Achieving the Best Compromise between Stability Margins and Disturbance Rejection Performance

Mohammadreza H. Mansur
mohammadreza.mansur@us.army.mil
Jeff A. Lusardi
jeff.lusardi@us.army.mil
Mark B. Tischler
mark.tischler@us.army.mil
Aeroflightdynamics Directorate (AMRDEC)
U.S. Army RDECOM
Moffett Field, CA, USA

Tom Berger
tom.berger@us.army.mil
UC Santa Cruz (UARC)
Moffett Field, CA, USA

Abstract

Optimization of control law parameters such that all stability, handling qualities, and robustness requirements are satisfied is a complex process, especially for rotorcraft, due to the competing nature of the various specifications. Of these, the trade-off between stability margins and disturbance rejection performance has been of significant recent interest. However, the question of the possible benefits of allowing a reduction in stability margin requirements to further increase disturbance rejection performance had never been evaluated with pilot ratings in flight. A research effort was therefore conducted at the U.S. Army Aeroflightdynamics Directorate (AFDD) to use multi-objective parametric optimization to generate families of optimized designs with increasing levels of disturbance rejection performance in exchange for reductions in stability margin requirements. Two designs, one satisfying the current stability margin requirements and another allowing some relaxation of these requirements, were then flight tested aboard AFDD's full authority fly-by-wire JUH-60 helicopter (RASCAL). Two pilots flew a set of ADS-33 hover and low speed Mission Task Elements (MTE's) using the two designs and pilot performance, Handling Qualities Ratings (HQRs), and comments were collected. Results of the limited flight test indicate that the design enforcing the current stability margin requirements felt more natural and predictable to the pilots while the design with increased disturbance rejection bandwidth and relaxed stability margins felt nervous, causing a slight tendency to over control.

Acronyms

ACAH	Attitude Command Attitude Hold
ARH	Armed Reconnaissance Helicopter
DM	Design Margin
DRB	Disturbance Rejection Bandwidth
GM	Gain Margin (dB)
HQR	Handling Qualities Rating
MTE	Mission Task Element
OLOP	Open-Loop Onset Point
PH	Position Hold
PI	Proportional/Integral
PID	Proportional/Integral/Derivative
PIO	Pilot Induced Oscillations
PM	Phase Margin (deg)
RASCAL	Rotorcraft Aircrew Systems Concepts Airborne Laboratory
RCDH	Rate Command Direction Hold
RM	Relaxed Margins
rps	radians per second
SM	Standard Margins
VH	Velocity Hold
VTUAV	Vertical Take-off Unmanned Aerial Vehicle

Presented at the American Helicopter Society 65th Annual Forum, Grapevine, Texas, May 27-29, 2009.

Introduction

The trade-off between system stability and system disturbance rejection performance is fundamental to flight control design. System stability requirements, as measured by stability margins, ensure robustness in the presence of uncertainties due to flight conditions, modeling accuracy, and normal life-cycle wear. These requirements have been specified in various requirement-documents, including the MIL-F-9490D [1], which has recently been superseded by the SAE AS94900 [2]. The current stability margin requirements call for 6 dB of gain margin (GM) and 45 deg of phase margin (PM), to be measured at the most critical design configuration and mission flight condition. A degradation of up to 50% from these levels is allowed when evaluating the effect of large uncertainties in aircraft response characteristics.

Disturbance rejection performance requirements, on the other hand, ensure that hold functionalities of the system satisfactorily reject atmospheric disturbances, a characteristic that is highly desirable in degraded weather and visibility. Unlike stability margin requirements, disturbance rejection performance requirements for rotorcraft have only recently been addressed. The test guide for ADS-33E-PRF [3] presents a new disturbance rejection bandwidth (DRB) specification, developed and flight-test validated by the U.S. Army's Aeroflightdynamics Directorate (AFDD), and proposes criteria values for attitude disturbances (ACAH response type). Still, currently no DRB criteria values exist for outer loops such as velocity hold (VH) and position hold (PH).

Stability margin and disturbance rejection performance are conflicting requirements because increasing performance generally requires increasing gains which in turn leads to lower margins. Achieving a trade-off between the two is therefore needed to achieve the best design, as considered in several recent projects including the development of modernized control laws for the AH-64D Longbow Apache [4], control law development efforts on the Fire Scout VTUAV [5], and the development of the Armed Reconnaissance Helicopter (ARH) [6]. These projects have all approached this trade-off by generating families of designs with increasing disturbance rejection performance but have always enforced the current stability margin requirements. However, recent flight test experience, for example with the CH-47F Digital Automatic Flight Control System (DAFCS) development program [7], have indicated that good handling qualities may be obtained even at stability margins lower than the current requirements. Therefore, a systematic look at the implications of relaxing the requirements on stability margins to allow further increases in disturbance rejection performance was carried out, based on the JUH-60 RASCAL (Fig. 1) [8] full-authority flight control system.

Figure 1 – The RASCAL JUH-60A variable stability helicopter

A multi-mode model following architecture was selected for this effort as it provides a 2D design space where the system bandwidth is mainly dependent on the command model. System bandwidth can, therefore, be kept effectively constant while stability margins and disturbance rejection performance are traded-off. The overall approach taken for this UH-60 trade-off study was to 1) start with two identical systems with identical sets of stability margins, handling qualities, disturbance rejection, and actuator activity requirements, 2) relax the stability margins requirements on one system, 3) optimize both systems for increasing disturbance rejection performance until maximum possible performance was reached for each system, resulting in a family of results for each case, 4) select one system from each family and carry through side-by-side flight testing aboard RASCAL, and, finally, 5) use pilot performance, Handling Qualities Ratings (HQRs), and comments gathered while flying ADS-33 hover and low speed Mission Task Elements (MTEs) to compare the two.

This paper presents the system used to carry out the trade-off study and the details of the analysis model. It then presents the results of the trade-off analysis followed by a brief overview of the flight test. Finally, the results of the flight test are presented including a summary of the pilot ratings and pilot comments.

System Description

The control laws in this study were full authority, multi-mode, and optimized for hover and low speed flight only. They provided an attitude command / attitude hold (ACAH) response type with ground speed (VH) and position (PH) hold modes in the longitudinal and lateral axes. The VH functionality was automatically activated with the cyclic stick in detent in each axis while the PH functionality was automatically activated when the ground velocity dropped below a prescribed threshold. In

