

Handling-Qualities Optimization and Trade-offs in Rotorcraft Flight Control Design

**Mark B. Tischler
Christina M. Ivler
M. Hossein Mansur**

**Aeroflightdynamics Directorate (AMRDEC)
US Army Research, Development, and Engineering Command
Ames Research Center
Moffett Field, CA
USA**

**Kenny K. Cheung
Tom Berger
Marcos Berrios**

**University Affiliated Research Center (UARC)
University of California, Santa Cruz
Moffett Field, CA
USA**

Abstract

Designing rotorcraft flight control systems to meet the handling qualities and stability margin requirements of Aeronautical Design Standard-33 (ADS-33E-PRF) and MIL-DTL-9490E ensures low-pilot workload, increased mission effectiveness and improved safety of operations in all weather and visibility conditions. The numerous requirements compete with one another and can result in a highly-constrained design space. Therefore, achieving a satisfactory design that makes best use of the limited available control power can be a very time consuming process. This paper demonstrates how multi-objective parametric optimization (CONDUIT[®]) can be used to optimize the many design parameters of a rotorcraft flight control system and meet the competing design requirements. Typical design trade-offs are demonstrated for a simple classical response-feedback flight control system for the XV-15 tilt-rotor aircraft. A more complex design study based on the UH-60 demonstrates how optimization methods are used for a modern multi-mode model-following system. Flight test data for the UH-60 RASCAL in-flight simulator validate the design models and predicted handling-qualities.

1 Introduction

The starting point of rotorcraft flight control system design is a detailed definition of the intended missions (*mission task elements*¹), visibility and vision aides (*usable cue environment*) and operational weather conditions. From here, the handling qualities and stability margin requirements of Aeronautical Design Standard-33 (ADS-33) (Anon 2000) and MIL-DTL-9490E (Anon 2008) specify the needed response types (e.g. attitude-command/attitude-hold (ACAH) vs. translational-rate command/position-hold (TRPH) and response characteristics (e.g., bandwidth) to ensure low-pilot workload, high mission effectiveness and safe operations. Compendiums of lessons learned and best practices in flight control design (AGARD 1987, AGARD 2000) further emphasize that “designing-in” good handling-qualities

¹Each key new term or concept is highlighted in *italics* the first time it is used.

characteristics at the start of the design reduces the time and cost of the development process, and minimizes the potential for costly and sometimes dangerous consequences when problems are uncovered only during flight testing. Achieving a design that meets the many and competing requirements is a challenging task for all air vehicles and perhaps more so for rotorcraft.

Four key challenges of air vehicle flight control system (FCS) design can be summarized based on an excellent review of recent projects (AGARD 2000). The first challenge is the multi-disciplinary aspect of the flight control problem. The control system designers must have a good understanding of a wide range of disciplines, including flight dynamics, structural dynamics, feedback control theory, simulation modeling, handling-qualities, sensors and actuators, redundancy management, verification and validation. Yet, most engineers are specialists and may have a mastery of only a few of these at best. Long and costly design cycles will arise if there is no common language or integrated design tool. A second challenge is that the design and evaluation of a flight control system requires checking numerous and competing design specifications – of the order of 50-100. This process is repeated for the tens (or even hundreds) of design (and off-nominal) points and several selectable modes that are evaluated for a full flight envelope control system. Finding a satisfactory design by manually “turning knobs” is not a practical approach for any more than a simple exercise. A third challenge is that the control system design engineer must continually update the design, integrating improvements in the mathematical models as hardware test data become available, implementing changes in design requirements, and incorporating pilot feedback from simulation and flight tests. A final challenge is the need for design tools that can facilitate the study of the trade-offs between competing specifications, hardware characteristics, and performance metrics, so that the final design may make the best use of available control authority. The failure to consider such trade-offs can compromise control system performance and handling qualities – especially for a *partial authority* system.

Clearly, sophisticated algorithms and associated interactive computational tools are needed to address the many aspects of the flight control design process. Beyond the general challenges of air vehicle flight control development, the literature cites significant flight control challenges on all of the recent upgrade and new-build rotorcraft programs. These flight control design challenges arise from the need for higher modes of augmentation (ACAH, ACVH) for low-speed/hovering flight in the degraded visual environment (DVE) and all-weather operations, combined with the much increased dynamics complexity of the rotorcraft due to: higher-order dynamics and strong inter-axis coupling of the bare-airframe response, multi-disciplinary nature of the rotor system analysis, very low signal-to-noise, and large response lags. Taken together these factors all greatly complicate the design problem and limit achievable flight control performance. Crawford (1998) estimated that UH-60 BlackHawk flight control development had accounted for about 37% of the overall flight test time. A comparable percentage was estimated for the RAH-66 Comanche. At a typical flight test cost of about \$50k/hr. for developmental flight testing, flight control system modifications/upgrades are considered a high financial risk proposal for most military Program Managers (PMs). A review of this experience emphasizes the need for control law architectures, design methods and analysis tools that offer the maximum transparency and insight for problem resolution. Integrated optimization tools can be a great help in achieving a good rotorcraft flight control design solution in an acceptable amount of time and flight testing.

There are three excellent compendiums of lessons learned and best practices in flight control design and development (AGARD 2000, Pratt 2000, Tischler 1996) that provide useful historical perspective and motivation for the rotorcraft flight control development method presented in this paper. These references indicate six key and repeating themes as important “do’s” of flight control:

- 1) Retain both the required (“1st tier”) and additional alternative (“2nd tier”) specifications in a multi-tier set of design specifications.
- 2) Transparency and simplicity of flight control architecture are very important and highly desirable for design (due to multidisciplinary aspects), failure (and redundancy) management, testing and troubleshooting. These considerations give a clear advantage to classical architectures (e.g., response feedback and model following) as compared to purely MIMO architectures (e.g., H_∞ and LQR).
- 3) Transparency is critical in the flight control system development process. The process must be systematic, understandable, and well documented.
- 4) Flight control system designs must allow for future growth. For example, the development of outer-loop (navigation and hold) modes. This provides an additional strength of the classical (nested loop) architecture.

- 5) Regardless of the design methods adopted, the system must meet the requirements. Key is the ability to quickly evaluate and tune the selected design architecture to meet the requirements. The multi-objective function parametric optimization method adopted herein is well suited to this consideration.
- 6) Actuator rate limiting has been a key contributor to many of the PIOs experienced in flight test programs and must be considered in the development and evaluation of the flight control system.

The importance of adhering to these six “do’s” can be appreciated by considering that 65% of the cost of a new flight control system is committed during design phase (AGARD 1987).

Flight control law design for a new aircraft, or a control system upgrade for an existing aircraft, rarely starts from a blank sheet of paper. There is usually a wealth of prior knowledge for a particular aircraft type (e.g., helicopter vs. fixed wing vs. tiltrotor) within the development group or company that establishes a launching point for each new design or upgrade. This launching point includes internal company design methods, rules of thumb, designer intuition, lessons learned, together with existing developmental building blocks and computational tools that are all brought to bear on a new project. So, it is quite common to adapt the same control law architecture for a progression of aircraft projects within a particular control system design group.

With the control law architecture or *structure* selected, the design task is focused on the selection of the *design parameters* to meet a set of (competing) requirements. Each control law structure has its particular design parameters or “tuning knobs.” For example, in a classical response-feedback structure, the design parameters are the regulator gains (proportional, integral, rate, or lead-lag compensation gains) and various feed-forward and crossfeed gains. In a model-following structure, the command model parameters (damping and natural frequency) are included as design parameters in addition to the regulator parameters.

The separation between the airframe model and controller structure on the one hand and optimization (tuning) of the design parameters on the other is illustrated in Fig. 1, and forms the basis of control system design using direct parametric optimization presented herein. The “unconstrained” design problem that is addressed by theoretically-based methods (e.g., LQR, H_∞ , eigenstructure assignment, etc.) is thus converted into a “constrained” design problem of optimizing design parameters within a fixed controller structure to meet competing requirements.

As noted by the proponents of the parametric optimization approach at the DFVLR (now DLR) (Grubel and Joos 1997), the focus of the design is correctly placed on the selection of a comprehensive set of design requirements (specifications). In this regard, a key advantage is that the specifications can now be given in physical terms of concern to flight vehicle application – for example, predicted handling-qualities, stability margins, or maneuvering capability. This provides a direct connectivity between the control system parameters and the design requirements expressed in terms of physically-meaningful metrics.

This paper reviews the key aspects of the flight control development process using multi-objective function parameter optimization and its benefits for rotorcraft control design. Typical design trade-offs are demonstrated for a simple classical response-feedback flight control system for the XV-15 tilt-rotor aircraft (lateral-directional dynamics). A more complex study based on a higher-order model of the UH-60 demonstrates how optimization methods are used for a modern multi-mode model-following system. Flight test data for the FBW UH-60 RASCAL validate the design models and predicted handling-qualities.

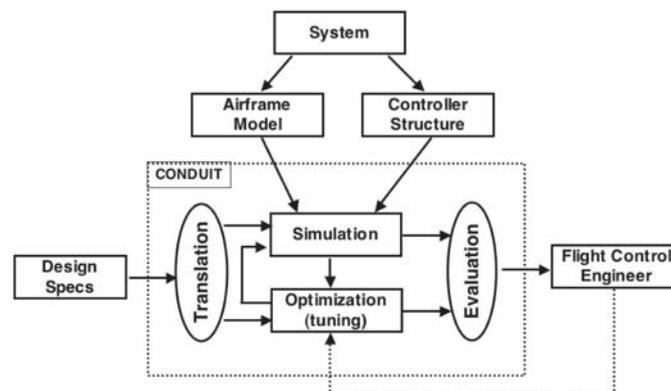


Fig. 1 Separation between the model/controller structure and design parameter tuning.

2 Rotorcraft Flight Control Design using Multi-Objective Parametric Optimization

The overall roadmap for the multi-objective parametric optimization design methodology is shown in Fig. 2. In the discussion that follows, each block of the roadmap is referred to in **bold**.

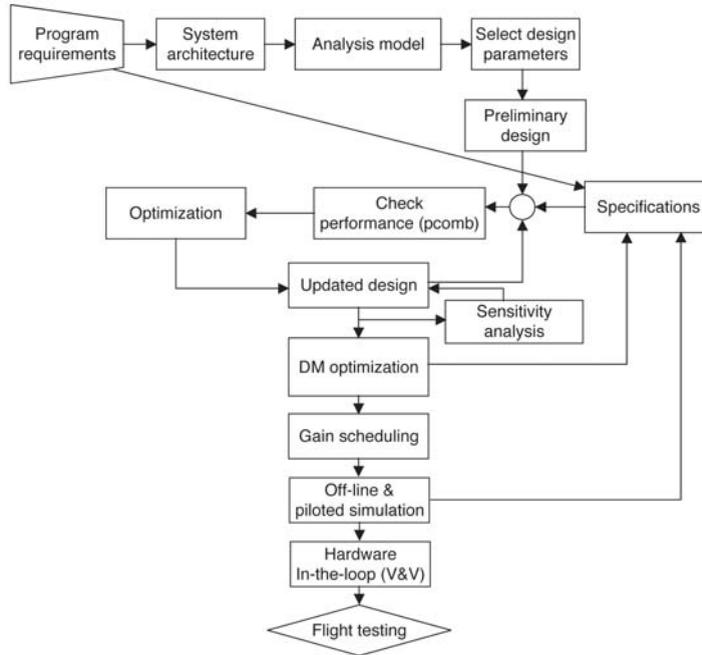


Fig. 2 Roadmap for multi-objective parametric optimization design methodology.

The entry point to rotorcraft flight control design is the **Program Requirements** and associated rotorcraft configuration, *operational environment* (i.e., weather, visibility), *operational mission*, *vision aids*, and definition of key flight conditions. The program requirements flow directly into the flight control design **Specifications**. The Aeronautical Design Standard for Handling-Qualities Requirements of Military Rotorcraft (ADS-33E-PRF, Anon 2000) contains a comprehensive set of quantitative (frequency-domain and time-domain) requirements for US military rotorcraft to ensure that satisfactory handling-qualities are achieved. Civil rotorcraft are often variants of military programs and may use ADS-33 as a design guide in addition to the (very limited) flight control design requirements set by the FAA. Flight control requirements in SAE-AS94900 (SAE 2007), the update/replacement to MIL-DTL-9490E (Anon 2008), set minimum stability margins for rigid body and structural dynamics and address uncertainty robustness issues and failures. Considering all axes and flight control modes, the design process will typically consider 50-100 specifications for each flight condition.

The **System Architecture** is defined by the flight control design requirements and provides the needed response character (e.g., rate-command/attitude-hold vs. translational-command/position-hold). At the highest level of architecture is the selection of a partial authority (mixed mechanical and electrical) flight control system vs. a full authority “fly-by-wire” architecture. Partial authority systems are well suited to upgrading legacy aircraft that often are already equipped with a stability augmentation system (SAS) that uses 10-20% of the available authority. Full authority fly-by-wire systems are the standard for new aircraft. The detailed architecture can range from classical implementations of response feedback (PID) and explicit model following, to modern multivariable concepts based on LQR and H_{∞} design methods. As stated earlier the choice depends largely on the experience at the company and specific needs of the program.

All flight control design methods require an accurate **Analysis Model** for linear and nonlinear response simulation. The model captures the dynamics responses of the bare-airframe and the various elements of the flight control system

hardware (actuators, sensors) and software (compensation elements, sampling, filters, memory blocks, and time delays). An accurate prediction of the linearized system frequency response for the loop broken at the actuator (*broken-loop response*) is especially important in the crossover frequency region (gain is approximately 0 dB), wherein small changes in magnitude and phase can have large effects on the closed-loop behavior. The key nonlinearities are associated with the actuator position limit and rate limit, and internal block diagram (port) limiters. Physics-based blade-element models such as GENHEL (Howlett 1981) or FLIGHTLAB[®] (ART 2001) can provide a good starting point especially for flight control design of new aircraft, but are subject to errors in modeling assumptions. Once the prototype test aircraft is available, the most accurate models are obtained using system identification (Tischler and Remple 2006, Ivler 2008). The complete model of the architecture, vehicle dynamics, and flight control system elements is represented in block diagram form (Simulink[®] is commonly used), with as many as 40,000 blocks and 500 dynamic states for a state-of-the-art rotorcraft.

The overall flight control system architecture typically has many tuning parameters. The gains, time constants, cross-feeds, etc. that are available to be adjusted manually or using optimization to meet the specifications are designated as **Design Parameters**. For two recent multi-mode flight control applications, these numbered some 20-40 parameters in the block diagram. There are many additional parameters that are fixed based on historical precedence or rules-of-thumb and may not contribute significantly to the specification compliance.

A reasonable **Preliminary Design** is needed as a starting point for all optimization-based design methods. Otherwise, there is no guarantee that the search engine will find a satisfactory solution that meets the many specifications. Preliminary design is usually accomplished by applying classical methods (e.g., root locus, loop shaping using Bode plots) to a simplified analysis model (often a reduced-order single-input/single-output representation). The MIMO theoretical methods (e.g., LQR) can also provide a useful starting point for key feedback gains. Then, a modern graphical user interface runs the analyses for the various specifications and allows the designer to **Check Performance** (as displayed in a horizontal bar chart or *performance comb*, **PCOMB**, Fan et al 1991) for the complete system. Trial and error testing can provide important physical insight into the key connectivities between the design parameters and specifications.

An automated **Optimization** procedure (engine) is used to tune the design parameters of the complete design problem to meet the many competing specifications. The key to automated optimization of the flight control system is the ability to determine a normalized (scaled) numerical grade for each of the metrics. The design parameters are iteratively updated by solving a multi-phase min-max optimization on the vector of requirements. This ensures that each individual specification is met and exposes trade-offs between the requirements. The **Updated Design** is rechecked against the design specifications and associated *supporting plots*. The optimization process is repeated until all of the requirements are met with a minimum level of feedback and control usage to do the job – thereby minimizing the “cost of feedback.”

Sensitivity Analysis of the converged design provides insight into the quality of convergence, allows a search of the landscape around the converged solution, evaluates the possibility of local minima, and provides an estimate of the accuracy (*Cramer-Rao Bounds*) of final design parameters. This ensures that the design problem is “well-posed” and that the solution trade-offs obtained are meaningful. Dependencies of the converged solution on initial design parameter guesses are evaluated to ensure (to the extent possible) that we have avoided local minima. The effects of modeling uncertainties are also evaluated.

While the complete design may meet all of the minimum requirements, it is usually wise to optimize for increased levels of performance to build in margins for off-nominal flight conditions and uncertainty—in exchange for a reasonable increase in control effort and reduction in stability margin. This is referred to as **Design Margin (DM) optimization**, and is achieved by progressively moving the specification boundary locations and re-optimizing the design. The locus of such solutions produces a family of perspective design points – which is a design-off trade curve. Several of the solutions from the trade-off curve are selected for final off-line and piloted evaluation.

The entire design process is repeated for other flight conditions and aircraft configurations to obtain a **Gain schedule** of flight control design parameters. Unlike fixed-wing aircraft, the dynamics of the helicopter are not a strong function of center-of-gravity or configuration. Typical look-up tables are limited to gain scheduling based on airspeed and altitude, and more recent concepts include scheduling based on external sling-load. Variations with other flight condi-

tion parameters (e.g., turn rate, rate of climb, external stores, etc.) and effects of aircraft degradation and variability are accommodated in uncertainty analyses in the off-line and piloted evaluation.

