp.384-388.

Table 1 – Alternative Design Methods
(RF = “response feedback”, CM = “command model”)

Method	**DOF’S**	**Tuning parameters (DP’s)**	**# of DP’s**
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

Table 2 – Comparison of Performance Metrics

Objective Metric	**Classical**	**LQR**	**Dynamic Inverse**	**H-infinity**
ω_c , aileron, rad/sec	2.89	3.10	0.344	4.12
ω_c , rudder, rad/sec	4.24	1.92	1.772	3.96
RMS, aileron, nondim.	0.0118	0.0137	0.0276	.0131
RMS, rudder, nondim.	0.1759	0.1702	0.1813	0.291

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