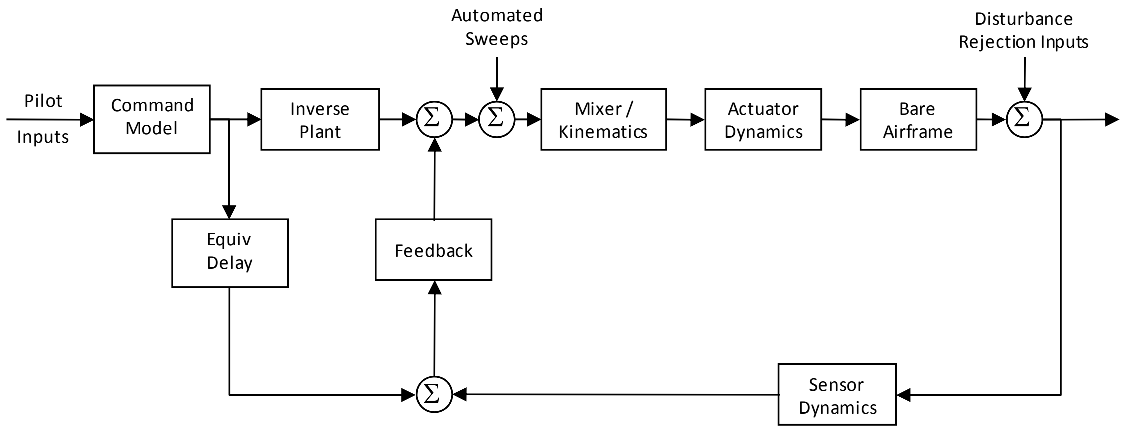

Figure 2 – Overview of control system

the directional axis, the control laws provided a rate command / direction hold (RCDH) response type, with the hold functionality automatically activated with the pedals in detent. Altitude hold functionality was also provided, but note that the heave channel used a simple response feedback architecture instead of model following, and the characteristics of the vertical axis response were constant throughout the trade-off study.

The basic feedback architecture for achieving ACAH in pitch and roll and RCDH in yaw was PID in all axes. The velocity hold functionality was implemented as outer loop PI architectures on the longitudinal and lateral ground velocities and the position hold was, in turn, wrapped around these velocity hold loops. An overview of the system is shown in Fig. 2 and its various components are discussed in the following sections. The analysis model for this system, which was based on a simplified set of Simulink® block diagrams implemented in the CONtrol Designers Unified InTerface (CONDUIT®) [9], is also discussed.

Bare Airframe Model --- The mathematical model of the aircraft constitutes the core of the control law analysis. Therefore, the availability of a model that accurately represents the bare airframe dynamics of the vehicle is critically important. In this work a highly accurate linear model of the RASCAL JUH-60A, identified from flight data using the Comprehensive Identification from FrEquency Responses (CIFER®) [10] tool, was used. The states of the model (22 total states) included the rigid body states, rotor flapping and lead lag, rotor RPM and torque, and rotor inflow states. The model described the dynamics of the RASCAL aircraft from swashplate input (4 total inputs) to aircraft responses. The accuracy of this model, both independently and as imbedded in the control laws, was validated against available flight data and noted discrepancies corrected using appropriate gain and time delay modifications, as will be discussed later in this paper.

Command Models --- The assumption inherent in the use of a model following architecture is that improved handling qualities can be achieved if the aircraft responses can be made to match the responses of low-order command models. The command models used for this work were 2nd order in pitch and roll attitudes and 1st order in yaw rate. Since in a model following implementation the bandwidths of the closed-loop system are primarily dependent on the bandwidths of the command models, the command models' parameters were selected to satisfy ADS-33 [11] bandwidth requirements. Additionally, to avoid overdriving unmodeled high-order dynamics not included in the low-order inverse (discussed next), time delays equivalent to the forward path time delays of the system were added to the commanded responses (feedback), as shown in Fig. 2.

Inverse Plant Model --- In a model following implementation, an inverse of the vehicle dynamics is used in the forward path to generate control inputs necessary to achieve the commanded responses. The accuracy to which the actual responses track the commanded responses is for the most part determined by the accuracy of this inverse. However, developing an accurate high-order inverse of the vehicle is generally not feasible and in practice only approximate low-order pseudo inverses are used which ignore the vehicle's higher-order dynamics. Therefore, feedback is needed to null out the resulting command following errors, in addition to providing stability and performance robustness. In the current implementation, simple first-order inverses were used in pitch, roll, and yaw. Note that as mentioned earlier the heave channel is not model following and uses only response feedback.

Actuator Models --- Since the RASCAL aircraft was the intended flight test platform for this trade-off study, the analysis was configured for this aircraft. Unlike a standard UH-60A Black Hawk, the RASCAL aircraft has two sets of actuators: 1) standard UH-60A primary actuators (Fwd, Aft, Lat, TR) and 2) RASCAL Research Servos (one per primary servo). All 8 actuators were modeled in the analysis as second-order systems with rate and position limiting. As discussed by Tischler et. al. [12] actuator rate saturation can have a significant detrimental effect on handling qualities and can directly lead to pilot induced oscillation tendency and therefore must be included in the analysis model and design requirements.

Sensor Models --- The analysis included representations of sensor filters used in the feedback path of the control laws to reduce signal noise and eliminate unwanted frequency content from the feedback data.

Analysis Model Validation

Just as it is critically important to start with an accurate model of the vehicle, it is also critically important to verify that the analysis model as a whole represents the real system accurately. In the case of this research, there was a unique opportunity to test the real system with a baseline set of gains and validate the analysis even before the optimization of the control laws was started. Since the characteristics of main interest were stability margins and disturbance rejection performance, it was important to ensure that the analysis predicted these characteristics of the actual system very accurately. Therefore, specialized flight tests were carried out during which automated sweeps were directly injected at the mixer input and into the sensors (Fig. 2) to calculate broken loop and disturbance rejection characteristics of the actual system, respectively, for analysis validation. The results of the subsequent validation work for the longitudinal, lateral, and directional axes are shown below. It should be noted that the vertical axis was also validated with comparable results. Finally, piloted sweeps were used to validate the end-to-end responses of the system, again with comparable results. These are not shown here for brevity.

Broken Loop

Automated sweeps at the mixer input (Fig. 2) were used in flight to generate time responses using a baseline set of gains. The pilots were instructed to minimize control inputs as much as possible and limit their inputs to pulse-type corrective inputs, as needed to prevent the aircraft from drifting too far away from trim. Therefore, the system was essentially in VH mode for these runs as the cyclic stick and pedals were mostly left in detent.

The resulting time responses were processed with CIFER® to generate broken loop frequency responses for each axis. These were then compared to frequency responses obtained from the analysis. Initial results indicated that even though the analysis model showed reasonable correlation, there was still room for further improvement. The slight discrepancies noted were determined to be partly due to inaccuracies in the mathematical model of the bare airframe and partly due to un-modeled dynamics. To eliminate these discrepancies as much as possible, CIFER® was used to calculate single gain and time delay values for each axis that would further improve the accuracy of the analysis. These gain and time delay values were then applied to the analysis to achieve improved correlation.

Figures 3 and 4 compare the broken loop responses of the corrected analysis model, in pitch and roll, with flight data. Note that two sets of flight data are shown to indicate the possible variations between different flight results due to aircraft state, flight conditions, and data gathering and processing effects. As may be seen, the gain and phase correlations are very good, especially in the region of crossover which is of most concern. The accuracy of the analysis can be further demonstrated by comparing the analysis crossover frequencies and gain and phase margins with those obtained from flight. This is done in Table 1, which indicates very good correlation for pitch and roll. For example, in the case of roll, the analysis predicts a crossover frequency of 3.0 rps compared to an actual value of 3.1 rps, a phase margin of 54.6 deg compared to 53.4 deg, and a gain margin of 7.3 dB compared to 7.0 dB.