Off-line and piloted evaluation analyses evaluate the effects of modeling uncertainties, configuration/flight condition variations (not included in look-up tables), aerodynamic nonlinearities and transient effects on the completed design. Monte-Carlo analysis for off-design conditions using the linear airframe model produces a scatter plot on the specification boundaries – degradation to Level 2 handling-qualities metrics may be allowable for some of the specifications and configurations. Off-line analysis in a full-blown nonlinear blade-element simulation model provides an important assessment of robustness to likely variations away from the linearized model dynamics. Finally, piloted simulation in a full immersion simulator environment provides *assigned handling-qualities ratings* for each of the individual MTEs and full missions to be compared with the predicted ratings. These various analyses can result in modifications to the design specifications, design parameters, design conditions, selected trade-off points or all of these – and the optimization process may be repeated.

The final steps in the flight control development process include **Hardware-in-the-Loop** testing that evaluates real actuator, sensor, and flight computer hardware effects. Software verification and validation (**V&V**) is a large and important topic unto itself. **Flight testing** of handling-qualities and control performance provides the ultimate evaluation of the control laws – such as the final assigned handling-qualities ratings and stability margins. Requirements for final modification/re-optimization of the control laws will be (hopefully) limited as these are costly to complete at this last stage of development. For example, industrial experience (Pratt 2000) indicates that fixing a flight control problem that is discovered in flight test is about 50 times more expensive than fixing it at the design step. Aspects of pilot feedback from the initial flight tests of FBW UH-60 RASCAL control laws and comparison with the design requirements are discussed later in this paper.

2.1 Why is This a Good Approach?

The multi-objective parametric optimization approach to flight control system design has four key advantages:

- i. Two-way interaction. At a high level, there is a two-way interaction between the engineer and the design method. The optimization results help the engineer understand trade-offs among competing requirements and the mapping from design parameters to specifications. The engineer guides the solution through the choice of specifications and changes in the specifications boundaries and controller architecture.
- ii. Fixed and well-vetted architectures. The use of fixed and well-vetted architectures, with a great deal of corporate knowledge and experience, enhances reliability and integration as demanded for aerospace applications. The number of design parameters is selected by the engineer and thus reduces complexity of gain scheduling (as compared to an H_∞ , for example) and further enhances control system reliability, verification, and validation. Also, the application-proven architectures are easily extended to add outer loops/features and clearly assess their effect on system performance and stability (Grubel and Joos 1997).
- iii. Fullest exploitation for given architecture. This approach allows the fullest exploitation of the available design space for a given architecture and control authority. Then when an initial architecture cannot adequately achieve the requirements, the engineer can evaluate the advantage of sequentially adding design complexity, such as different control effectors, more sensors, more complex compensation schemes, and increased control authority. Comparison of optimized designs for alternative architecture allows a direct and unbiased assessment based on a common set of requirements.
- iv. Establish tuning parameters of theoretically-based methods. This design approach can be used indirectly to establish the tuning parameters of the theoretically-based methods to meet the physical requirements (e.g., bandwidth, gust rejection, etc.) while retaining the inherent theoretical properties of the methods. For example, Q and R matrix values can be selected as design parameters for an LQR design that retains “stability-margin guarantees.” This allows a clear understanding of the mapping of the tuning parameters to the requirements. The direct comparison of methods optimized to the same set of requirements resolves what would otherwise be considered as “apples vs. oranges” in many of the published case studies, since each method has its own primary design goals. In fact, Tischler et al (2002) have shown that the various methods produce little significant differences or relative advantages when optimized to a common set of requirements.

2.2 Software Tools for Flight Control Design Using Multi-Objective Parametric Optimization

Control system design based on multi-objective parametric optimization was first proposed by Kreisselmeier and Steinhauser (Kreisselmeier and Steinhauser 1979). They proposed casting the problem in the min-max framework and suggested to evaluate design trade-offs by progressively tightening the design objectives (directly akin to design margin optimization herein). In a follow-on paper in the International Journal of Control these authors (1983) provided a detailed explanation of the design approach and a flight control case study based on the F-4C. This work at the DFVLR (now DLR) evolved into the Multi-Objective Parametric Synthesis (MOPS) tool (Grubel and Joos 1997), which has most recently been used and reported on for application to the Eurofighter program (Moritz and Osterhuber 2006).

In the US, Nye made a similar case for control system design using multi-objective parametric optimization in his 1983 PhD dissertation at UC Berkeley and embodied in the DELIGHT package (Nye 1983). This work was further evolved at the University of Maryland into CONSOL-OPTCAD by Fan et al (1991). There are many similarities between the German and US implementation, both using the same core sequential quadratic programming algorithms and a multi-stage min-max optimization. However, there are some important differences in application strategy. Under a long-term Science and Technology effort that evolved from research collaborations with the University of Maryland team, the US Army Aeroflightdynamics Directorate (AFDD) has developed the Control Designer's Unified Interface (CONDUIT[®], Tischler et al 1999) for aircraft and rotorcraft flight control design analysis and optimization. Some key features of CONDUIT[®] are:

- Comprehensive analysis and optimization environment in a “one-stop shop” that integrates many aspects of design.
- Integrated and user-friendly graphical interface that is focused on the aircraft and rotorcraft applications (Fig. 3).
- Graphical libraries of all key rotorcraft and fixed-wing handling-qualities and flight control requirements.
- Seamless integration with Simulink[®] for control system and airframe modeling.
- Seamless integration with system identification tool CIPHER[®].
- Specialized tools for analysis of control system response and robustness to uncertainty.
- Specialized tools for assessing specification compliance based on flight test data.
- Theoretical accuracy metrics based on numerical sensitivity analysis to determine accuracy of design parameters and level of parameter correlation.
- Batch job and databasing of results for large-scale industrial applications.
- Comprehensive User's Manual (Fig. 4).

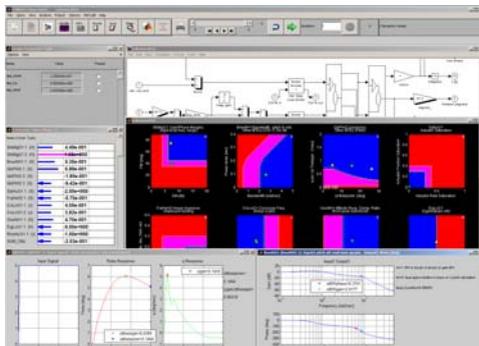


Fig. 3 CONDUIT[®] graphical interface.



Fig. 4 CONDUIT[®] User's Manual.

3 Multi-Phase Design Optimization Method

This section presents the key elements of rotorcraft flight control design using multi-objective function multi-phase optimization. The key specifications that must be considered in the design are reviewed first. Then we explain how individual metrics are normalized and “scored” based on the handling-qualities boundaries as a precursor to numerical optimization using the min-max cost function.

3.1 Need and Challenge of Numerical Optimization of Flight Control Design

The complete control system block diagram for a modern rotorcraft is typically very complex and is comprised of many layers and modes. In a recent rotorcraft fly-by-wire flight control application, this amounted to some 450 dynamic states and 34,000 blocks. There are many tunable design parameters (35 in the recent application) that are “scattered” throughout the block diagram layers (i.e., beyond the simple attitude and rate gains) and affect the overall response in subtle ways. So, the available “wiggle room” in the design is very small indeed and will not be well exploited when the design parameters are tuned manually or constrained based on a priori rules-of-thumb. In many applications, the inner-loop attitude (and rate) gains are held constant for the higher levels of augmentation – such as translational rate command and position hold – thereby requiring compromises that further complicates manual tuning. So, inevitably we need to consider analysis and optimization trade-offs of the complete MIMO problem and associated specifications. The many and varied specifications (86 for the hover/low-speed flight modes in the recent application) are given in different units and relative scaling, they compete with one another, do not reflect distinct behavior, and are affected by more than a single design parameter. Taken together, flight control design presents a highly non-orthogonal and non-convex problem in numerical optimization – finding a good solution is very challenging and can be computationally intensive.

The approach herein uses feasible sequential quadratic programming (FSQP) to solve a min-max optimization of the vector of multiple objective functions (specifications) (Fan et al 1991). The core (QL) algorithm is due to Schittkowski (2005). The solution is achieved by dividing the problem into 3 stages. Each stage addresses specs that achieve increasing refinement of the flight control design. The final solution is one that meets all of the specifications with minimum control usage, thereby minimizing the “cost of feedback.” This approach has been found to be very well suited to the difficult problem of rotorcraft multi-mode flight control system design.

3.2 Selection of Design Specifications

As discussed earlier, the “driver” of the parametric optimization design process is the set of system requirements or specifications. The ADS-33E-PRF performance specification (Anon 2000) contains a comprehensive set of quantitative requirements for US military rotorcraft that is a good starting point to ensure that satisfactory handling-qualities are achieved. The ADS-33 document defines response type to piloted inputs (e.g., angular rate, attitude response, translational rate) as a function of the *usable cue environment* (UCE). The response type sets the loop architecture and command model order (for a model following implementation). The rotorcraft *category* (e.g., utility, vs. attack) sets the *mission task elements* (MTEs) that comprise the overall mission. Then boundary locations are provided as a function of mission task elements (MTEs), such as “target acquisition and tracking” (most aggressive) to “full attended operations” (least aggressive). There are several areas where the current version of ADS-33 is deficient – most notably in criteria for disturbance rejection and actuator saturation. The US Army Aeroflightdynamics Directorate (AFDD) has developed/adapted specifications to address these important aspects that are primary drivers and limitations of control system design (Harding et al 2006, Blanken et al 2008). Flight control requirements in SAE-AS94900 (SAE 2007), the update/replacement to MIL-DTL-9490E (Anon 2008) set minimum stability margins for rigid body and structural dynamics and address uncertainty robustness issues and failures. The key requirements that drive the flight control design for hover are summarized here based on evaluation for: piloted inputs, disturbance rejection, stability margins, robustness, and control usage.

Handling-qualities criteria for piloted inputs (applied to stick):

- Bandwidth (ADS-33). Short-term small amplitude criterion that characterizes the speed of response to pilot stick inputs – defined from a Bode plot of attitude response.

- Quickness (ADS-33). Moderate amplitude attitude change criterion is identical to bandwidth for small inputs and characterizes the required response under actuator saturation – not very useful or practical for control system design, better used in post-design analysis.
- Damping ratio (ADS-33). Based on time response evaluation of overshoot or damping characteristic for a pilot input.
- Inter-axis coupling (ADS-33). Based on a comparison of the off-axis vs. on-axis time response for pilot input. For aggressive maneuvering this characteristic is evaluated based on the off-to-on axis frequency response.
- Translational rate command time constant (ADS-33). Speed of response for TRC mode is based on the rise time for pilot inputs.
- Model following accuracy (AFDD). Mismatch cost function based on comparison of frequency responses for command model and aircraft to ensure that satisfactory model following is achieved (Tischler and Remple 2006).
- Open-loop Operating Point (OLOP, DLR). Frequency-domain spec that assesses possibility of limit cycle response due to actuator rate saturation following a large pilot input (Duda 1997).

Disturbance rejection:

- Settling time (ADS-33). Ensures that aircraft response returns quickly to trim following an upset based on a pulse input to the actuator.
- Disturbance rejection bandwidth (AFDD). The disturbance rejection bandwidth (Harding et al 2006, Blanken et al 2008) is directly analogous to pilot response bandwidth, only the excitation is at the output. The criteria is based on the frequency for -3dB on the magnitude frequency response.
- Disturbance rejection peak (AFDD). Ensures against a large peak (overshoot) to disturbance inputs as determined from the peak magnitude of the disturbance frequency response.

Stability and stability margin criteria:

- Eigenvalue (generic). All closed-loop poles must have eigenvalues in the left-half plane (or within a specified tolerance).
- Stability margin (SAE 2007). Gain and phase margin requirements for the broken-loop response.
- Nichols stability margin (Magni et al 1997, Moritz and Osterhuber 2006). Defines stability margin in terms of an exclusion zone on the broken-loop gain-phase response (Nichols plot).

Robustness:

- Performance robustness (Yaniv 1999). Evaluates the variation in closed-loop response of off nominal flight/mass configuration cases relative to the nominal design point and quantifies this in terms of the LOES mismatch cost.
- Stability robustness (Magni et al 1997, Moritz and Osterhuber 2006). Ensures that the off-nominal conditions maintain a (reduced) stability margin in terms of the Nichols margin exclusion zone.

Control Usage:

- Actuator RMS for pilot inputs (AFDD). Evaluates actuator RMS of closed-loop system for inputs at the pilot stick.
- Actuator RMS for disturbance inputs (AFDD). Evaluates actuator RMS of closed-loop system for disturbance inputs to the actuator.
- Crossover frequency (generic). 0dB frequency as determined from broken feedback-loop response. This is a good measure of feedback system bandwidth.

3.3 Numerical Scores for the Specifications

The central requirement for automated design optimization is a capability to return numerical scores of system performance for a given set of design parameter values. In the current case, the requirement is to convert the handling-qualities or flight control performance as plotted on each specification (i) to a *numerical score* (f_i). The relative scaling (or weighting) of the scores is critical to the optimization process. So, how can a score for piloted control bandwidth (e.g., 2.2 rad/sec) be compared with the score for phase margin (e.g., 48 deg) or settling time (e.g., 4.7 sec)?

3.3.1 General Concept of Normalization

Nye and Tits (1986) proposed an excellent solution to achieving a uniform normalization of the numerical scores. The values of a specification corresponding to two *levels of satisfaction* are defined as *good* and *bad* values. Then, the *normalized* specification rating (f_i) can be defined as:

$$\hat{f}_i \equiv \frac{f_i - \text{good}_i}{\text{bad}_i - \text{good}_i} \quad (1)$$

where we see that a normalized value of $\hat{f}_i = 0$ always corresponds to a good score and a normalized value of $\hat{f}_i = 1$ always corresponds to a bad score. Based on this scheme, the optimization will always work to decrease the normalized scores of the specifications to a value of $\hat{f}_i \leq 0$. Successful design optimization using numerical min-max methods requires that the choice of good and bad values obey the *uniform satisfaction/dissatisfaction rule* (Nye and Tits 1986):

“having all of the various specifications achieve their corresponding good values should provide the same level of ‘satisfaction’ to the designer for each, while achieving the bad values should provide the same level of ‘dissatisfaction.’”

Nye and Tits (1986) point out that a key advantage of this approach to normalization is that the good and bad values have the same physical units as the specifications themselves. Now, we need to define good and bad values for all of the specifications.

3.3.2 Normalization of numerical scores based on predicted HQR ratings

The handling-qualities/flight control specification boundaries that separate the Level 1 handling-qualities region from the Level 2 handling-qualities region (*Level 1 boundary*) and Level 2 handling-qualities region from the Level 3 handling-qualities region (*Level 2 boundary*) provide for a natural and automatic definition of the good and bad values. For example, consider the bandwidth/phase-delay handling-qualities data boundaries of Fig. 5 as gathered from flight test studies (Mitchell et al 1989). For each flight control configuration (with an associated bandwidth (ω_{BW}) and phase delay ($\tau_{p\phi}$)), pilots flew a variety of mission task elements and provided a rating using the Cooper-Harper scale as shown in the figure. Then boundaries (splines) were fared through the data to separate configurations into 3 regions.

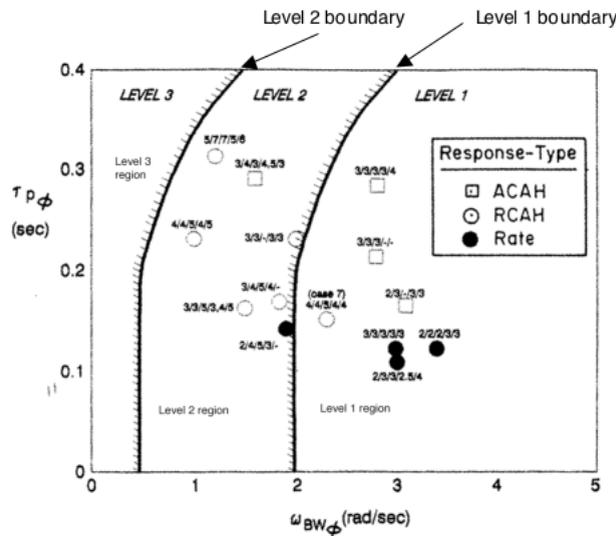


Fig. 5 Bandwidth/phase-delay handling-qualities data and specification boundaries.

The Level 1 boundary marks the limits of the *Level 1 region* that provides for a handling-qualities rating HQR = 1-3 (“satisfactory without improvement”). All configurations lying exactly on the Level 1 boundary will receive a pre-

dicted $HQR=3.5$, which we can consider as “good” – and thus can be considered as having an equal level of designer satisfaction. The Level 2 boundary marks the limits of the *Level 2 region* that provides for a handling-qualities rating $HQR = 4-6$ (“deficiencies warrant improvement”). All configurations lying on the Level 2 boundary will receive a predicted $HQR=6.5$, which we can consider as “bad” – and thus can be considered as having an equal level of designer dissatisfaction.