Figure 5 compares the broken loop response of the corrected analysis in yaw with flight data results. In this case there is some discrepancy in the analysis results, as is further indicated by the mismatch of yaw crossover frequency and margins shown in Table 1. This discrepancy was traced to a dynamic mode not modeled by the analysis and could not be directly addressed. However, as Fig. 5 shows, the correlation between the analysis and flight is quite good for frequencies between 1.5 and 4.5 rps. Therefore, the phase margin analysis would still be highly accurate as long as the yaw crossover frequencies for the final optimized systems remained below 4.5 rps.

Table 1 – Phase and gain margin correlation results

		Flight Test Results (Average)	Analysis Results (Corrected)
Pitch	GM	9.2 dB @ 9.4 rps	11.4 dB @ 10.5 rps
	PM	49.8 deg @ 3.5 rps	48.9 deg @ 3.3 rps
Roll	GM	7.0 dB @ 8.9 rps	7.3 dB @ 8.7 rps
	PM	53.4 deg @ 3.1 rps	54.6 deg @ 3.0 ros
Yaw	GM	5.5 dB @ 17.1 rps	13.2 dB @ 19.4 rps
	PM	47.1 deg @ 6.3 rps	56.1 deg @ 5.5 rps

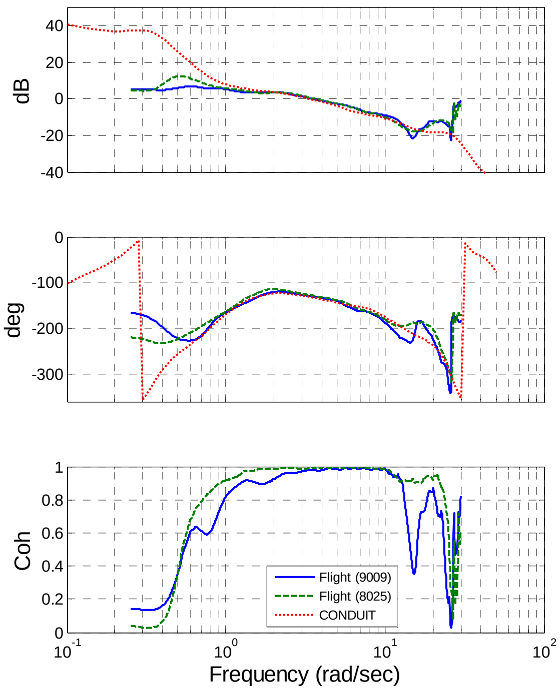

Figure 3 – Longitudinal broken loop response comparison between analysis and flight

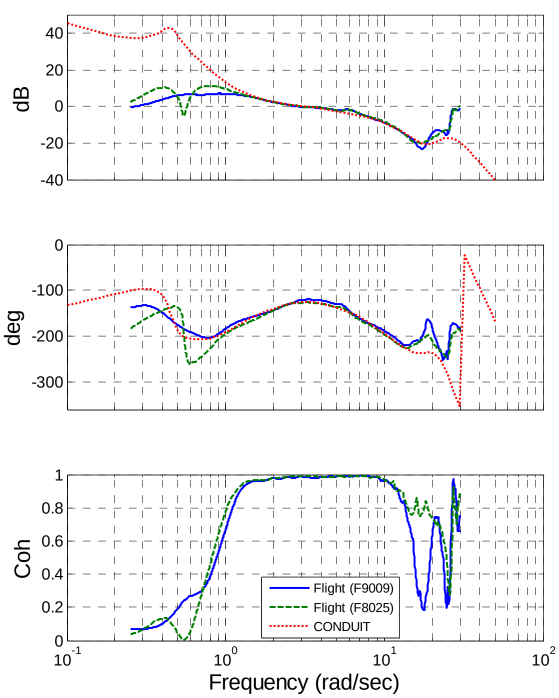

Figure 4 – Lateral broken loop response comparison between analysis and flight

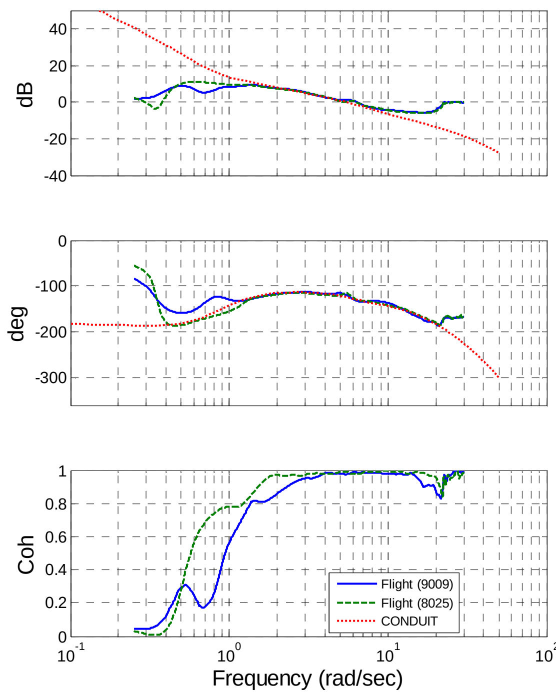

Figure 5 – Directional broken loop response comparison between analysis and flight

Disturbance Rejection

Automated sweeps into the sensors (Fig. 2) were used in flight to generate time responses using a baseline set of gains. As before, the pilots were instructed to minimize control inputs as much as possible and limit their inputs to pulse-type corrective inputs as needed to prevent the aircraft from drifting too far away from trim. Unlike the broken loop runs, however, in the case of the disturbance runs the system remained in ACAH mode even with the stick in detent.

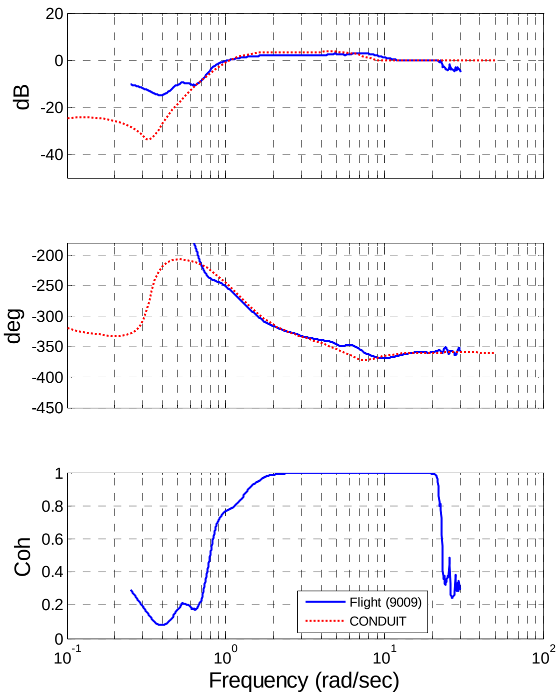

Figure 6 – ACAH lateral disturbance rejection response comparison between analysis and flight

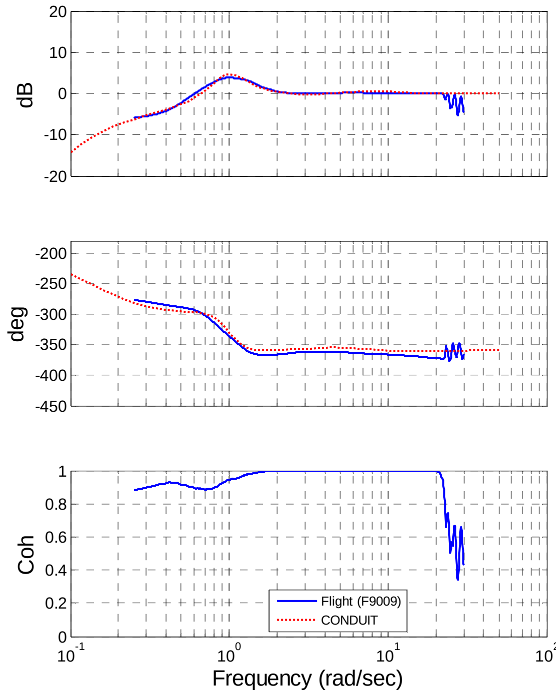

Figure 7 – ACVH longitudinal ground speed disturbance rejection response comparison between analysis and flight

The resulting time responses were processed with CIFER® to generate disturbance rejection frequency responses for attitude and ground speed hold modes. These were then compared to frequency responses obtained from the analysis. Results for roll attitude and longitudinal ground speed are shown in Figs. 6 and 7, respectively. As may be seen, the analysis results match the flight results quite well. This same level of accuracy was seen in the pitch and yaw attitudes and in the lateral ground speed as well (they are not shown here for brevity).

The accuracy of the analysis model can be further demonstrated by comparing the analysis DRB values with those obtained from flight. This is done in Table 2, which indicates very good correlation between analysis and flight. For example, in the case of roll attitude disturbance the analysis predicts a DRB of 0.88 rps which is an excellent match of the actual value of 0.84 rps.

Table 2 – Disturbance rejection bandwidth (DRB) correlation results

	Flight Test (Flight 9009)	Analysis
Roll Att DRB	0.84 rps	0.88 rps
Pitch Att DRB	0.59 rps	0.55 rps
Yaw Att DRB	0.80 rps	0.74 rps
Lon Grnd Spd DRB	0.47 rps	0.49 rps
Lat Grnd Spd DRB	0.82 rps	0.92 rps

Multi-Objective Parameter Optimization

The main goal of the optimization process was to carry out a trade-off between stability margin and DRB. However, since the final optimized systems were intended for flight-testing to gather pilot performance and opinion data, all optimized systems had to meet Level 1 handling qualities requirements. Ensuring Level 1 handling qualities in the optimization results would insure that the pilots would be able to concentrate on, and discern differences in, response characteristics resulting mainly from the variations in stability margins and disturbance rejection. Therefore, a complete set of requirements for adequate stability, handling qualities, cross coupling, actuator saturation, disturbance rejection, and actuator activity were included in the analysis, as will be discussed shortly.

As mentioned earlier, in a model following implementation the piloted bandwidth is largely dependent on the command model and independent of feedback. As such, it would be possible for the optimization process to lower the crossover frequencies to undesirably low values while satisfying all stability and handling qualities requirements. Excessively low crossover frequencies reduce the performance robustness of the system for off-nominal conditions (where inverse plant performance is degraded). Therefore, while the systematic trade-off between stability margin and DRB is carried out crossover frequency minimums must be maintained. These minimums can be determined in preliminary design based on model-following accuracy and performance robustness considerations [12]. An example analysis from reference [12] is shown in Fig. 8, indicating that a minimum crossover frequency value of around 2.5 rps is needed to satisfy model following and performance robustness requirements of the UH-60 aircraft.

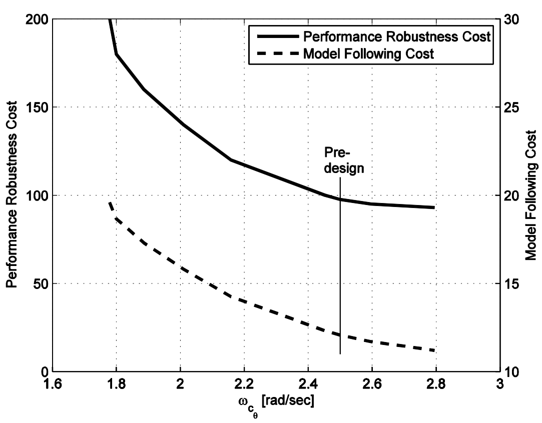

Figure 8 – Variations in model following and performance robustness with changes in crossover frequency (from reference [12])

Several approaches could have been followed to ensure that the crossover frequency minimums, determined in preliminary design, were achieved. The direct approach is to optimize the system at all off-nominal conditions. However, for all but the simplest of systems the computational cost of this direct approach would be prohibitive. Another approach is to constrain the crossover frequencies to fixed minimums, based on preliminary design results, while disturbance rejection performance is increased. A third approach is to start the minimum crossover constraints at the preliminary design values but increase the minimum constraints as the DRB requirements are increased. The third approach was employed in this work in order to arrive at solutions that not only provided improved disturbance rejection performance but also improved performance robustness for off-nominal conditions.

Optimization Approach

The Control Designer's Unified Interface (CONDUIT®) tool [9] was used to carry out the analysis and optimization work for this trade-off study. The optimization started with two identical CONDUIT® cases with identical specifications. The complete list of specifications is shown in Table 3. One of the two cases was then modified by relaxing the stability margin requirements from the current GM = 6 dB and PM = 45 deg down to GM = 4 dB and PM = 35 deg. The CONDUIT® case with the *relaxed* stability *margin* requirements will henceforth be referred to as the "RM" case while the case with the *standard* stability *margin* requirements will be referred to as the "SM" case. All other specifications were exactly the same for both cases.

Note that the specifications are divided into 4 categories, as seen in Table 3. "Hard constraints" (H-type in Table 3) are specs that relate to the stability of the system and have to be satisfied ahead of all others. "Soft constraints" (S-type in Table 3) are specs that relate to handling qualities of the system and are satisfied after all hard constraints are satisfied. When all the stability-related (hard) and handling-qualities-related (soft) constraints are satisfied, a viable, though not yet optimal, system is reached. Further optimization to ensure satisfaction of all the requirements without overdesign and with minimum achievable gains (cost of feedback) is then carried out by minimizing the sum of all "summed objectives" (J-type in Table 3). The final type of specification used is the "check only" type (C-type in Table 3) which are specifications that are not considered as part of the optimization but whose values are calculated and presented. More detail on the optimization process is given in references [9] and [12].

Referring to Table 3, two specifications warrant further discussion. First is the "open-loop onset point" spec [13] for category II pilot-induced oscillations (PIO), which ensures acceptable actuator saturation characteristics and low PIO tendency. This spec accounts for actuator rate and position limits, which are otherwise often ignored in linear analysis and can lead to significant actuator saturation on the real system.

The second is the new "normalized off-axes drift" spec, which was developed specifically for this effort to address off-axes drift concerns raised by the pilots in preliminary flight tests. This spec is an extension of the ADS-33 attitude hold specification and applies to the off-axes responses. The spec requires that the magnitudes of the off-axes responses should never exceed 25% of the maximum magnitude of the on axis response and should all reduce to less than 10% of the maximum on-axis response within 5 seconds. In practice, this spec controls the off-axes drift through governing the integral gains of the control system, which in the past have often been the system parameters of concern because of a lack of appropriate specifications for their control.