The handling-qualities requirements document (ADS-33) is comprised of a large set of such individual specifications. Each and every specification (e.g., bandwidth/phase-delay, damping ratio, settling time, etc.) must lie in the Level 1 region for the pilot to rate the overall aircraft handling-qualities as “Level 1” – satisfactory as is. If any single specification falls within the Level 2 region (with the rest being in the Level 1), the pilot would rate the overall aircraft handling-qualities as “Level 2” – improvement is warranted. Thus, when scored as a predicted HQR , all specifications have an equal weighting in terms of pilot acceptance of the aircraft handling-qualities – and so the uniform satisfaction/dissatisfaction rule is met. We can assign the Level 1 boundary as corresponding to the locus of “good” values for all specs. Similarly we assign the Level 2 boundary as corresponding to the locus of “bad” values for all specs. The geometric distance between the Level 1 and Level 2 boundaries always corresponds to 3 HQR rating points – thus providing for a common and automatic scaling for all of the specs. The raw score f_i of Eq. (1) is simply the predicted HQR Level as obtained based on the geometric location of the configuration relative to the Level 1 and Level 2 boundaries – this score is obtained by means of the *distance algorithm*. Thus in Eq. (1), a “good” raw value will be $good_i \equiv 1.0$ for all specs and a “bad” raw value will be $bad_i \equiv 2.0$ for all specs.

3.3.3 Distance Algorithm

The key approach of the distance algorithm stems from the desire to conceptualize a specification as lines of constant numerical ratings. By using the gradient information defined from the normalized scaling of the Level 2 region, a “scoring” surface can be generated as depicted in Fig. 6. Ultimately, for visualization purposes, the specification can be viewed as a contour map of scoring values. The handling-qualities and flight control specifications implemented in CONDUIT[®] can be categorized into two distinct types: point type specifications and line type specifications. A point specification is one that evaluates to a single design value (Fig. 7). This design value is scored and a single numerical rating is assigned. On the other hand, a line specification is one that evaluates to a vector of data on the specification (Fig. 8). Each value is scored and the worst (maximum) score is returned as the final rating.

An example of a more complex boundary geometries is the damping ratio spec of Fig. 9. For these types of specs, a sophisticated and generic surface fitting algorithm is required to compute the rating. In CONDUIT[®], the biharmonic spline interpolation algorithm introduced by Sandwell (1987) is used to determine the specification scoring for this family of specs, as illustrated in the example of Fig. 9.

The distance algorithm provides a unified and consistent mechanism for determining the numerical ratings of all specs, which are the essential weighting elements to the optimization process.

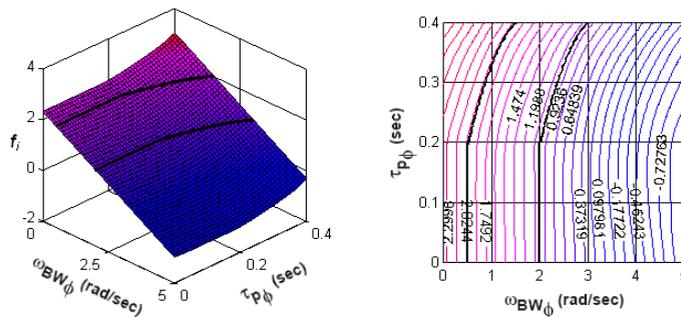


Fig. 6 Surface plot and contour map for the bandwidth/phase-delay specification.

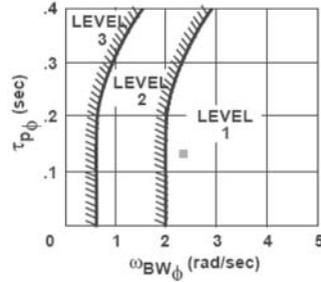


Fig. 7 Example of a point spec (bandwidth/phase-delay).

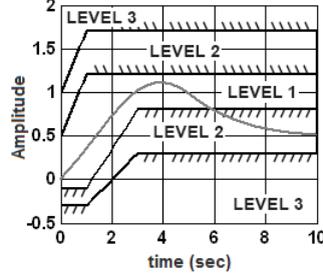


Fig. 8 Example of a line spec (time-response envelope).

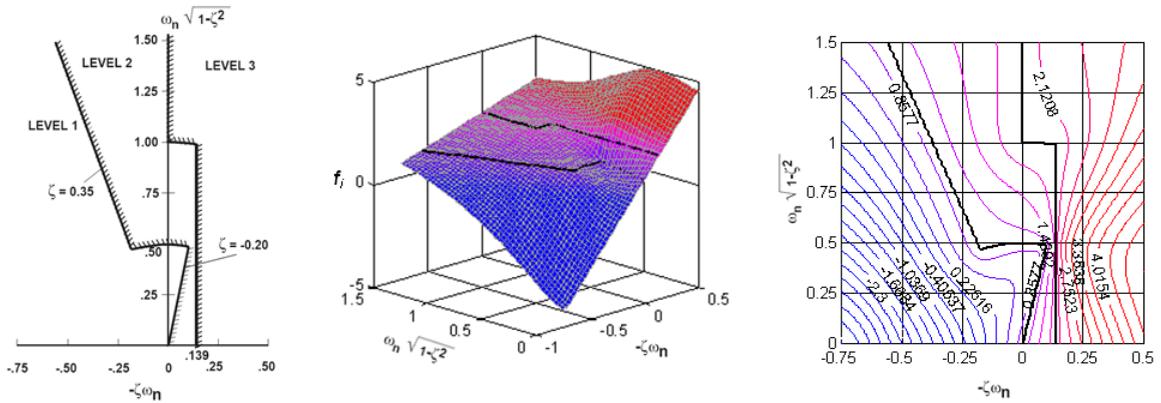


Fig. 9 Example of a generic, non-uniform spec (damping ratio requirement).

3.4 Numerical Optimization of the Design

The overall objective of the numerical optimization process is to tune the vector of design parameters (\mathbf{dp}), n_{dp} in number, until the specifications (n_{spec} in number) are all in the Level 1 region, with minimum *over-design* – that is, minimum use of the available control power as needed to meet the requirements.

3.4.1 Multi-Objective Function Parametric Optimization using the FSQP solver

Feasible Sequential Quadratic Programming (FSQP, Fan et al 1991) is one of several approaches to multi-objective function parametric optimization. This constrained parametric optimization solver is well suited to the flight control system design problem and is used herein and in CONDUIT[®]. FSQP allows the various specifications to be grouped and achieved in phases according to their relative importance. The FSQP search direction is based on a quadratic approximation of performance measures and a linear approximation to the constraint boundary. At the core of the FSQP solver is the convex quadratic programming (QL) (Schittkowski 2005) and Powell (1983) search algorithms, the latter also used in the development of the MOPS tool (Kreisselmeier and Steinhauser 1983, Grubel and Joos 1997). FSQP has been found to be a fast and efficient solver for many complex flight control design problems.

The handling-qualities and flight control requirements are designated by the designer as belonging to one of four possible categories. The first category is referred to as comprising the *hard* (“H”) constraints ($f_{hard_1}, f_{hard_2}, \dots, f_{hard_{n_{hard}}}$). These are the few (n_{hard}) requirements that the designer decides must be met as an initial basis for the design and enforced as the design optimization progresses. The second category, referred to as *soft* (“S”) constraints ($f_{soft_1}, f_{soft_2}, \dots, f_{soft_{n_{soft}}}$), comprises the many (n_{soft}) handling-qualities and flight control requirements. These should

be met to the extent possible. The designer will want to assess if these specifications can be met within the block diagram architecture selected, or if for example additional feedbacks are needed to allow certain requirements to be met. We may accept a design in which some requirements only reach the Level 2 region if the reward for the other requirements is great enough. The third category, referred to as *check-only* (“C”) *constraints* are to be monitored during the design process, but not enforced as part of the optimization algorithm. The performance measures or *objectives* (“O”) form the final category ($f_{obj_1}, f_{obj_2}, \dots, f_{obj_{n_{obj}}}$) and are the (n_{obj}) relative measures of control system effort which are combined into a single *summed objective* (“J”).

Using these requirements groups, the FSQP solver updates the design parameters in three optimization phases. In *Phase 1* only the hard constraints are considered, while in *Phase 2* the soft constraints are considered while still enforcing the hard constraints. Finally in *Phase 3*, the performance measures are minimized while enforcing the hard and soft constraints. Central to the optimization process is the definition of an appropriate overall normalized *cost function* (f) for each Phase based on the individual normalized numerical scores, f_i .

3.4.2 Maximum Value Cost Function

The maximum-value cost function is a good choice for numerical design optimization to achieve satisfaction of the many handling-qualities and flight control specifications in Phase 1 and 2:

$$\tilde{f} = \max\{\tilde{f}_1, \tilde{f}_2, \dots, \tilde{f}_{n_{spec}}\} \quad (2)$$

The numerous handling-qualities and flight control design specifications are design *constraints* – they are either satisfied or not. A satisfactory design requires that *each and every* handling-qualities and flight control specification be in the Level 1 region (i.e., raw scores $f_i \leq 1.0$ or normalized scores $f_i \leq 0$ for *all* i). Further, the pilot opinion of the aircraft response is determined by the worst performing metric. By adopting the max value (i.e., *worst case specification*) as the cost function, the actual multi-objective function optimization problem is thus replaced by a min-max (or *minimax*) scalar solution. The design is improved by pushing the worst case specification toward the Level 1 region, until a design solution is found that satisfies all of the specs (constraints). This solution approach adapts well to a large number of non-orthogonal and competing (i.e., *tightly coupled*, Boyd and Barratt 1991) requirements. Design problems involving tens of specifications have been successfully solved using CONDUIT®.

3.4.3 Phase 1 – Satisfying the Hard Constraints

The hard (“H”) constraints are requirements that must be satisfied (Level 1) and enforced regardless of the ultimate impact on the remaining specs. The absolute stability spec (eigenvalues in the left-half plane) and stability margin spec (gain and phase margins) for all of the control axes are generally selected as hard constraints. This ensures that the remaining handling-qualities metrics will be defined and that the design space will be well-behaved for the remainder of the design optimization. In *Phase 1*, the cost function is defined as:

$$\tilde{f} = \max\{\tilde{f}_{hard_1}, \tilde{f}_{hard_2}, \dots, \tilde{f}_{hard_{n_{hard}}}\} \quad (3)$$

and the FSQP algorithm searches over the vector of design parameters (\mathbf{dp}) to solve the optimization problem:

$$\min_{\mathbf{dp}} \left(\max\{\tilde{f}_{hard_1}(\mathbf{dp}), \tilde{f}_{hard_2}(\mathbf{dp}), \dots, \tilde{f}_{hard_{n_{hard}}}(\mathbf{dp})\} \right) \quad (4)$$

As the design parameter guesses are updated, the hard constraints are progressively moved toward the Level 1 region. The solution remains in Phase 1 as long as at least one of the hard constraints fails to meet Level 1 requirements. The solution progress from Phase 1 to Phase 2 as soon as *all* hard constraints are in the Level 1 region, if this is possible within the defined control system architecture. If the hard constraints cannot be met the solution will stop at the

design result where the worst hard spec is closest to the Level 1 boundary. Soft constraints and performance metrics are ignored entirely in Phase 1.

3.4.4 Phase 2 – Satisfying the Soft Constraints

In Phase 2, the cost function is defined as:

$$\hat{f} = \max \left\{ \hat{f}_{\text{soft}_1}, \hat{f}_{\text{soft}_2}, \dots, \hat{f}_{\text{soft}_{n_{\text{soft}}}} \right\} \quad (5)$$

and the FSQP algorithm searches over the vector of design parameters (\mathbf{dp}) to solve the optimization problem:

$$\min_{\mathbf{dp}} \left(\max \left\{ \hat{f}_{\text{soft}_1}(\mathbf{dp}), \hat{f}_{\text{soft}_2}(\mathbf{dp}), \dots, \hat{f}_{\text{soft}_{n_{\text{soft}}}}(\mathbf{dp}) \right\} \right) \quad (6)$$

subject to the condition that all of the hard (Phase 1) constraints remain satisfied ($\hat{f}_{\text{hard}_i} \leq 0$).

So for example, in Phase 2, the solver will only search among the stable designs. The many handling-qualities and flight control requirements for all of the various axes and modes are selected as soft (“S”) constraints. As the design parameter values are updated, the soft constraints are progressively moved toward the Level 1 region, while the hard constraints are kept within the Level 1 region.

The solution remains in Phase 2 as long as at least one of the many soft constraints fails to meet Level 1 requirements. Here there are difficult trade-offs to be worked between the tightly coupled dynamic response metrics and many potential trial guesses are made before each design update. The solution progresses from Phase 2 to Phase 3 as soon as *all* soft constraints are in the Level 1 region, again if this is possible within the defined control system architecture. If the soft constraints cannot be met, the solution will stop at the design result where the worst soft constraint is closest to the Level 1 boundary. The performance metrics are treated as soft constraints when the solution is in Phase 2.

If a design is found that meets all hard and soft constraints, then all of the handling-qualities and control design requirements are achieved and Phase 3 has been reached. At this point we have a *feasible* design, also called an *achievable* design solution (Boyd and Barratt 1991) or *attainable* design solution (Maciejowski 1989). The solution is not unique as many such feasible solutions can usually be found. Some design metrics may be just inside the Level 1 region and others deep inside this region, so we may be able to improve one or more metrics without incurring a trade-off or crossing over the Level 1/Level 2 boundary.

3.4.5 Phase 3 – Selecting the best solution from among the feasible sets.

Given two feasible designs with all metrics in the Level 1 region, how is the better one selected? The design with the smaller *performance measure* is judged as the better choice. Good performance measures for flight control are the actuator usage for pilot stick inputs and broken-loop crossover frequency (i.e., feedback system bandwidth).

There are many reasons for wanting to minimize these performance measures:

- 1) Minimize potential for *spillover* of control energy to high-frequency modes that can excite unmodelled structural dynamics (Bryson 1994).
- 2) Minimize structural coupling to extend fatigue life of aircraft components (Pratt 2000, Rozak and Ray 1998).
- 3) Improve stability margin robustness to uncertainty in aircraft dynamics.
- 4) Minimize closed-loop sensitivity to noise.
- 5) Minimize potential for actuator saturation and associated susceptibility of aircraft to pilot-induced oscillation (Duda 1997, AGARD 2000). This will be an especially important consideration for limited authority control system implementations.

For the same reasons as above, achieving a design that meets the requirements with minimum control usage is also the stated goal by Kreisselmeier and Steinhauser (1983) which is the basis for the MOPS design tool (Grubel and

Joos 1997). Horowitz (1963) refers to items i-iv as the “cost of feedback.” Good classical design practice also seeks to minimize this cost by using only enough feedback to achieve the design goals.

Experience with CONDUIT® indicates that the best performance measures are the:

- RMS value for actuator position response to piloted inputs (one value for each loop; 4 in total).
- Crossover frequency value for the feedback loop broken at the mixer input (one value for each loop; 4 in total) resulting in a total of 8 performance measures for the 4 control loops (roll, pitch, yaw, heave).

Performance measures or *objectives* differ from the hard and soft constraints in that here the relative value in the Level 1 region always matters – as long as all of the handling-qualities constraints are satisfied, a lower performance metric (i.e., less actuator usage or lower crossover frequency) is always better. In *Phase 3*, the cost function is defined as:

$$\hat{f} = \left(\frac{1}{n_{\text{obj}}} \right) (\hat{f}_{\text{obj}_1} + \hat{f}_{\text{obj}_2} + \dots + \hat{f}_{\text{obj}_{n_{\text{obj}}}}) \quad (7)$$

and the FSQP algorithm searches over the vector of design parameters (**dp**) to solve the optimization problem:

$$\min_{\mathbf{dp}} \left(\left(\frac{1}{n_{\text{obj}}} \right) (\hat{f}_{\text{obj}_1}(\mathbf{dp}) + \hat{f}_{\text{obj}_2}(\mathbf{dp}) + \dots + \hat{f}_{\text{obj}_{n_{\text{obj}}}}(\mathbf{dp})) \right) \quad (8)$$

subject to the condition that all of the hard (Phase 1) and soft (Phase 2) constraints remain satisfied ($\hat{f}_{\text{hard}_i} \leq 0$; $\hat{f}_{\text{soft}_i} \leq 0$).

The choice of the (scaled) *summed* objective (“J”) cost function in Phase 3 rather than the max value allows some individual performance measures to be reduced even if the worst one cannot be reduced any further. As an example, consider a design solution where the pitch actuator usage has been reduced to a minimum value that allows the flight control design to *just meet* Level 1 handling-qualities requirements – further reductions in pitch actuator usage will cause a violation of one (or more) of the pitch specifications. If the pitch actuator usage exceeds that of roll, a min-max solution would stop at this point, even if roll actuator usage could be reduced further. Since the pitch and roll performance measures are largely decoupled (or nearly *independent*, Boyd and Barratt 1991), the use of the summed objective will drive the solution to minimize roll actuator usage further until the roll constraints are also just met. By using the scale factor ($1/n_{\text{obj}}$), the Phase 3 cost function is taken as the average value of the component objective functions (“J”), and so will have the same interpretation of good and bad values.