Table 3 – Specifications

Name	Description	Type	Comments
EigLoG1	Eigenvalues	H	Ensure stability
StbMgG1	Gain/Phase Margins (rigid-body frequency range)	H	Ensure adequate stability margins (MIL-F-9490D) Margins have to be checked at various points so multiple copies of this spec are used
EigDpG1	Generic Damping Ratio	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	Yaw Bandwidth. Other MTEs (Yaw)	S	Short term yaw response requirement (ADS-33D)
CrsMnG2	Min. Crossover Freq. (linear scale)	S	Ensure acceptable crossover frequencies
OlpOpG1	Open Loop Operating Point Rate Limit Saturation Spec.	S	Ensure acceptable actuator saturation characteristics and low PIO tendency
DstBwG1	Disturbance Rejection Bandwidth (linear scale)	S	Ensure satisfactory disturbance rejection bandwidth
DstLoG1	Disturbance Rejection Peak Magnitude (Low Freq.)	S	Ensure good damping of disturbance response
DstNmG2	Generic Disturbance Response	S	Control steady state 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	Pitch-Roll Coupling Frequency Domain	C	Ensure good pitch/roll coupling (ADS-33D)
FrqHeH1	Heave Response Hover/Low-Speed	C	Ensure good heave dynamics (ADS-33D)

Table 4 – Design Parameters for trade-off optimization

	Long	Lat	Dir
Angular Rate Proportional Gain	K_q	K_p	K_r
Angular Attitude Proportional Gain	K_tht	K_phi	K_psi
Angular Attitude Integral Gain	K_tht_i	K_phi_i	K_psi_i
Ground Speed Proportional Gain	K_Vx	K_Vy	
Ground Speed Integral Gain	K_Vx_i	K_Vy_i	
Position Proportional Gain	K_X	K_Y	

In CONDUIT®, the system parameters that are designated as "tuning knobs" in the optimization process are referred to as "Design Parameters" (DPs). Note that not every system parameter is a DP as most systems contain many additional parameters which remain constant throughout the optimization. A total of 15 system parameters were designated as DPs in the current analysis, as listed in Table 4. Note that of these only 13 DPs were allowed to vary freely during the optimization. The integral gains on longitudinal and lateral ground velocities (Vx and Vy) were constrained to their corresponding proportional gains using $\frac{K\_Vx,y\_i}{K\_Vx,y} = \frac{\omega_c}{10}$, where ω_c is the corresponding nominal crossover frequency.

The concept of a "Design Margin," available in CONDUIT®, was used to systematically increase the DRB and crossover frequency requirements for the SM and RM cases. In CONDUIT® all specs are divided into 3 regions which roughly correspond to the 3 handling qualities levels of the Cooper Harper rating scale [14], namely Level 1: Acceptable without improvement, Level 2: Deficiencies warrant improvement, and Level 3: Deficiencies require improvement. The goal of the optimization is then to determine Design Parameters that would allow all specifications to be in Level 1 with minimum overdesign. The concept of a Design Margin is based on these levels. A non-zero positive Design Margin in effect moves the boundary between Levels 1 and 2 in the direction of the Level 1 region, making the requirements more stringent and more difficult to satisfy. The amount by which the Level 1 / Level 2 boundary is moved into the Level 1 region is equal to the product of the Design Margin and the width of the Level 2 region of the spec. So, for example, a Design Margin of 0.2 means that the Level 1 / Level 2 boundary has moved 20% of the width of the Level 2 region into the Level 1 region. This is depicted in Fig. 9. Note that Design Margin can be activated on a per spec basis and was only applied to DRB and minimum crossover frequency specs for this effort.

Design Margin Optimization was then used to optimize both the SM and RM cases with incrementally increasing requirements on DRBs and crossover frequencies. Design Margin Optimization is a CONDUIT® capability that automates the process of systematically varying Design Margin values and optimizing the system for these values in a batch process. The tool also automatically documents all the intermediate and final results and produces comparison charts.

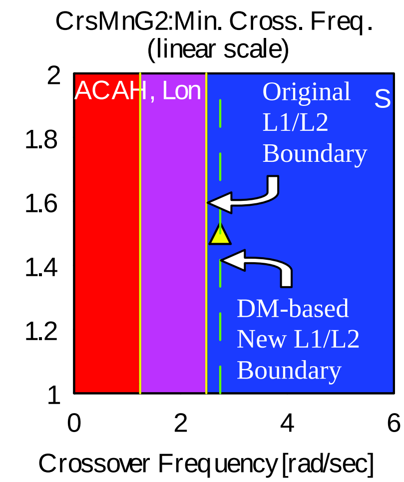

Figure 9 – Effect of Design Margin

Optimization Results

The SM case was successfully optimized for Design Margins between 0.0 and 0.4 but optimization could not proceed past a Design Margin of 0.4 without violating the stability margin requirements. For the RM case the optimization was successful for Design Margins up to 0.6 but could not proceed further. Note that since the two cases were identical (except for the stability margin requirements) the interpretation of the Design Margin was also identical between the two.

A partial view of the final CONDUIT® HQ Window for the 0.2 Design Margin case (SM family) is shown in Fig. 10. As may be seen, all the specifications are satisfied in Level 1. The crossover frequencies have been optimized to the Design-Margin-augmented boundaries and the model following spec shows very low costs for all axes, indicating that the actual responses follow the commanded responses very closely. The yaw drift due to