3.4.6 Converged Phase 3 Design is a Pareto Optimum Solution

Over some range of design parameters values for a design in Phase 3, actuator usage can be reduced and the handling-qualities metrics will not penetrate the Level 1/Level 2 boundary – reduced actuator usage may in fact provide improvements in some of the handling-qualities metrics (e.g., OLOP, Sec. 4.2). However, at some point further reductions in pitch and/or roll actuator usage will result in an *unachievable solution* – one that does not meet one or more of the handling-qualities constraints. For example, further reductions in actuator usage will cause a violation of the disturbance rejection requirements.

This solution at the boundary between the achievable and the unachievable solution space is referred to as a *Pareto optimum* solution as explained by Boyd and Barratt (1991) and Maciejowski (1989) and depicted in the sketch of Fig. 10. At this solution, further improvements in any one criteria can only be achieved by sacrificing (degrading) one or more other criteria. By moving the specifications boundaries, we can explore the entire *trade-off* curve or family of Pareto optimum solutions, as also depicted in Fig. 10. Boyd and Barratt suggest that the Pareto optimal solutions are wise design choices, and that we need only to evaluate (interactively) the possible solutions and choose the best from among these.

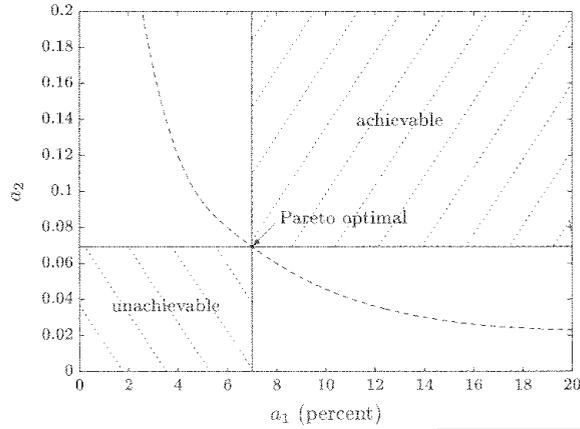


Fig. 10 Pareto optimum solution (reproduced from Boyd and Barratt 1991).

In CONDUIT[®], the boundaries are moved relative to the specified locations (e.g., as given in ADS-33) either directly by “dragging” the specification boundary splines, or equivalently by definition of a *design margin* (DM). A positive design margin ($DM > 0$) moves the Level 1/Level 2 boundary a fixed percentage into the Level 1, thus building in a margin with some over-design for robustness to modeling uncertainty, or improved performance relative to the ADS-33 minimums. The design margin is easily implemented by defining the good “raw” value:

$$\text{good}_i \equiv 1.0 - DM \quad (9)$$

A design margin of 10% ($DM = 0.1$), corresponding to dragging the Level 1/Level 2 boundary 10% of the width of the Level 2 region, is achieved by setting the “good” raw values to $\text{good}_i \equiv 0.9$. Zero design margin $DM = 0$ results in a design that just meets the Level 1 requirements, with no margin for uncertainty. Ultimately, the design engineer should select a design based on an evaluation and comparison of the design solutions and trade-offs, including a careful examination of the supporting analysis plots.

Sections 4 and 5 that follow present detailed case studies of design optimization and exploration of design trade-offs. The first case study is a simple example of a PID controller design for the roll-yaw hover dynamics of the XV-15 tilt-rotor aircraft. The second case study presents control system optimization and trade-offs for a complex multi-mode model-following system for the UH-60Mu helicopter.

4 Design Trade-offs for a Simple Application

A simple notional flight control problem based on the lateral-directional dynamics of the hovering XV-15 tilt-rotor aircraft (Fig. 11) was used to understand key flight control trade-offs.



Fig. 11 XV-15 in hover.

The hovering XV-15 is an ideal case for a simple helicopter flight control study because the bare airframe lateral/directional dynamics are decoupled from the longitudinal dynamics (because of the side-by-side rotor configuration) allowing for simplicity, but still exhibiting the unstable hovering cubic characteristic of a helicopter.

4.1 Analysis Model

An accurate model of the flight dynamics is needed for flight control design. The state-space model of the XV-15 was identified directly from flight data using the CIFER[®] system identification technique (Tischler and Remple 2006):

$$\begin{bmatrix} \dot{v} \\ \dot{p} \\ \dot{r} \\ \dot{\phi} \end{bmatrix} = \begin{bmatrix} Y_v & Y_p & Y_r & g \\ L_v & L_p & L_r & 0 \\ N_v & N_p & N_r & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} v \\ p \\ r \\ \phi \end{bmatrix} + \begin{bmatrix} Y_{\delta_{ail}} & Y_{\delta_{rud}} \\ L_{\delta_{ail}} & L_{\delta_{rud}} \\ N_{\delta_{ail}} & N_{\delta_{rud}} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \delta_{ail} \\ \delta_{rud} \end{bmatrix} \quad (10)$$

and matches the frequency-response flight data closely for frequencies up to 10 rad/sec. The identification process also provides uncertainty estimates for the frequency-response data (*random error*, ϵ_r) and for the individual state-space parameters (*Cramer-Rao bounds*, (\overline{CR}_i)).

A simple XV-15 block diagram was developed to demonstrate fundamental flight control design. This notional block diagram is meant to be realistic, but is not representative of the actual stability augmentation system on the aircraft. These example control laws consist of standard PID control in a response-feedback type configuration. The components of the block diagram include the command filter, stick sensitivity gain, lateral stability augmentation system (SAS), directional SAS, actuator dynamics, sensor dynamics, and filters.

The lateral SAS has an attitude-command/attitude-hold (ACAH) response type. The roll controller has roll attitude and integral feedback, as well as roll rate feedback in a PID-type controller architecture as shown in Fig. 12. Since this study focused on the short term response, the integral gain was constrained as a ratio of the proportional gain based on classical design rules of thumb ($(k_{I\phi}/k_{\phi}) = 0.2\omega_c$), rather than being a free design parameter. This ensures that the integrator gain is effective, without overly degrading the phase margin. The directional controller is a rate-command/attitude-hold (RCAH) response type architecture with proportional gains on yaw rate and heading angle.

The actuator dynamics are modeled as second-order systems with rate and position limiting. A sampling delay is included in the sensor dynamics block to account for the delay in digital sampling. Noise filters on the measurements are included in the block diagram.

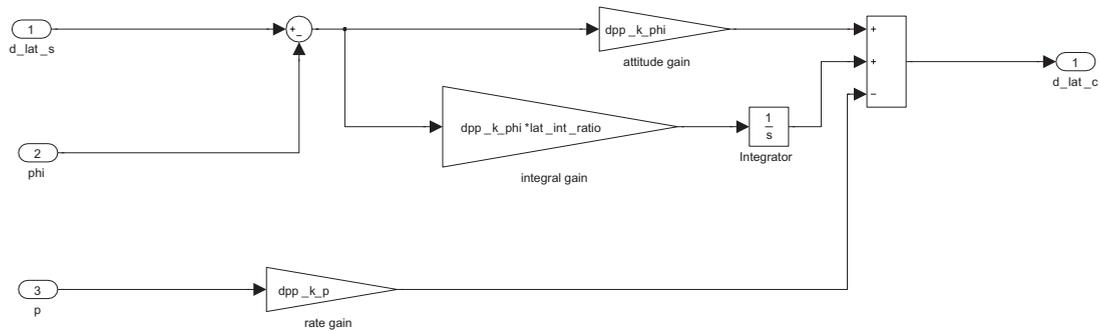


Fig. 12 Lateral stability augmentation

4.2 Summary of Specifications for XV-15

Table 1 shows the specifications chosen for the XV-15 optimization. The choices for Hard Constraints, Soft constraints, and Summed Objectives are also shown in the Table. To summarize the methodology: specifications that involve stability criteria are optimized as hard constraints (H), specifications that involve handling qualities are optimized as soft constraints (S), and specifications that involve performance/actuator usage are optimized as summed objectives (J).

Table 1 XV-15 Specifications

Spec Name	Description	Axes	Type
EigLcG1	Eigenvalues	All	Hard
StbMgG1	Stability Margin (9490)	Roll, Yaw	Hard
NicMgG2	Nichols Margin	Roll, Yaw	Hard
BnwAtH1	Roll Bandwidth, All other MTE s - UCE =1 and Fully Attended Operations (ADS-33)	Roll	Soft
BnwYaH2	Yaw Bandwidth, other MTE s (ADS-33)	Yaw	Soft
EigDpG1	Generic Damping Ratio		Soft
DstBwG1	Disturbance Rejection Bandwidth	Roll, Yaw	Soft
DstPkG1	Disturbance Rejection Magnitude Peak	Roll, Yaw	Soft
OlpOPG1	Open Loop Operating Point Rate Limit Saturation Spec.	Roll, Yaw	Soft
CrsLnG1	Crossover Frequency	Roll, Yaw	Summed Objective
RmsAcG1	Actuator RMS	Roll, Yaw	Summed Objective

Open-Loop Operating Point Specification (OlpOPG1). An interesting specification that addresses actuator rate saturation and is well suited for use in the optimization is the open loop operating point (OLOP) for Category II pilot-induced oscillation (PIO) criteria (Duda 1997). Category II PIOs are characterized by rate-or-position limiting. With linear analysis methods, such as based on the Simulink[®] “linmod” function, non-linear effects of rate and position limiting are ignored. It is possible to push the design well beyond the limits of the real aircraft in the linear world (where limiters are ignored), but when implemented on a nonlinear system the effects of position and rate limiting become apparent. Rate limiting is particularly dangerous as it has been associated with most recent PIO events. The OLOP specification is based on frequency domain describing function concepts, but does not require the user to apply the describing function technique.

The details of the OLOP specification are given by Duus (2000). An example analysis of the OLOP criteria for the XV-15 in the directional axis is given in Fig. 13(a) and Fig. 13(b). First, the onset frequency of the rate limiting is determined as shown in Fig. 13(a). The broken-loop response is then calculated with the loop broken at the actuator as shown in Fig. 13(b). The amplitude and phase of the broken-loop are recorded at the onset frequency (1.28 rad/s) and then plotted against the criteria. In this example case, the point, OLOP = 5.25dB, -97.27 deg, is in the “PIO unlikely area” of the OLOP specification shown in Fig. 14.

The implementation of OLOP shown above is for the aircraft-FCS loop alone, without the pilot model that is used in some of the references. Using the specification in this way ensures that the aircraft system does not become unstable when the rate limiter is initially activated (for unstable bare airframes). It does not consider the pilot’s reaction to the unexpected change in dynamics from rate limiting, which can sometimes cause PIO as well. However, we have found that the FCS-only implementation of OLOP is a good compromise for the time being, as the use of a pilot model tends to severely (and unrealistically) limit predicted performance.

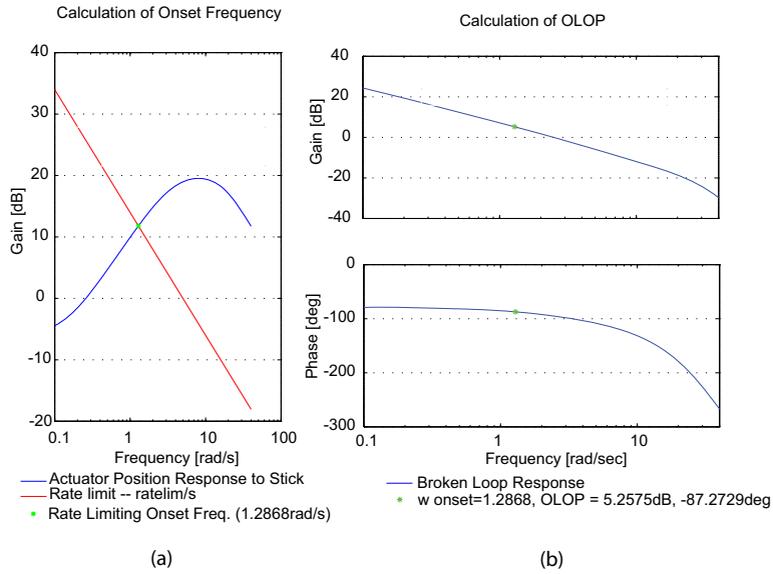


Fig. 13 Calculation of rate limiting onset frequency for XV-15 yaw axis.

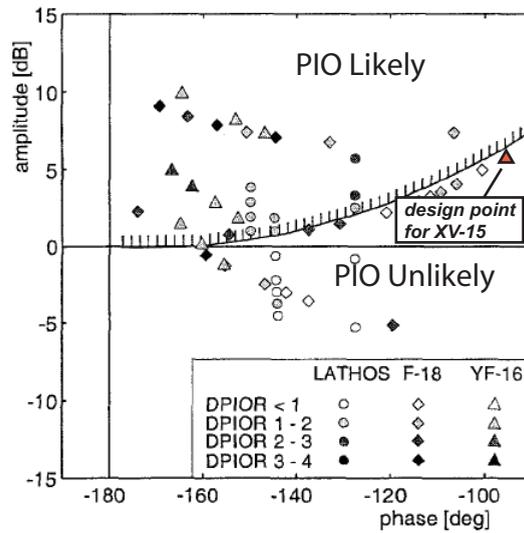


Fig. 14 OLOP Criteria and XV-15 design point.

4.3 Design Optimization

The four design parameters freed in the optimization were the attitude and rate gains for roll and yaw (k_ϕ , k_p , k_ψ , k_r). In this and many other cases, the dynamic response specifications of ADS-33 do not sufficiently constrain the low-frequency (steady-state) response which defines the integral gain requirement, so alternatively (as mentioned above) the integral gain was fixed as a proportion of the attitude gain based on classical rules-of-thumb. The initial control gains were chosen such that the design would be in Level 3 to start out (Phase 1). The initial design is intentionally poor to demonstrate the optimization. The movement of the design parameters needed to achieve a final optimized design is shown in Fig. 15. The optimization is finished after all specifications have moved into Level 1; the summed objective has been minimized and the design parameters have not changed in 3 iterations.

Fig. 16 shows the final optimized results displayed in the *handling qualities (HQ) window*. All specifications meet the Level 1 requirements. It can be inferred that the gains have been minimized because the bandwidths (d, e) and disturbance rejection specifications (f, g) are against the boundaries. This indicates that further reducing the gains would move the system into Level 2, which is not allowed by the optimization. A 20% design margin (DM=0.2) was specified on the bandwidth requirements as shown by the dashed lines, which moves the effective “Level 1” optimization requirement 20% of the width of the Level 2 region away from the actual Level 1/Level 2 boundary, providing some margin to uncertainty.

Fig. 17 shows an example of the stability margin as determined from the broken-loop response of the final design for the lateral axis. The figure shows that there is satisfactory gain and phase margin respectively, as also indicated by the HQ window of Fig. 16, that the crossover frequency is about $\omega_c = 2.5$ rad/sec, and that the shape of the broken-loop response has a desirable (k/s) characteristic.

A step response in the lateral axis is shown in Fig. 18. This serves as an example that the time responses are reasonable. Overall, the step response shows that the aircraft response is smooth and has no residual oscillations. The response meets bandwidth requirements, and has little saturation for large inputs. The overshoot and slow settling time are related to the low control power of the XV-15, as turning up the gains high enough to further improve these characteristics causes serious rate and position limiting of the actuators for all but very small inputs. The overshoot is approximately 30%, which corresponds to a damping ratio of ~ 0.35 for a second order system, which is just at the boundary of the required Level 1 damping for hover/low speed.

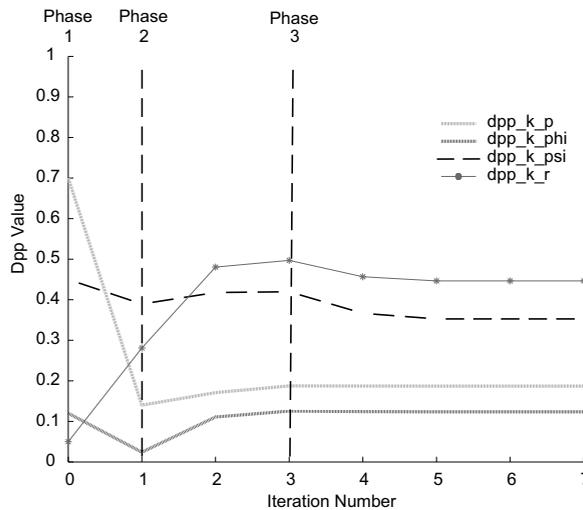


Fig. 15 Iteration history for design parameters.

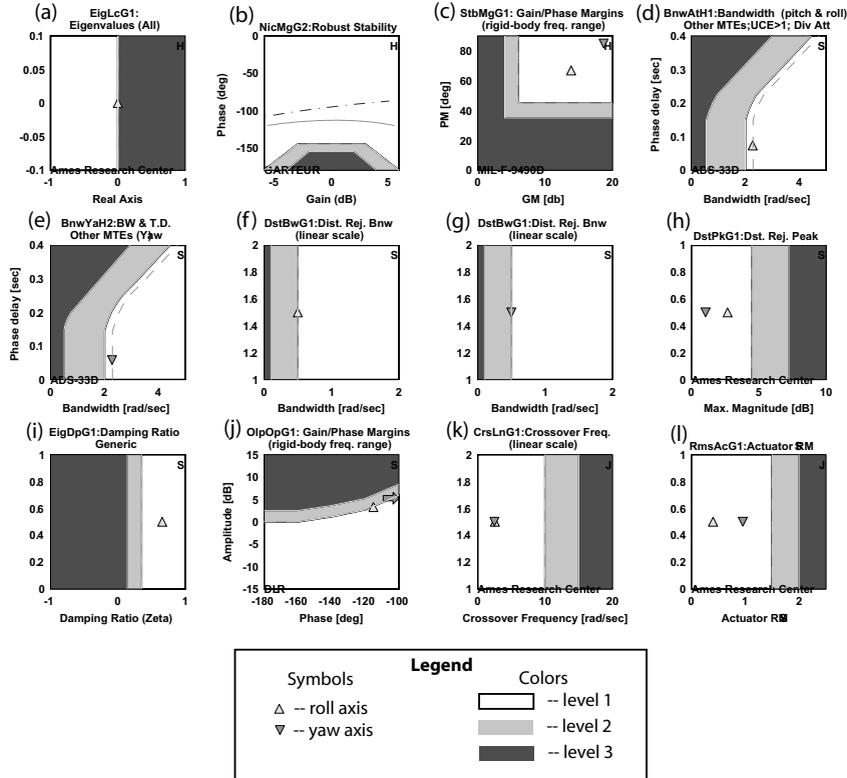


Fig. 16 HQ window for optimized XV-15 design.

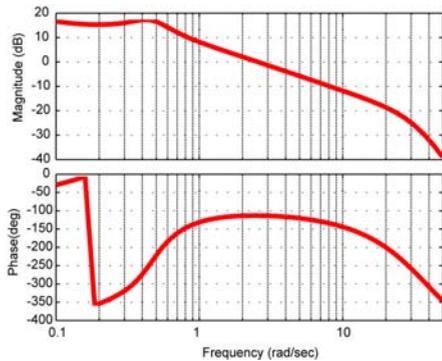


Fig. 17 Lateral broken-loop response.

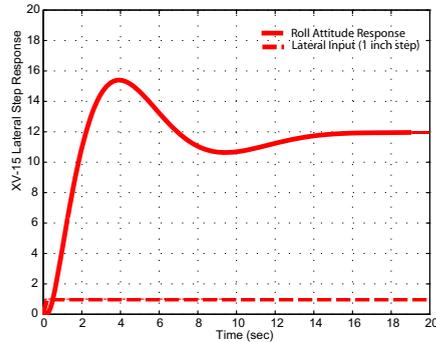


Fig. 18 Closed-loop step response for XV-15 control system solution.

4.4 Uniqueness of the Solution

A sensitivity analysis was performed on the converged solution. The resulting Cramer-Rao (CR) bounds for the design parameters had values all less than 20%, indicating a well-posed problem and a well-defined optimum. In order to test the uniqueness of the solution with respect to initial conditions, 100 cases were run with the initial conditions randomly varied by +/- 50%. The result was that all 100 optimizations converged to the same final gain set. The initial and final gain values for the roll rate gain are shown for all 100 cases in Fig. 19. The figure indicates that the

initial gains were widely varied but the final gains after optimization are identical. The same results were seen for the other freed design parameters but are not shown for brevity. This result gives great confidence that the problem was setup such that a unique solution could be determined.

By looking at some selected iteration histories for the roll rate gain as presented in Fig. 20, it is clear that the design parameters take different paths during the optimization, but converge at the same optimized solution in the end. The iteration histories for the other design parameters show similar behavior, but are not depicted in the paper. The excellent convergence properties are due in large part to the smooth behavior of the distance algorithm Sec. 3.3.3

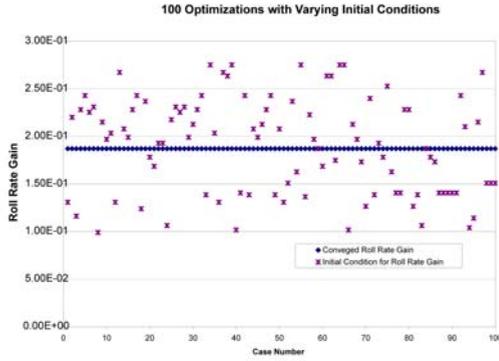


Fig. 19 Optimized results for dpp_k_p for 100 cases with varying initial conditions.

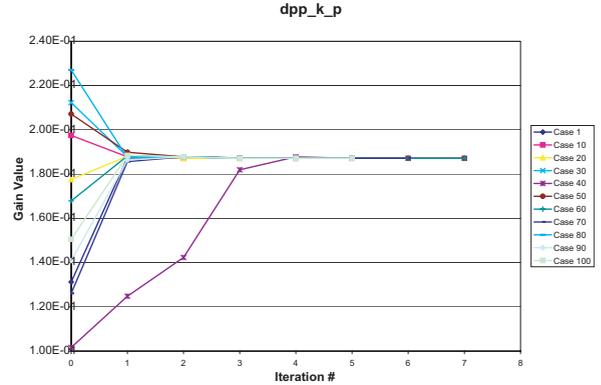


Fig. 20 Iteration history for dpp_k_p (selected cases).

4.5 Robustness of the Optimized Solution to Model Uncertainty

The two types of identification models that are used for flight control design are parametric and non-parametric models, each of which have their own uncertainty bounds. The non-parametric model is also known as an identified frequency response from flight data and has a confidence range of magnitude and phase estimates. The second type of model is a parametric model, or state-space model, and has uncertainty bounds for each parameter value. These uncertainties can be included in the analysis in order to determine how robust the control system is to uncertainty in the plant dynamics.

4.5.1 Non-parametric Uncertainty

The frequency response uncertainty (95% confidence limits) can be estimated based on random error ϵ_r in the frequency response estimate (Ivler 2008, Bendat and Piersol 1986). The magnitude uncertainty is:

$$\left| \hat{H} \right| (1 - 2\epsilon_r) \leq |H| \leq \left| \hat{H} \right| (1 + 2\epsilon_r) \quad (11)$$

where $|H|$ is the true magnitude and $\left| \hat{H} \right|$ is the estimated magnitude. The associated phase uncertainty is:

$$\hat{\phi} - 2\epsilon_r \leq \phi \leq \hat{\phi} + 2\epsilon_r \quad (\text{rad}) \text{ with } 95\% \text{ confidence} \quad (12)$$

where ϕ is the true phase and $\hat{\phi}$ is the estimated phase.

By replacing the state-space model with the identified frequency response in the on-axis and then applying the error bounds, the effect of the random error on the flight control system can be measured. To generate the broken-loop

response, the frequency response data is multiplied with the control system data. The analysis is validated by comparing the broken loop response using this method to the broken-loop response using the state-space model.

The next step is to determine how much the stability margins change when using the identified frequency response plus the estimated error bounds due to random error. The results for the analysis in the roll axis are shown in Fig. 21. For the worst case, which is the combination of the lower magnitude and lower phase bound, the low frequency gain margin decreases from 13.9 dB to 8.2 dB - this large variation is due to the poor coherence (accuracy) of the identified response at low frequency. The phase margin changes no more than a few degrees for any combination, because of the good coherence around crossover.

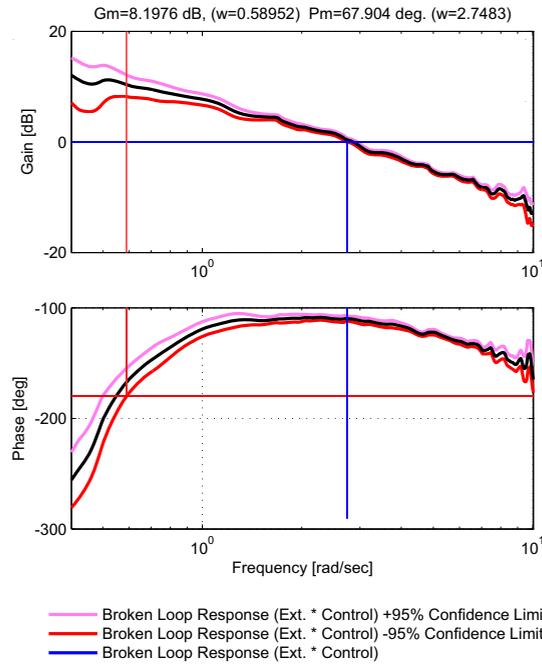


Fig. 21 Non-parametric uncertainty analysis for roll axis.

4.5.2 Parametric Uncertainty

Looking at the robustness of the converged design to parametric modeling uncertainty is another important consideration for a flight control design. Once the optimization has converged on a solution that the engineer thinks provides good performance, and meets all the specifications required, a Monte Carlo type parametric uncertainty analysis is performed. Identified models will have an associated uncertainty (Cramer-Rao) bound with each identified derivative (Tischler and Remple 2006). The Cramer-Rao (CR) bounds provided by CIFER[®] are scaled to represent the expected standard deviation in the identified parameters:

$$(CR_i)_{CIFER} \approx \sigma_i \quad (13)$$

so $\pm 3(CR_i)_{CIFER}$ corresponds to the 99% confidence interval for each derivative. The value of the Cramer-Rao bound is normalized by the parameter value to yield an accuracy in percent (CR_i) .

For example, the Cramer-Rao bound (and associated standard deviation) for L_p is 15.41%. The test shown in Fig. 22 uses a random positive or negative 3σ variation on all the derivatives, and evaluates how the handling qualities are changed with uncertainty. The results indicate that the stability and Nichols margins remain good and that by building in a 20% design margin on the nominal case ensures that bandwidth performance remains within the ADS-33 bounds, even with a $\pm 3\sigma$ uncertainty in the identified stability derivatives. The stability margins in the roll axis vary by 10 degrees phase margin and 8 dB gain margin across all the cases shown. The stability margins in the yaw axis vary by

about 5 deg and 5 dB. This is a small variation in the stability margin, especially when the margins are not near the boundary, as in the case of the XV-15. These results show that the design is robust and that including a conventional stability margin and a design margin builds in a good tolerance for uncertainty.

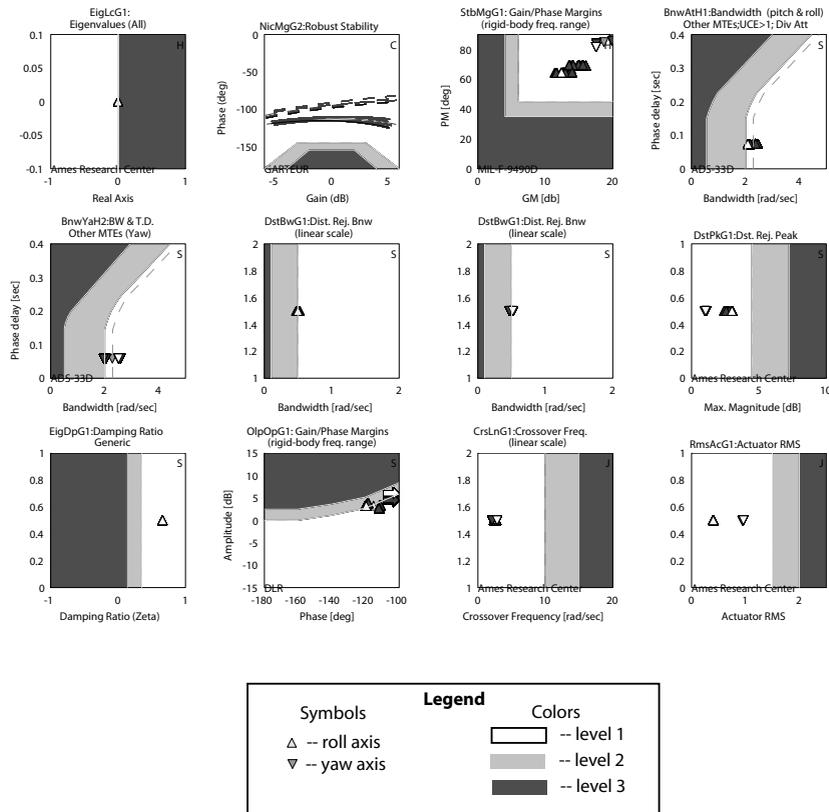


Fig. 22 100 perturbations at 3σ .

4.6 Design Trade-Offs

A study was conducted to examine fundamental trade-offs in the controller design. For this study, the design margin was placed only on the disturbance rejection bandwidth (DRB) specifications (roll and yaw). The results presented in Fig. 23-Fig. 27 display smooth (monotonic) trends.

The disturbance rejection bandwidth is mainly set by the attitude gain. Increased disturbance rejection bandwidth is achieved with increased attitude gains in both roll and yaw. Fig. 23 shows that as the disturbance rejection bandwidth increases the piloted roll bandwidth does not change. This is due to the optimization trying to find the minimum crossover frequency and actuator RMS (actuator usage), which occurs when the piloted bandwidth is against the Level 1 boundary for this aircraft and the chosen set of specifications.

As shown in Fig. 24, the crossover frequency for the roll axis decreases with increasing disturbance rejection bandwidth, while the yaw axis crossover frequency increases – not an immediately intuitive result. However, since the roll axis is an ACAH architecture, the increased attitude gain also increases the piloted bandwidth and thus allows a reduction in the crossover frequency to maintain the roll bandwidth at the boundary. The yaw axis is RCAH, so the increased attitude gain causes a reduction in the piloted bandwidth, and requires an increase in the crossover frequency to still meet the required yaw bandwidth. Fig. 25 shows that the increase in the disturbance rejection bandwidth (and associated higher attitude gains) also creates a downward trend in the phase margin in both axes. The actuator RMS is also increasing in both channels as would be expected with the higher attitude gains as needed to provide the improved disturbance rejection bandwidth (Fig. 26).

Fig. 27 shows that the OLOP specification moves towards the boundary as the disturbance rejection bandwidth increases. This makes sense because the actuators are becoming more active (as shown by the actuator RMS plot) as the disturbance rejection bandwidth increases. The OLOP spec is the limiting factor in the design margin optimization. If this specification was relaxed, then the optimization would be able to meet all the Level 1 requirements for higher disturbance rejection bandwidth design margins.

The trade-off results are not immediately intuitive even for this fairly simple case study, and they illustrate the significant value of the optimization-based approach to allow a clear understanding of the inter-relationship between the requirements, the design parameters, and the associated control system performance.

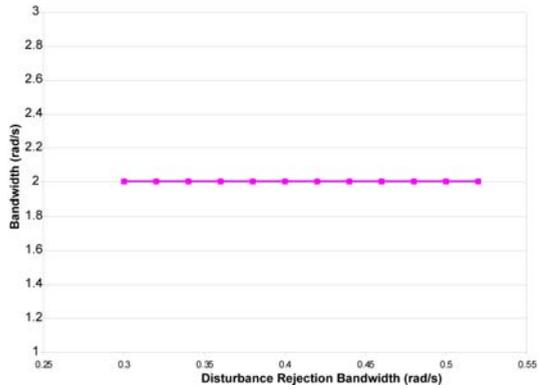


Fig. 23 Piloted roll bandwidth vs. disturbance rejection bandwidth.

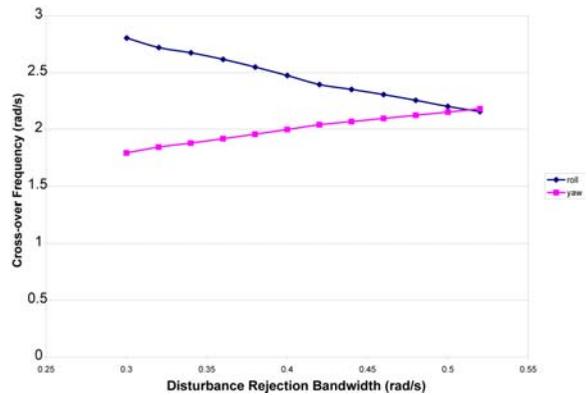


Fig. 24 Crossover frequency vs. Disturbance Rejection Bandwidth

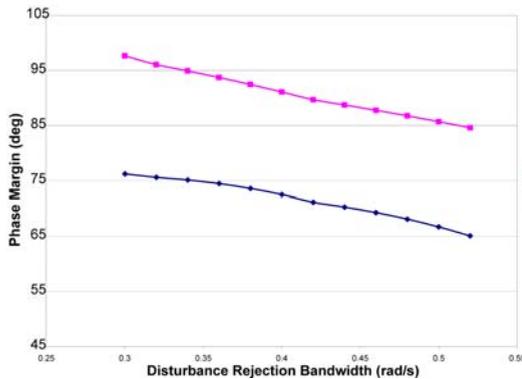


Fig. 25 Phase margin vs. disturbance rejection bandwidth.

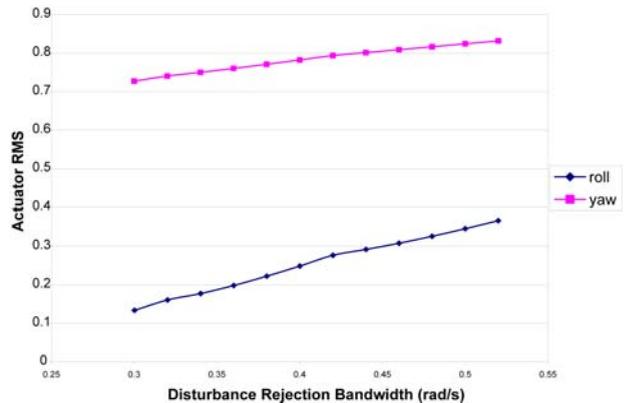


Fig. 26 Actuator RMS vs. disturbance rejection bandwidth.