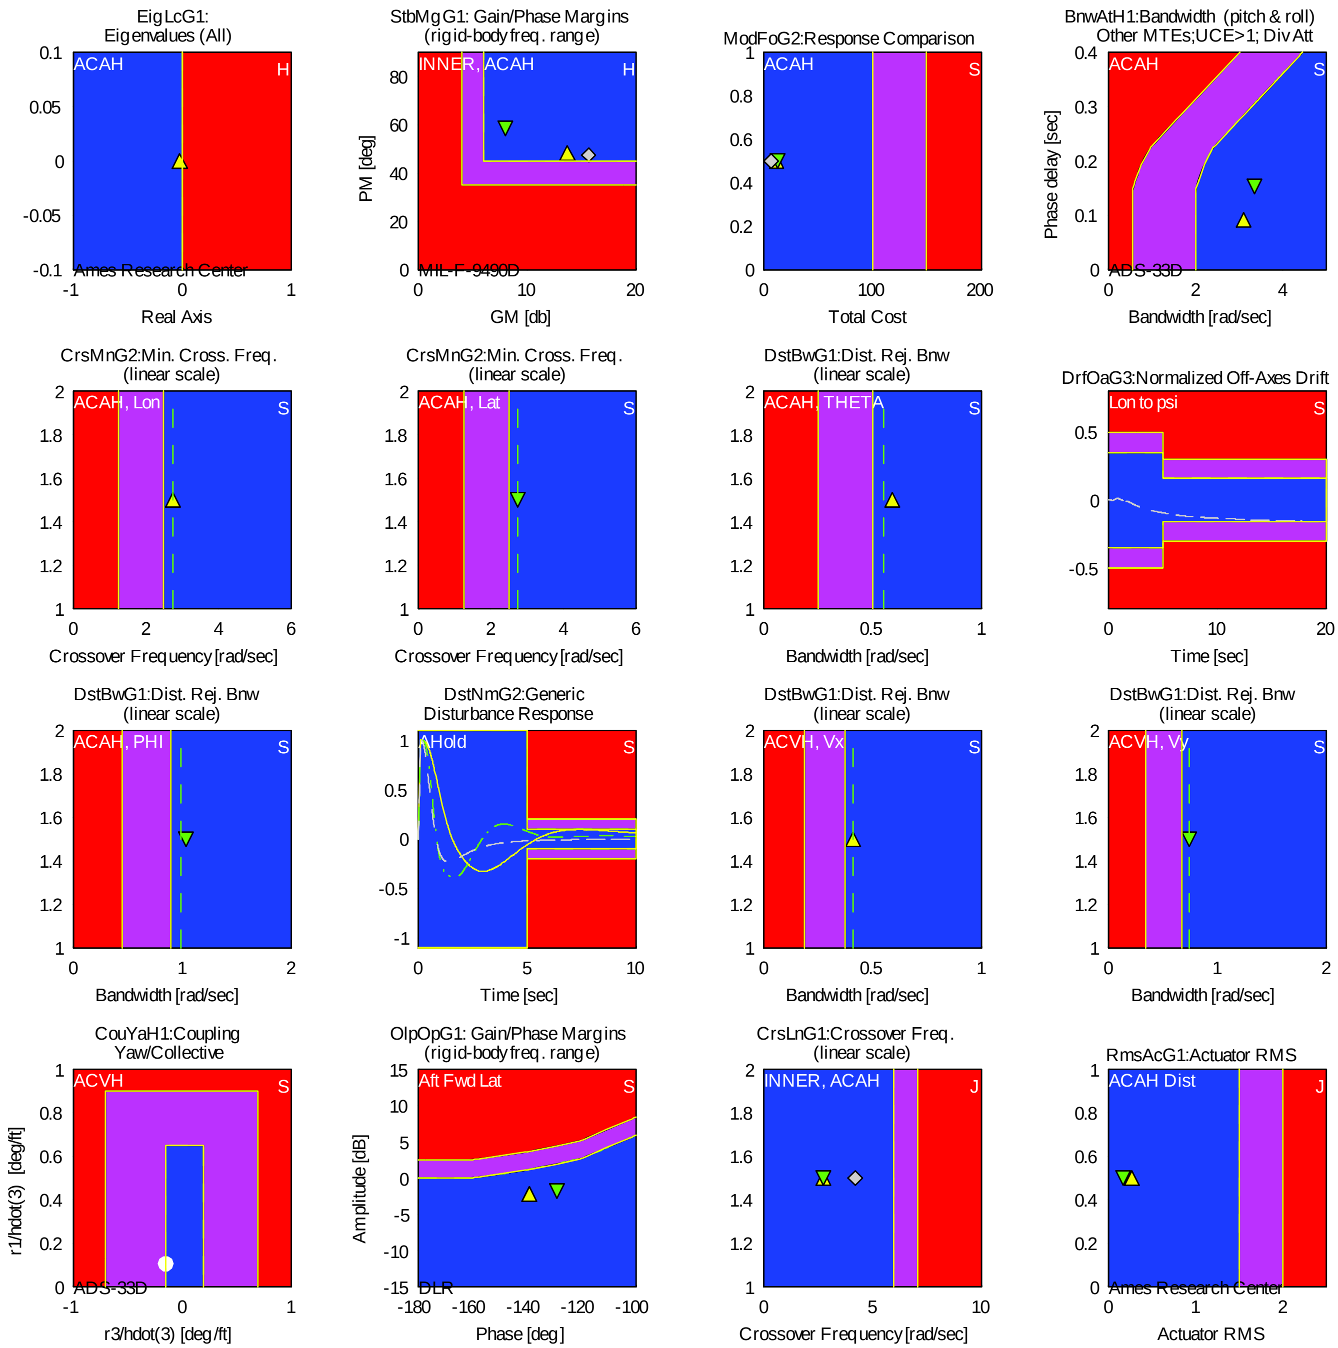

Figure 10 – CONDUIT® HQ Window for 0.2 Design Margin case of standard margins family of results

a longitudinal input is successfully controlled by the normalized off-axes drift spec and the yaw to collective response is in Level 1. Finally, the OLOP spec indicates that the system has low actuator rate saturation and no PIO tendencies. The details of the variations of the key metrics for each axis as a function of the variations in the Design Margin are discussed next. Note that the two cases eventually flight tested, as discussed later, are marked with filled symbols in Figs. 11–19.

The variations of the longitudinal phase margin, crossover frequency, and pitch attitude DRB with Design Margin are shown in Fig. 11, while the variations of longitudinal ground speed (Vx) and longitudinal position (X) DRBs are shown in Fig. 12. Note that stability margins were evaluated for all the axes and all the various modes (ACAH, VH, PH) at the inner and, if applicable, outer loops (Fig. 13). The plots of crossover frequency and phase margin in Fig. 11 are for ACAH and are shown as examples. Corresponding plots for other modes have not been shown for brevity. As may be seen, the crossover frequencies have optimized to the Design-Margin-augmented boundary for every point (Fig. 11a). Since crossover frequency is a performance objective, which is minimized to lower the cost of feedback and prevent over-design, this indicates that the optimization had been successful for all points.

As may also be seen from Fig. 11b, the pitch attitude DRB initially starts above the Design-Margin-augmented boundary. This is likely because a higher pitch attitude gain was needed to support the DRB requirements of the outer velocity and position hold loops (Fig. 12). As Fig. 12 shows, the longitudinal ground velocity and longitudinal position DRB values appear on the corresponding Design-Margin-augmented DRB boundaries for every optimization point. This indicates that the outer loop may to some extent have been a driver of the inner loop gains. In the current effort all design parameters (inner and outer loops) were optimized at the same time, therefore making it possible for the outer loop requirements to drive the inner loop gains as necessary. However, this "one-shot" optimization approach makes the problem larger and more difficult to process. In general, a "nested optimization" approach, where inner loops are optimized first and corresponding gains frozen before moving to the next outer loop, may be the more efficient design approach. However, the latter approach would eliminate the ability of the outer loop requirements to drive the inner loop gains, as was the case in the current work. As such, the "nested optimization" approach may not result in the true optimized solution that delivers the best system performance possible. A judgment of which approach to take has to be made based on available time and computational resources.