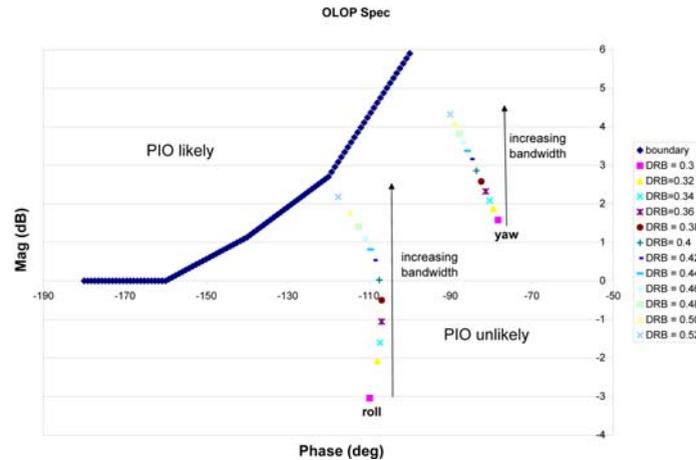


Fig. 27 OLOP Spec for disturbance rejection bandwidth design margin optimization.

5 Model Following Control Laws and Trade-offs for UH-60M

The XV-15 case is a simple example used to demonstrate general concepts. To demonstrate the application of the methods to a complex real world problem, we now look at a system which follows the general structure of the new UH-60Mu control laws as implemented on the RASCAL aircraft (Fig. 28) for risk reduction flight tests (see Fletcher et al 2008). In contrast to the actual UH-60Mu control laws, however, the example here only considers hover and low speed flight and implements a subset of the response types available in the full-up system.



Fig. 28 AFDD RASCAL JUH-60A in-flight simulator.

Mimicking the more complex UH-60Mu control laws, the system being presented is a full authority, model-following, set of control laws providing attitude-command/attitude-hold (ACAH) response characteristic in pitch and roll maneuvers with the stick out of detent (per axis). Velocity hold (VH) functionality is activated when the stick is returned to detent (again, per axis). In yaw, the control laws provide rate command functionality with the pedals out of detent with heading hold functionality activated when the pedals are returned to detent. The development concentrates mainly on pitch, roll, and yaw and assumes a direct connection for the collective stick with no augmentation.

5.1 Analysis Model

Fig. 29 depicts the general structure of the model following system.

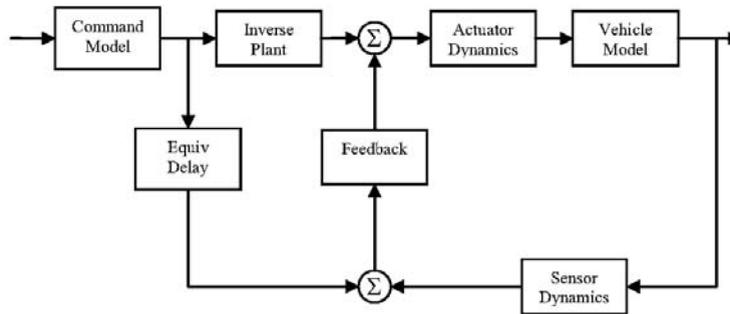


Fig. 29 Overview of the control system.

The basic feedback architecture for achieving ACAH/VH in pitch and roll and RCAH in yaw is PID in all axes. The command models (2nd order in roll and pitch attitude and 1st order in yaw rate) generate desired responses to piloted inputs and are designed to meet ADS-33 bandwidth requirements. A first-order pseudo-inverse of the aircraft generates actuator commands to the vehicle which would approximately result in the commanded responses. The use of simple first-order inverse dynamics as compared to the higher-order response of the aircraft results in model-following errors that are nulled via the feedback. The RASCAL aircraft employs two sets of actuators; one set being the standard UH-60A primary actuators and the other the research servos. In the current analysis, both sets were modeled as simple second order systems, including rate and position limiting which is very important in the analysis and optimization.

The availability of a linear model that accurately represents the bare airframe dynamics is of utmost importance. This study was based on a 22 state model identified from frequency sweep flight data using CIFER[®] and validated against additional flight data not used in the identification process. The states of the model were: 3 body translational rates (u, v, w), 3 body angular rates (p, q, r), 2 Euler angles (ϕ, θ), rotor RPM, thrust coefficient, torque, fuel flow, rotor azimuth, lateral and longitudinal flapping, coning rate and angle, inflow, 2 lateral lead lag, and 2 longitudinal lead lag. The controls can be referenced to equivalent pilot stick inches at the mixer input ($\delta_{lon}, \delta_{lat}, \delta_{col},$ and δ_{ped}).

Delayed values of commanded responses are compared with the actual responses and the errors are nulled via the feedback path. The time delay block, equivalent to the forward path time delays of the system, is included to avoid overdriving high-order dynamics not included in the lower-order pseudo-inverse. The feedback system is a PID architecture on the attitude responses and a PI architecture on the longitudinal and lateral ground velocities. The analysis included representations of the sensors filters used in the feedback path of the control laws meant to reduce signal noise and eliminate unwanted frequency content from the feedback data.

For this effort, there were $n_{dp} = 13$ design parameters consisting of proportional gains on angular rates and proportional plus integral gains on Euler angles for pitch, roll, and yaw attitude, and proportional plus integral gains on longitudinal and lateral components of ground speed. It is important to determine reasonable initial values for these parameters at the beginning of the analysis. Therefore preliminary design analysis should be carried out as will be discussed later in this paper.

5.2 Broken-Loop Model Validation

To make absolutely certain that the analysis was a good representation of the actual system, a few hours of preliminary flight testing were conducted in June 2008 to gather automated frequency sweep data for analysis validation. The broken loop responses of the analysis were compared to this flight data and slight discrepancies noted, some related to the mathematical model of the bare airframe and others due to unmodeled effects. To correct the analysis and eliminate as much of the observed discrepancies as possible, gain and time delay modifications were calculated and

applied. Fig. 30 shows the roll axis broken-loop comparison between the corrected analysis and flight. As can be seen, the match in both gain and phase is very good. Results from other axes were similar and are not shown for brevity. Table 2 compares the gain and phase margins obtained using the same nominal set of gains in the analysis and in flight, and shows that the analysis accurately predicts the key characteristics of the feedback system.

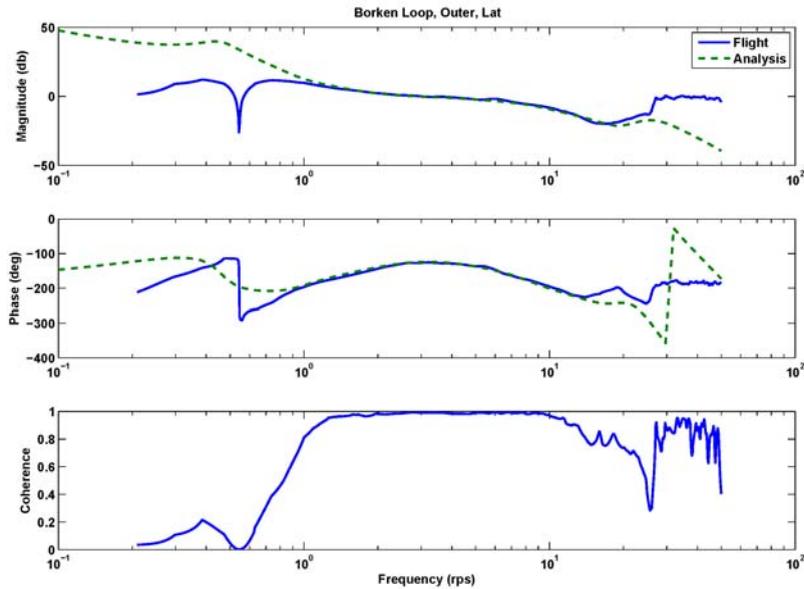


Fig. 30 Lateral broken-loop comparison between analysis and flight.

Table 2 - Gain and phase margin comparisons between flight and analysis

		Flight	CONDUIT (Fixed)	% Diff
Longitudinal	GM	10.26 @ 10.06 rps	11.35 @ 10.54 rps	10.6
	PM	51.09 @ 3.56 rps	47.82 @ 3.34 rps	-6.4
Lateral	GM	6.25 @ 8.52 rps	7.35 @ 8.67 rps	17.6
	PM	53.26 @ 2.98 rps	55.35 @ 3.00 rps	3.9
Directional	GM	N/A	13.21 @ 19.37 rps	
	PM	50.53 @ 6.34 rps	56.06 @ 5.44 rps	10.9
Collective	GM	6.15 @ 6.70 rps	6.91 @ 8.08 rps	12.4
	PM	75.91 @ 2.20 rps	68.42 @ 1.84 rps	-9.9

5.3 Specifications

Table 3 lists the specifications used for the analysis and optimization of the current system. These include requirements covering stability margins, handling qualities, disturbance rejection capabilities, and off-axes coupling. As discussed in detail earlier in the XV-15 example section, actuator rate saturation can have a significant detrimental effect on the handling qualities of the vehicle and directly lead to pilot induced oscillation tendency. Therefore, the OLOP spec, discussed earlier, was included in the current analysis to evaluate and control the predicted PIO tendency of the design during the optimization process.

Of significant interest among these specifications is the minimum crossover frequency (CrMnG2) requirement. The XV-15 example discussed earlier uses a response feedback architecture which closely ties the broken loop crossover and piloted bandwidth values. Therefore, ensuring good piloted bandwidth (an ADS-33D requirement) also ensures good system crossover. This is not the case with the UH-60 example which uses a model following architecture. In a model following architecture the piloted bandwidth is mostly set by the command model being used and is independent of the broken loop crossover frequency. In a model-following system, therefore, it is possible to have good piloted bandwidth with fairly low crossover frequencies. Although on first consideration, these low crossover frequencies seem desirable because they minimize the “cost of feedback,” there are important additional considerations

that establish a need for higher crossover frequencies – good performance robustness in the face of model parameter uncertainties, and good model-following (command model tracking) with an imprecise pseudo-inverse. It is also recognized that a designer often sets the crossover frequencies based on historical precedence (for upgrade programs) and company design practice.

In this study, various methods of ensuring reasonable values of crossover frequencies in the model following architecture were considered – a good approach is to use the performance robustness analysis in preliminary design to calculate a minimum value of crossover needed. This method is further discussed next in Sec. 5.4 (“Preliminary Design”). The minimum crossover frequency specification is then used during the full analysis to ensure that the optimized crossover frequency is at least the predesign value (or higher)

Table 3 Specifications used in the analysis

Name	Description	Type	Comments
EigLoG1	Eigenvalues (All)	H	Ensure stability
StbMgG1	Gain/Phase Margins (rigid-body frequency range)	H	Ensure adequate stability margins (MIL-STD_9490D) Margins have to be checked at various points so multiple copies of this spec is used
EigDpG1	Damping Ratio Generic	H	Ensures that all eigenvalues in frequency range of interest have sufficient damping
ModFoG2	Response Comparison	S	Ensure responses of aircraft closely match responses of command model
BnwAtH1	Bandwidth (pitch & roll) Other MTEs; UCE>1; Div Att	S	Short term pitch/roll response requirement (ADS-33D)
BnwYaH2	BW & T.D. Other MTEs (Yaw)	S	Short term yaw response requirement (ADS-33D)
CrsMnG2	Min. Cross Freq. (linear scale)	S	Ensure acceptable crossover frequencies
OlpOpG1	Open-loop onset frequency	S	Ensure acceptable actuator saturation characteristics and low PIO tendency
DstBwG1	Dist. Rej. Bnw (linear scale)	S	Ensure satisfactory disturbance rejection bandwidth
DstLoG1	Dst. Rej. Peak (Low Freq.)	S	Ensure good damping of disturbance response
DrfOaG3	Normalized Off-Axes Drift	S	Ensure minimal off-axes drift due to on-axis control input
CouYaH1	Coupling Yaw/Collective	S	Ensure good yaw/collective coupling (ADS-33D)
RmsAcG1	Actuator RMS	J	Control over design
CrsLnG1	Crossover Freq. (linear scale)	J	Control over design
CouPRH2 (check)	Pitch-Roll Coupling Frequency Domain	C	Ensure good pitch/roll coupling (ADS-33D)
FrqHeH1 (check)	Heave Response Hover/Low-Speed	C	Ensure good heave dynamics (ADS-33D)

5.4 Preliminary design

For complex control design problems such as the UH-60Mu flight control system, it is helpful to conduct a preliminary design analysis before beginning the full design and optimization process. This ensures a reasonable initial estimation of feedback gain and essential handling qualities values. There are several approaches to performing preliminary design including classical control methods, loop shaping, and rules of thumb. In all cases, the full control design problem is generally reduced to decoupled axes and a simplified block diagram. The preferred method used herein for model simplification/decoupling is the constrained variable, or coupling numerator, approach (Cheng et al 1995). This method assumes that tight feedback loops are used to constrain the off-axis responses. This is a reasonable assumption because the final control system will stabilize the off-axis responses. Here, an example of one method for preliminary design analysis will be shown for the longitudinal axis.

For preliminary design, the constrained longitudinal response can be well approximated in the crossover frequency range by the simple form $(M_{\delta_{\text{on}}} e^{-\tau_0 s})/s^2$. Using classical control techniques and rules of thumb, the following rela-

tionships between the maximum achievable pitch channel crossover frequency, ω_{c_θ} , the effective time delay, τ_θ , the control derivative, $M_{\delta_{lon}}$, and the feedback gains, k_θ and k_q , can be established (Tischler 1987):

$$\omega_{c_\theta} = \frac{0.370}{\tau_\theta}; \frac{k_\theta}{k_q} = 0.442\omega_{c_\theta}; k_\theta = \frac{\omega_{c_\theta}^2}{2.48M_{\delta_{lon}}} \quad (14)$$

These relationships set the maximum achievable crossover frequency as a function of the time delay of the system to achieve a phase margin of 45° at the point of maximum phase. In the lateral axis, where the aircraft transfer function model is closer to a $(L_{\delta_{lat}} e^{-\tau_\theta s})/s$ form in the crossover frequency range, there is an analogous set of expressions. Using the values for equivalent time delay and response gain determined from the coupling numerator response, the maximum achievable crossover frequency for the longitudinal channel was estimated at $\omega_{c_\theta} = 2.3$ rad/sec, and the initial values of the attitude and rate gains were determined. For the actual pitch response, a somewhat higher crossover frequency can be achieved (see Fig. 35), because the actual pitch rate damping (M_q) is in fact not zero, and provides extra lead in the loop.

Two key determinants of the broken-loop crossover frequency are model-following accuracy for the nominal flight condition and the performance robustness for off-nominal conditions - both generally improve with increasing crossover frequency. A preliminary design optimization of the feedback gains using CONDUIT[®] is useful to determine the crossover frequency values that ensure adequate model-following accuracy and performance robustness, with acceptable stability margins. In this preliminary design analysis, an integral gain set as a fraction of the proportional gain (as in the XV-15 case) is included in the compensator.

Model-following accuracy is measured by a mismatch cost function (Tischler and Remple 2006) between the closed-loop pitch rate and the commanded pitch rate (Model Following Cost). Fig. 31 shows an overlay of the commanded response and actual response for a nominal feedback design. In this case, Model Following Cost is very low (12); values of less than 100 indicate reasonable agreement of the aircraft response and command mode and values of less than 50 indicate nearly perfect agreement (Tischler and Remple 2006). Performance robustness is measured using the same mismatch cost functions, but now calculated between the nominal closed-loop pitch attitude frequency response and off-nominal cases, with the spec taking the largest value (Performance Robustness Cost). Fig. 32 shows the variation in the pitch attitude frequency response for different bare airframe models, with a nominal feedback design. Fig. 33 shows the improvement in both Model Following and Performance Robustness Costs for increasing crossover frequency in the longitudinal channel design. The final preliminary design configuration was determined by setting the Performance Robustness Cost to be no greater than 100, with an associated desired crossover frequency of about 2.5 rad/sec.

The preliminary design analysis provided a good starting estimate for the feedback gains, as well as a reasonable value for the minimum crossover frequency, for use in the full design phase. Preliminary design analysis in the lateral and directional axes produced desired crossover frequencies of about 3.0 rad/sec and 4.0 rad/sec, respectively. Calculating the Performance Robustness Cost requires the input of additional bare airframe models at many off-nominal conditions and is a strong function of the models used. Therefore, it is recommended to establish a desired crossover frequency in the preliminary design phase and then constrain this as a lower bound using the minimum crossover frequency specification.

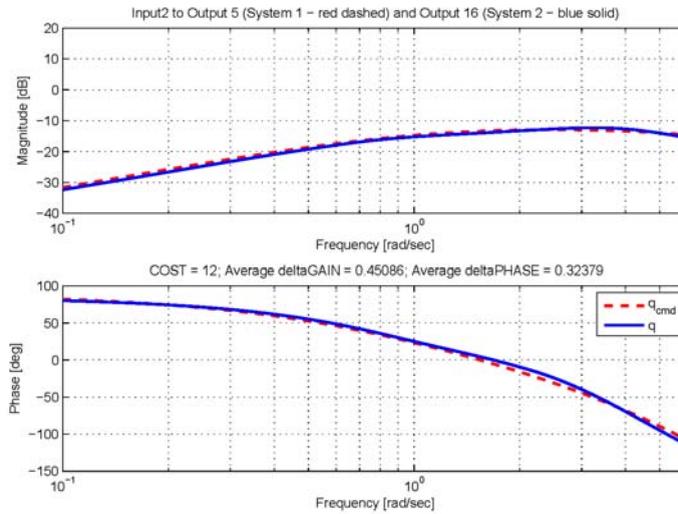


Fig. 31 Model-following accuracy for longitudinal preliminary design (nominal model).

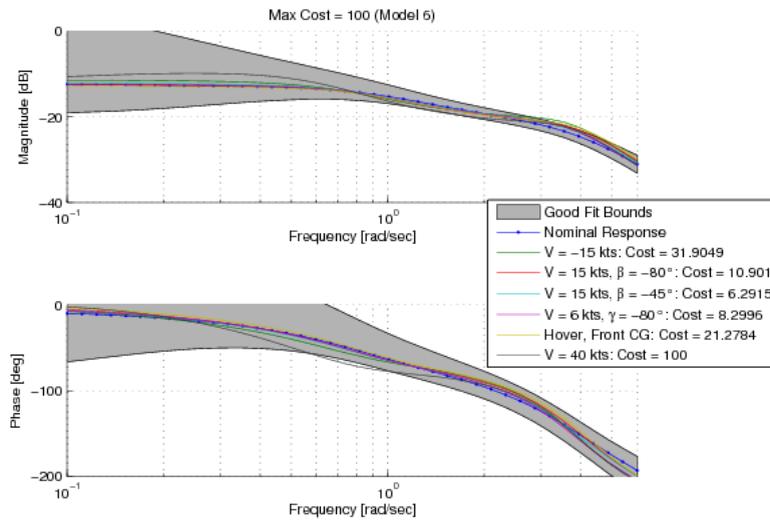


Fig. 32 Closed-loop performance robustness for longitudinal preliminary design.

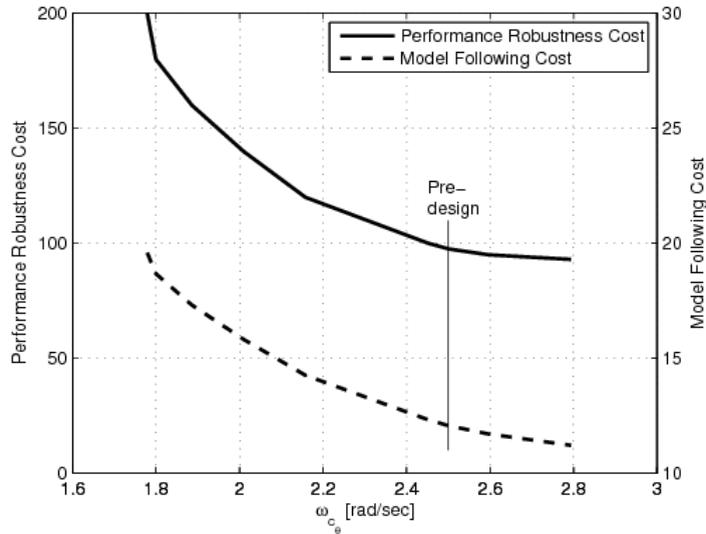


Fig. 33 Variation in model following and performance robustness costs with change in crossover frequency.

5.5 Optimization Approach

The goal of the current optimization effort was to explore the effects of increasing the performance of the system, as measured by disturbance rejection bandwidth (DRB) capabilities, on stability margins and other system characteristics. As shown by the XV-15 example, DRB and stability margins are competing requirements since improving disturbance rejection performance tends to require an increase in feedback gains and reduction in stability margins.

Whereas the requirements for stability margins are well defined in existing specification documents (e.g. MIL-DTL-9490D), boundaries for disturbance rejection bandwidth have not yet been established. Efforts are underway to determine minimum ADS-33 requirements for DRB for different axes and different modes. Starting from these minimums, DRB capability of the system can then be used as a tuning knob and varied to achieve the tightest disturbance rejection possible. While this systematic increase of DRB is carried out, the crossover frequency can either be maintained (using the minimum crossover frequency spec) at the value prescribed by preliminary design, or it can be increased in unison with the DRB values. Since the values obtained through performance robustness analysis in preliminary design are “minimums,” it was decided to increase the minimum crossover requirements along with the DRB values to arrive at solutions that not only provide improved disturbance rejection performance but also improved performance robustness. The systematic tightening of the specification boundaries is accomplished by means of the design margin (DM).

Of the control law gains selected as design parameters for optimization, the integral gains were of elevated interest. This was because pilot comments from preliminary flight tests (mentioned earlier and conducted to gather analysis validation data and preliminary pilot comments) indicated a concern about off-axes drift characteristics of the system. Prior to these flight tests the integral gain in each axis was tied to the corresponding proportional gain and crossover frequency, based on classical rules-of-thumb, as was done for the XV-15 example. Pilot comments made it clear that for a complex problem as this, a more direct approach had to be devised. Therefore, a new specification which mimics the ADS-33 attitude hold spec, but is applied to the off-axes responses, was developed. The ADS-33 attitude hold spec (shown in Fig. 34) ensures that following a pulse-type control input the aircraft will return to within 10% of the maximum response within 10 sec, as required by ADS-33D for attitude hold systems. The new spec (also shown in Fig. 34), in comparison, ensures that the magnitude of the off-axes responses resulting from an on-axes input should never exceed 25% of the on-axis response magnitude and should all reduce to less than 10% of the maximum on-axes response magnitude within 5 seconds.

A family of designs maintaining the current stability margin requirements (GM= 6db and PM=45 deg) with ever increasing DRB capability and minimum crossover frequency was then generated and analyzed.

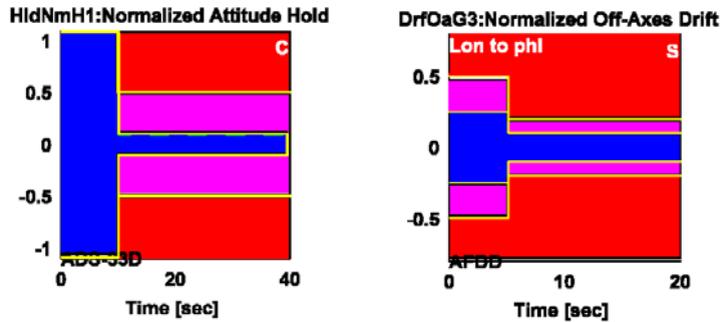


Fig. 34 ADS-33D attitude hold spec (left) and the new off-axis drift spec (right).

5.6 Analysis Results

Even though this system is not overly complex, it is a good example for demonstrating the difficulty in finding a set of design parameters that can satisfy all the competing requirements that are typical in a modern design. On the surface, the trade-off being looked at is between stability margin and disturbance rejection performance. However, this trade-off is made much more complicated by the fact that a slew of other requirements have to be satisfied at the same time.

The complete system consists of a total of 155 states with 13 design parameters that are varied to satisfy a total of 74 specifications. Clearly, not all the specifications are sensitive to all the design parameters. However, in many cases there are a number of specifications that are sensitive to the same design parameter. Moreover, in many cases increasing a given design parameter can improve the system's compliance with a given specification while at the same time deteriorating the system's compliance with another specification. Therefore, a myriad of secondary trade-offs are concurrently taking place in conjunction with the trade-off between stability margins and disturbance rejection.

Design margin optimization was used to carry out the trade-off between stability margin and DRB. An initial base optimization was carried out to establish baseline values for DRB and crossover frequency boundaries. These boundaries were then set at 6 points both above and below this baseline case and the system fully optimized where possible (the points are marked with a design margin value from DM=0.0 to DM=1.0 at 0.2 intervals, i.e., 20% increments in increased performance in the DRB and minimum crossover requirements). Note that the range of crossover frequency values used covered values both above and below the values suggested by preliminary design. Even though the lower values would normally be discarded (based on the results of preliminary design), they were kept to allow a thorough look at the effects of variations in crossover frequency and DRB.

Maintaining the 6 db gain margin and 45 degrees phase margin requirement, the optimization was successful for 4 of the 6 cases. The results showed the system to be more limited in the pitch axis due to the pitch phase margin consistently being around the required limit of 45 degrees while the phase margin in roll was much higher. The following discussions will therefore concentrate on the pitch axis and will only refer to roll axis results when necessary. Fig. 35-Fig. 37 show a compilation of the results for pitch ACAH and ACVH, with Fig. 35 depicting the variation of optimized crossover frequencies, Fig. 36 depicting the variation of optimized DRBs, and Fig. 37 depicting the optimized phase margins.

For each design point, the optimization engine pushes the crossover frequencies and actuator RMS as low as possible, while still satisfying all requirements. As discussed earlier, in this analysis minimum crossover frequency specifications were used to ensure satisfactory values in each axis. It would, therefore, be expected that the crossover frequencies would end up on the minimum specified boundaries for a fully optimized case. Looking at the first point in the crossover frequency plots of Fig. 35, however, it can be seen that the inner ACAH loop crossover frequency was not driven to the allowable boundary while the outer ACVH loop crossover frequency was. This indicates that for the outer ACVH loop to achieve the required crossover frequency, a higher level of performance (than required by the specified boundaries) needs to be achieved by the inner ACAH loop. The same conclusion can be made by looking at the first points in the disturbance rejection plots of Fig. 36 where, again, the DRB of the inner ACAH loop was not driven to the allowable boundary while the DRB of the outer ACVH loop was. This shows that the optimization pro-

cedure is successfully balancing the requirements of the various modes and ensuring that all requirements are satisfied. From Fig. 37, it can also be noted that an additional constraint on the inner ACAH loop is the phase margin requirements. The figure shows that the inner ACAH loop phase margin stays on the 45 deg. requirement boundary throughout the design margin optimization.

As design margin increases, higher and higher performance is demanded from the system by increasing the requirements on crossover frequency and DRB in both the inner and outer loops. Initially, the increased performance of the outer ACVH loop is achieved by increasing the proportional gain on velocity, as seen in Fig. 38, without changing the inner loop gains. This keeps the inner ACAH loop crossover frequency and DRB mostly flat. However, at DM=0.4 the inner ACAH loop crossover frequency reaches the increasing boundary and gains have to increase for this spec to remain in Level 1. As may be seen from the design parameter variation plots of Fig. 38, the rate and attitude gain values significantly increase between DM=0.4 and DM=0.6. Additionally, note the significant concurrent drop in the velocity integral gain that is needed to allow the increasing DRB of the outer ACVH loop. At DM=0.6 the crossover and DRB values for both inner and outer loops are on the boundaries and increases in all design parameters are needed to support further increases in performance. This, however, would not be possible without violating the 45 degree phase margin requirement on the inner ACAH loop and the design margin optimization can not proceed further. Therefore, the DM=0.6 case has the best performance achievable given the current stability margin requirements.

As mentioned earlier, the range of crossover frequencies used for design optimization encompassed the minimum value suggested by the preliminary analysis. Even though most of the optimized points had roll crossover frequencies above the preliminary design minimum, only the last point (DM=0.6) had a higher pitch crossover frequency. Therefore, DM = 0.6 on the plots would be the selected design from the set of optimized solutions shown here as it not only has the best performance but it is the only solution with acceptable performance robustness.

Fig. 39 shows a partial view of the CONDUIT[®] HQ window for the DM=0.6 case. All requirements are in the Level 1 region, except for yaw-axis bandwidth which cannot be achieved for the UH-60 helicopter based on its installed control power (therefore this spec is designated as “check only”). Notice that some of the specifications are on the Level 1 boundary, which is always the case for the optimized results. The results show that the off-axis drift has been effectively controlled and the off-axis responses all remain within the boundaries of the new spec while the optimized values of the integral gains compare well with nominal values based on corresponding crossover frequencies and proportional gains. Note that in the case of yaw attitude response to a longitudinal input (“Lon to psi” in Fig. 39) the original Level 1 boundaries of the spec turned out to be too stringent for the system to achieve. The boundary had to be expanded as shown in the figure.

As a further example of the effectiveness of the off-axes drift specs, the normalized insensitivity value (\bar{I}_i , Tischler and Remple 2006) of the yaw integral gain was compared with and without the yaw attitude response to longitudinal input drift spec (“Lon to psi” in Fig. 39). An insensitivity value ($\bar{I}_i < 20\%$ is desirable and $\bar{I}_i < 10\%$ is ideal and indicates that the specifications are sensitive to changes in that design parameter (Tischler and Remple 2006). With the spec included in the analysis, the yaw integral gain has a good insensitivity value of around $I = 10\%$, while removing the spec results in the insensitivity value to increase to an undesirable value of $I = 59\%$. This indicates that the new off-axis drift specs are an effective means of guiding the value of the integral gains during the optimization. The OLOP spec was included in this case study (though not included in the figure), but with the large control power of the UH-60 (as compared to the XV-15) was not found to be a limiting factor for any of the designs.

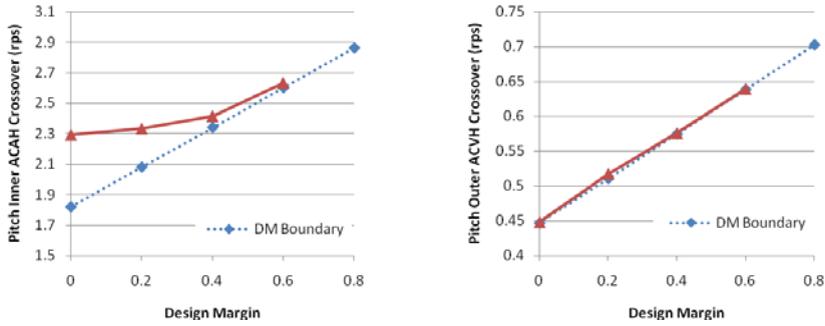


Fig. 35 Variation of crossover frequency with design margin.

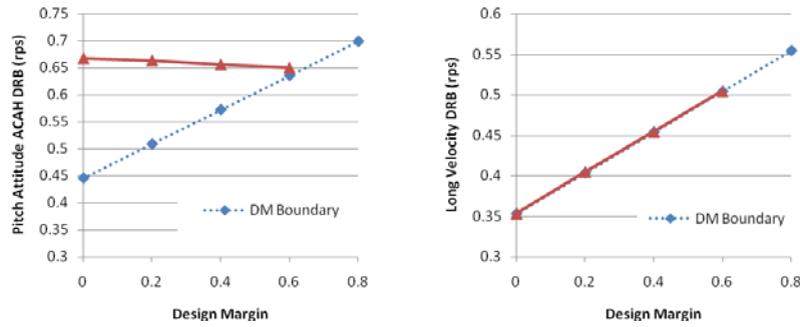


Fig. 36 Variation of disturbance rejection bandwidth with design margin.

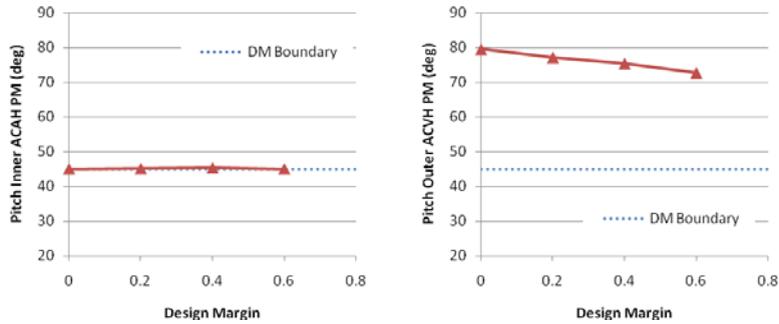


Fig. 37 Variation of phase margin with design margin.

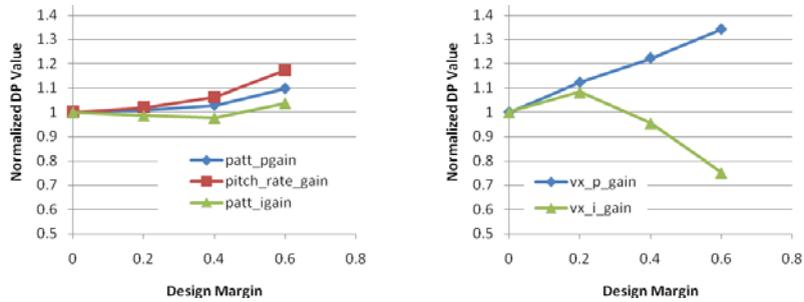


Fig. 38 Variation of design parameters with DM, normalized to first point.