Clearly, as crossover and DRB increase the phase margin is reduced (Fig. 11c). At the same time, Fig. 14 shows that both OLOP and actuator RMS values monotonically increase. This is an indication of the increasing cost of feedback. In fact, a look at the OLOP results plotted on the OLOP spec itself, as shown in Fig. 15, clearly shows that not only the OLOP values are increasing with Design Margin, but they are approaching the Level 1 / Level 2 boundary at a 0.6 Design Margin, above which the optimization fails to complete. Also note that the rates of

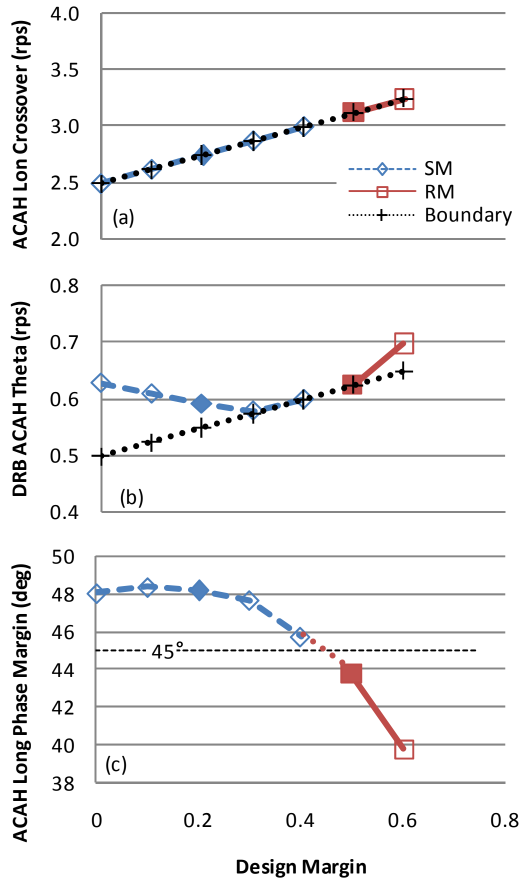

Figure 11 – Longitudinal inner loop results (flight cases filled)

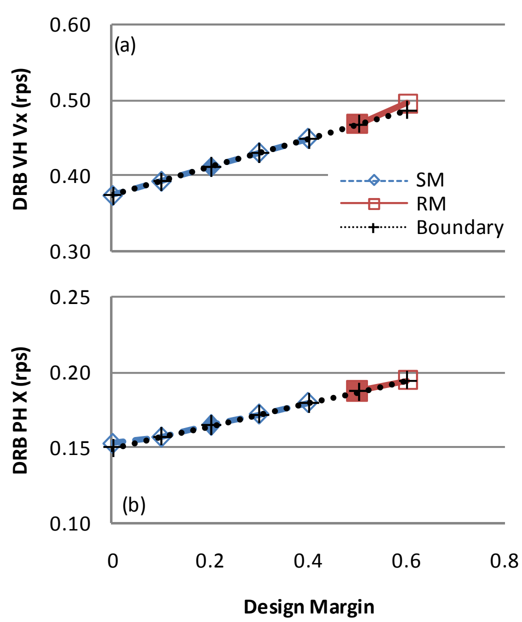

Figure 12 – Longitudinal outer loop results (flight cases filled)

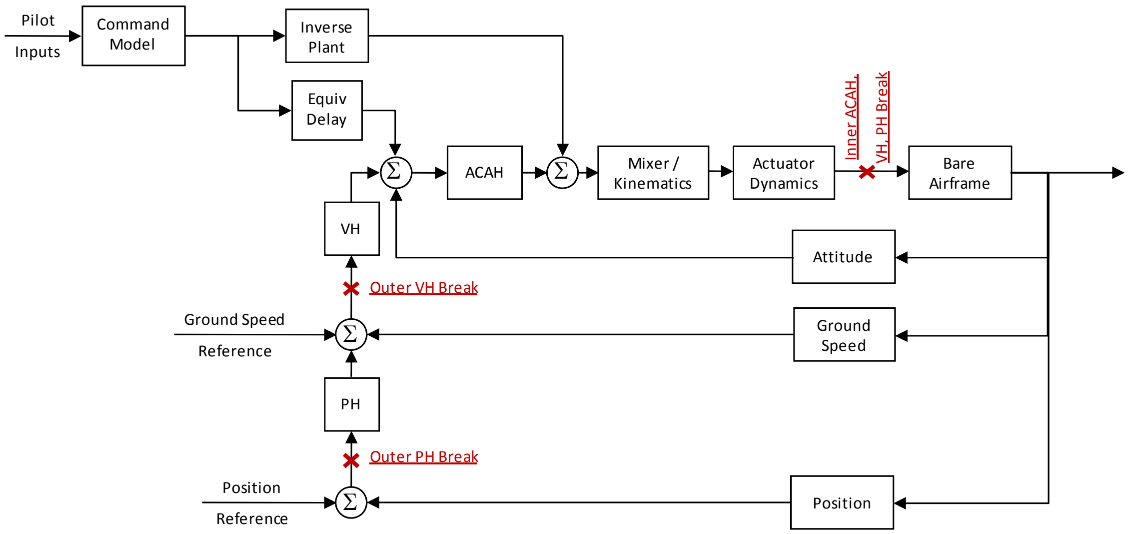

Figure 13 – Inner and outer loop break locations for stability margin calculations

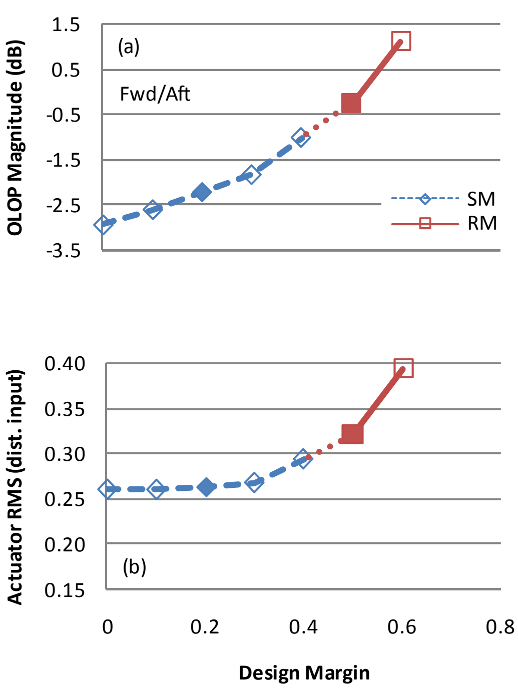

Figure 14 – Longitudinal OLOP and actuator RMS (for disturbance inputs) results (flight cases filled)