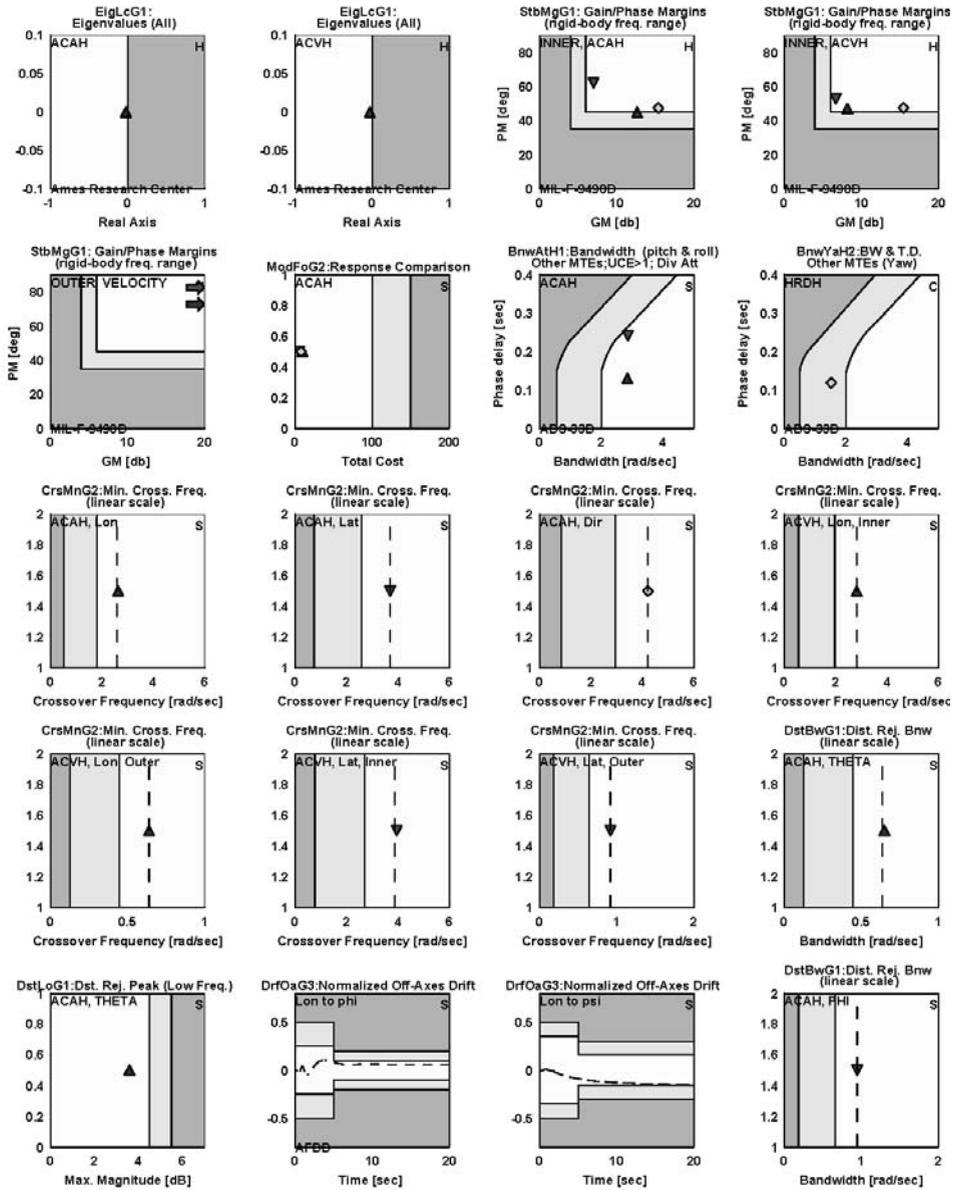


Fig. 39 CONDUIT[®] HQ Window (partial).

5.7 Preliminary Flight Tests

As mentioned earlier in Sec. 5.2, several hours of flight testing were conducted during June 2008 to collect frequency sweep data for analysis validation. Very preliminary piloted evaluations of the optimized gains, as they existed then were also conducted at the same time. The design has been matured and the current optimized gains are substantially different than those flown at that time, and therefore no flight data will be presented here. Formal flight testing of the latest designs is scheduled for early 2009. Nevertheless, it is worth mentioning that pilot comments indicated the preliminary optimized gains to be at least as good as gains hand-tuned for a previous effort. Pilot comments indicated the new gains, obtained automatically using multi-objective parametric optimization, to be better behaved, more predictable, and more precise.

6 Recent Application Pay-offs

The literature reports considerable payoffs in flight control performance, development time, and cost accrued from the use of the direct parametric-optimization design approach and associated tools. In the Garter case study (Grubel and Joos 1997), the researchers indicate a 30% improvement in control system performance achieved with MOPS as compared to manual tuning of the design parameters. This tool has been used as the basis for the successful flight control system design and development of the Eurofighter program (Moritz and Osterhuber 2006).

The UH-60 MCLAWS program (Tischler et al 2005) demonstrated significant improvements in handling-qualities and closed-loop performance relative to the baseline UH-60 control laws in only 6 hours of flight test development. Bell Helicopter engineers used direct design optimization (CONDUIT[®]) to tune AFCS gains on the Armed Reconnaissance Helicopter (ARH-70A) program to meet ADS-33 Level 1 requirements prior to initial flight testing. They credit this with being “able to clear the ARH-70A in less than one third of the flight test time that had been originally planned for AFCS development” (Christensen et al 2007).

Boeing Helicopter authors (Miller et al 2008) cite that for the V-22 and RAH-66 programs “some 60% of control law design and analytical integration activity occurred after first flight of the aircraft.” This resulted in labor and flight test intensive flight development programs. Under the HACT program, the integrated tools for design optimization (CONDUIT[®]) and the block diagram-to-flight autocoder (VISVEC[®]) were shown to provide a 240% improvement in design efficiency relative to the historical data for the V-22.

7 Conclusions

This paper has presented the motivation and essential elements of rotorcraft control system design using direct multi-objective function parametric optimization. The method provides a direct means to optimize complex flight control system architectures to meet numerous and competing design requirements. Specific conclusions from this paper are:

- 1) Based on the accumulated experience, when the design problem is properly posed (reasonable initial guesses for gains, specifications are selected properly and resulting Insensitivities and Cramer-Rao bounds of the design parameters are within the guidelines, the optimized designs are not sensitive to initial conditions.
- 2) Enforcing traditional stability margins and Nichols margins together with a design margin produces an optimized design that is robust to expected uncertainty.
- 3) There is an inherent trade-off between disturbance rejection and stability margin. Maximum gains and crossover frequency may well be limited by actuator rate saturation even when stability margins are still acceptable. The OLOP spec is an excellent means for monitoring actuator saturation in the context of the frequency-domain analysis and optimization.
- 4) Predesign analysis of performance robustness and model following performance provides a good means of establishing minimum crossover frequency for use in a complete model-following design optimization.
- 5) Integrator gains are well determined when both on- and off-axis drift specifications are included.
- 6) Multi-objective function optimization methods are well suited to the design of state-of-the-art multi-mode model-following rotorcraft flight control. The design margin optimization method produces a family of solutions of increasing performance, while coping with the complex interaction between the many competing requirements.
- 7) Published experience with multi-objective function optimization in the US and Europe has proven a benefit in improved performance and reduced cost.

References

- Anon (1987), "Flight Vehicle Development Time and Cost Reduction," AGARD Conference Proceedings No. 424, May.
- AGARD (2000), "Flight Control Design – Best Practices," NATO-RTO-029, December.
- Anon (1997), "Flying Qualities of Piloted Aircraft," USAF MIL-HDBK-1797, Dept. of Defense Interface Standard.17 December.
- Anon (2000), "Aeronautical Design Standard, Handling Qualities Requirements for Military Rotorcraft," USAAMCOM ADS-33E-PRF, U.S. Army Aviation and Missile Command, Huntsville, Alabama, March.
- Anon (2005), CONDUIT[®] User's Guide, Version 5.0, University Affiliated Research Center, Moffett Field, CA.
- Anon (2007), "Aerospace - Flight Control Systems - Design, Installation and Test of Piloted Military Aircraft, General Specification For," SAE-AS94900, July.
- Anon (2008), "Flight Control Systems - Design, Installation And Test Of Piloted Aircraft, General Specification For," MIL-DTL-9490E, Department of Defence, April.
- Advanced Rotorcraft Technology, Inc. (2001)., "FLIGHTLAB Theory Manual", April.
- Bendat, J. S., and Piersol, A. G. (1986), Random Data: Analysis and Measurement Procedures, 2nd ed., Wiley, New York.
- Blanken, C. et al (2008), "Test Guide for ADS-33E-PRF," AMRDEC Special Report AMR-AF-08-07, July.
- Blakelock (1991), J. H., Automatic Control of Aircraft and Missiles, 2nd ed., Wiley, New York.
- Blight, J. D., Dailey, R. L. and Gangsaas, D. (1996), "Practical Control Law Design for Aircraft Using Multivariable Techniques, Advances in Aircraft Flight Control, Taylor and Francis, London, pp. 231-267.
- Bryson (1994), Control of Spacecraft and Aircraft, Princeton University Press, Princeton, New Jersey.
- Boyd, S. P., Barratt, C. H. (1991), Linear Controller Design: Limits of Performance, Prentice Hall, Englewood Cliffs, New Jersey.
- Cheng, R. Tischler, M. B., Biezad, D. (1995), "Rotorcraft Flight Control Design Using Quantitative Feedback Theory," Symposium on Quantitative and Parametric Feedback Theory, Purdue University, West Lafayette, Indiana, 2-4 August.
- Christensen, Kevin T., Campbell, Kip G., Griffith, Carl D., Ivler, Christina M., Tischler, Mark B., Harding, Jeffrey W. (2007), "Flight Control Development for the ARH-70 Armed Reconnaissance Helicopter Program," American Helicopter Society 63rd Annual Forum, Virginia Beach, Virginia, May.
- Crawford, C., C. (1998), "Potential for Enhancing the Rotorcraft Development/Qualification Process Using System Identification," Presented at NATO RTO SCI Symposium on "System Identification for Integrated Aircraft Development and Flight Testing, Madrid, Spain, 5-7 May.
- Dryfoos, J. B., Kothmann, B. D., Mayo (1999), J., "An Approach to Reducing Rotor-Body Coupled Roll Oscillations on the RAH-66 Comanche Using Modified Roll Rate Feedback," American Helicopter Society 55th Annual Forum, Montréal, Canada, 25-27 May. pp 1127-1140.
- Duda, H. (1997), "Prediction of Pilot-in-the-Loop Oscillations due to Rate Saturation", Journal of Guidance, Navigation, and Control, Vol. 20, No. 3, May-June.
- Duus, Gunnar (2000), "SCARLET 3 - A Flight Experiment Considering Rate Saturation", AIAA Atmospheric Flight Mechanics Conference and Exhibit, Denver, Colorado, 14-17 August.

- “ADS-33E Predicted and Assigned Low-Speed Handling Qualities of the CH-47F with Digital AFCS,” American Helicopter Society 63th Annual Forum, Virginia Beach, Va, 1-3 May.
- Fan, M. K. H., Tits, A. L., Zhou, J. L., Wang, L. S., and Koninckx, J. “CONSOL-OPTCAD , User’s Manual,” Technical Report 87-212r2a, Dept. of Electrical Engineering and Institute for Systems Research, University of Maryland, College Park, Md.
- Fletcher et al (2008), “UH-60M Upgrade Fly-By-Wire Flight Control Risk Reduction using the RASCAL JUH-60A In-Flight Simulator,” American Helicopter Society 64th Annual Forum, Montréal, Canada, April 29 - May 1.
- Franklin, J. A. (2002), Dynamics, Control and Flying Qualities of V/STOL Aircraft, AIAA Education Series.
- Grubel, G. and Joos, H-D (1997), “Multi-Objective Parameter Synthesis (MOPS),” Lecture Notes in Control and Information Sciences 224 (Magni and Bennani and Terlouw, Editors), pp 13-21.
- Harding, J. W. et al (2006), “Development of Modern Control Laws for the AH-64D in Hover/Low Speed Flight,” American Helicopter Society 62nd Annual Forum, Phoenix, AZ, May 9-11.
- Hodgkinson, J. (1998), Aircraft Handling Qualities, AIAA Educational Series, AIAA, Reston, VA.
- Horowitz, I. M., Synthesis of Feedback Systems, Academic Press.
- Howlett, J. J. (1981), “UH-60A Black Hawk Engineering Simulation Program: Volume I – Mathematical Model,” NASA CR-166309, December.
- Ivler, C. (2008), System Identification Modelling for Flight Control Design, Rotorcraft Handling Qualities Conference, Liverpool, U.K., Nov.
- Kreisselmeier, G., Steihauser, R. (1979), “Systematic Control Design by Optimizing a Vector Performance Index,” Proceedings of IFAC Symposium on Computer-Aided Design of Control Systems, Zurich, pp 113-117.
- Kreisselmeier, G., Steihauser, R. (1983), “Application of Vector Performance Optimization to a Robust Control Loop Design for a Fighter Aircraft,” International Journal of Control, Vol 37, No. 2, pp 251-284.
- Landis, K. H., Davis, J. M., Dabundo, C., and Keller, J. F.(1996), “Advanced Flight Control Research and Development at Boeing Helicopters,” Advances in Aircraft Flight Control, Taylor and Francis, London, pp. 103–142.
- Maciejowski, J. M. Multivariable Feedback Design, Addison-Wesley Publishing Company.
- Magni, J-F, Bennani, S., Terlouw, J (Eds) (1997), “Robust Flight Control – A Design challenge,” Lecture Notes in Control and Information Sciences 224.
- McLean, D. (1990), Automatic Flight Control Systems, Prentice–Hall International, Hertfordshire, UK.
- McRuer, D. T., Ashkenas, I. L., and Graham, D. (1973), Aircraft Dynamics and Automatic Control, Princeton Univ. Press, NJ.
- Miller, D. G., White, J. F., Taylor, J. E., Lukes, G. P., Gradle, R. P., Segner, D. (2008), “HACT Program Technology Transfer,” American Helicopter Society 64th Annual Forum, Montréal, Canada, April 29 - May 1.
- Mitchell, D. G., Hoh, R. H., Morgan, J. M. (1989), “Flight Investigation of Helicopter Low-Speed Response Requirements,” Journal of Guidance, Control, and Dynamics, Vol. 12, No. 5, Oct-Sept, pp. 623-630.
- Moritz, M., Osterhuber, R. (2006), “Three-Stage Gradient-Based Optimization Scheme in Design of Feedback Gains within Eurofighter Primary Control Laws,” AIAA Guidance, Navigation, and Control Conference and Exhibit, Keystone Colorado, 21-24 August.
- Nye, W. T. (1983), DELIGHT: An Interactive System for Optimization-Based Engineering Design, Ph.D Thesis, Dept of EECS, University of California, Berkeley.

- Nye, W. T., Tits, A., L. (1986), "An Application-oriented, optimization-based methodology for interactive design of engineering systems," *International Journal of Control*, Vol. 43, No. 6, pp1693-1721.
- Padfield, G. D. (1996), *Helicopter Flight Dynamics: The Theory and Application of Flying Qualities and Simulation Modeling*, AIAA, Reston, VA, p. 137.
- Powell, M. J. D. (1978), "A Fast Algorithm for Nonlinearly Constrained Optimization Calculations" in *Numerical Analysis*, G. A. Watson, Ed, Lecture Notes in Mathematics, vol. 630, Springer, Berlin.
- M.J.D. Powell (1983), On the quadratic programming algorithm of Goldfarb and Idnani, Report DAMTP 1983/Na 19, University of Cambridge, Cambridge.
- Pratt, R. W. (Organizing Editor) (2000), *Flight Control Systems*, AIAA Progress in Astronautics and Aeronautics, Vol 184.
- Rozak, J.N. and Ray, A. (1998), "Robust Multivariable Control of Rotorcraft in Forward Flight: Impact of Bandwidth on Fatigue Life," *J. of the American Helicopter Society*, Vol. 43, No. 3, pp. 195-201, July.
- Sandwell, D. T., (1987), "Biharmonic Spline Interpolation of Geos-3 and Seasat Altimeter Data," *Geophysical Research Letters*, Vol. 14, No. 2, pp 139-142.
- Schittkowski, K. (2005) "QL: A Fortran Code for Convex Quadratic Programming - User's Guide, Version 2.11," Department of Computer Science, University of Bayreuth.
- Tischler, M. B. (1987), "Digital Control of Highly Augmented Combat Rotorcraft," NASA-TM 88346.
- Tischler, M. B. (organizing ed.) (2007), *Advances in Aircraft Flight Control*, Taylor and Francis, London.
- Tischler, M. B. and Remple, R. K. (2006), *Aircraft and Rotorcraft System Identification: Engineering Methods With Flight Test Examples*, AIAA, Aug.
- Tischler, M. B., et al (1999), "A Multidisciplinary Flight Control Development Environment and Its Application to a Helicopter," *IEEE Control Systems Magazine*, Vol. 19, No. 4, pg. 22-33, August.
- Tischler, M., B., Lee, J. A., Colbourne, J. D. (2002) "Comparison Of Alternative Flight Control System Design Methods Using the CONDUIT® Design Tool," *AIAA Journal of Guidance, Control, and Dynamics*, Volume 25, No. 3, pg. 482-493, May-June.
- Tischler, M. B., Blanken, C. L., Cheung, K. K., Swei, S. S. M., Sahasrabudhe, V., and Faynberg, A. (2005), "Modernized Control Laws for UH-60 BLACK HAWK: Optimization and Flight Test Results," *AIAA Journal of Guidance, Control, and Dynamics*. Vol 28, No. 5, Sept-Oct, pp 964-978.
- Yaniv, O. (1999), *Quantitative Feedback Design of Linear and Nonlinear Control Systems*, Kluwer Academic Publishers, Mass.