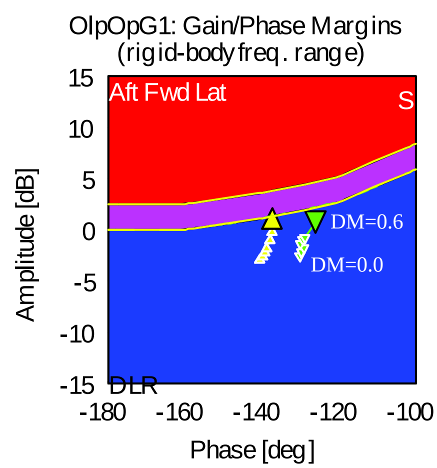

Figure 15 – OLOP results on OLOP spec

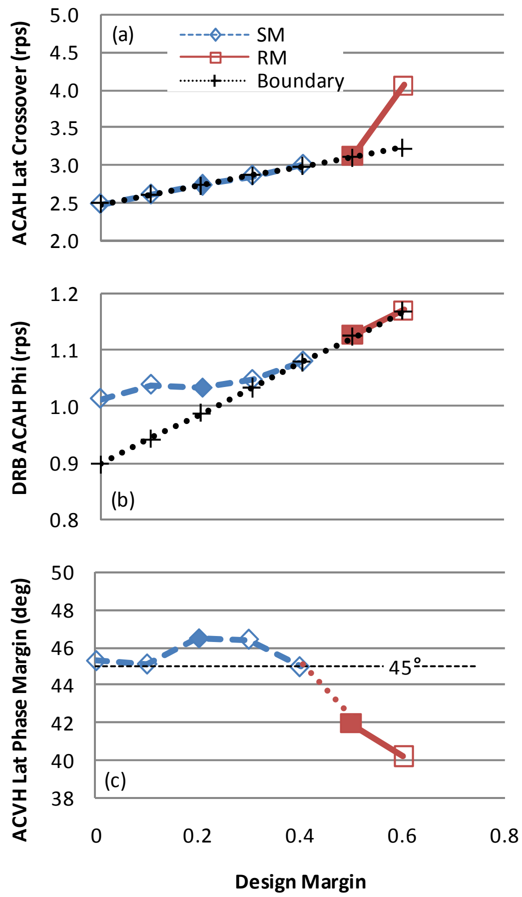

Figure 16 – Lateral inner loop results (flight cases filled)

increase in OLOP and RMS are higher for higher Design Margin values, indicating increasing cost of higher disturbance rejection performance.

The variations of the lateral phase margin, crossover frequency, and roll attitude DRB with Design Margin are shown in Fig. 16 while the variations of lateral ground speed (Vy) and lateral position (Y) are shown in Fig. 17. As may be seen, the crossover frequencies (Fig. 16a) have again optimized to the Design-Margin-augmented boundary for almost every point, indicating successfully completed optimizations. The only exception is the 0.6 Design Margin point which shows an optimized crossover frequency significantly above that required by the boundary.

Figure 16b shows that the roll attitude DRB, as in pitch, initially starts above the Design-Margin-augmented boundary, this time because higher roll attitude gains are needed to maintain stability margins. As may be seen from Fig. 16c, for Design Margins of 0.0 and 0.1 the phase margin is almost on the 45 deg boundary (SM case) and then initially moves higher as the crossover frequency

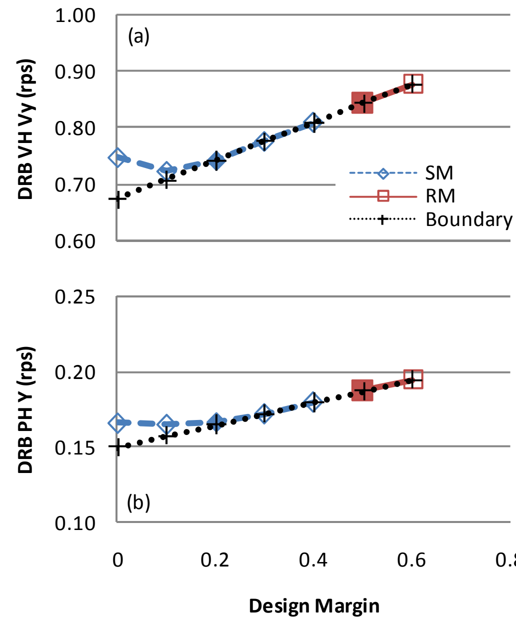

Figure 17 – Lateral outer loop results (flight cases filled)

increases, before returning to the expected trade-off trend of dropping with increasing DRB. Note from Fig. 17 that the outer loop DRBs are initially away from the boundaries so unlike in pitch, the outer loop requirements in roll are not driving the inner loop gains, at least not for the first two Design Margin points.

As in pitch, Fig. 18 shows that both OLOP and actuator RMS values monotonically increase for roll. This, again, is an indication of the increasing cost of feedback. Also, again the rates of increase in OLOP and RMS are higher for higher Design Margin values, indicating increasing cost of higher disturbance rejection performance. Moreover, note the significant jump between Design Margins 0.5 and 0.6, which is consistent with the 0.6 Design Margin case being the highest limit of performance achieved before further increase would violate one or more of the specifications.

The variations of the directional phase margin, crossover frequency, and yaw attitude DRB with Design Margin are shown in Fig. 19. Unlike longitudinal and lateral, the directional crossover frequency (Fig. 19a) starts and stays above the Design-Margin-augmented boundary until a Design Margin of 0.4. The yaw attitude DRB (Fig. 19b) also starts significantly above its corresponding boundary and remains flat until the directional crossover matches the increasing crossover requirement. The high yaw attitude gain that is behind the high crossover and DRB values appears to be driven by the yaw-to-collective coupling and the yaw drift to longitudinal input,

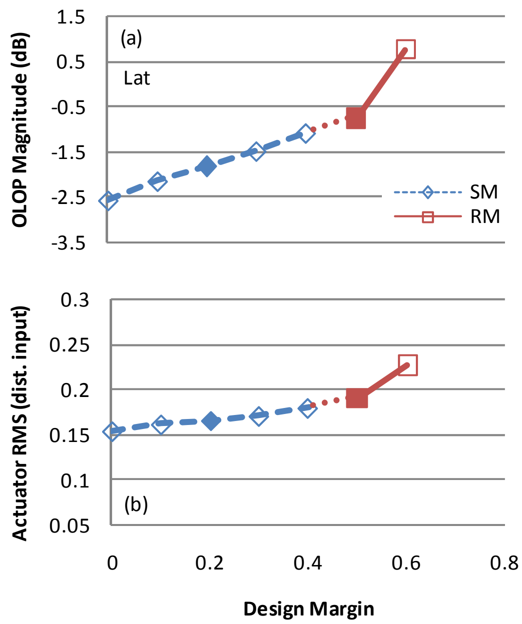

Figure 18 – Lateral OLOP and actuator RMS (for disturbance inputs) results (flight cases filled)

requirements that are otherwise unrelated to the ongoing trade-off between stability margins and disturbance rejection performance. As a result, the expected variation of stability margins (Fig. 19c) with increasing DRB can not be discerned here. Higher initial values of yaw crossover frequency and yaw attitude DRB would possibly have allowed a more clear trade-off result in the yaw axis. However, increasing the initial requirements would probably have resulted in the optimization not reaching the higher Design Margin values achieved here.

Selection of Flight Test Configurations

From each of the resulting two families of designs (SM and RM) one design was selected for flight testing. Looking at the designs in the RM set it was noted that even though the 0.6 Design Margin case had completed, there seemed to be a clear knee in the various curves at the 0.5 Design Margin case, accompanied by a sharp increase in some of the crossover frequencies and DRBs. Also, there was an increase in the slope of the OLOP-magnitude and actuator RMS results as Design Margin increased from 0.5 to 0.6. It was therefore decided to use the 0.5 Design Margin case as the selection from the RM family of results. Then, to allow for reasonable separation between the two cases, the 0.2 design margin case was selected as the candidate system from the SM family of results.

Given the model following architecture of the control laws and the fact that both gain sets were optimized to

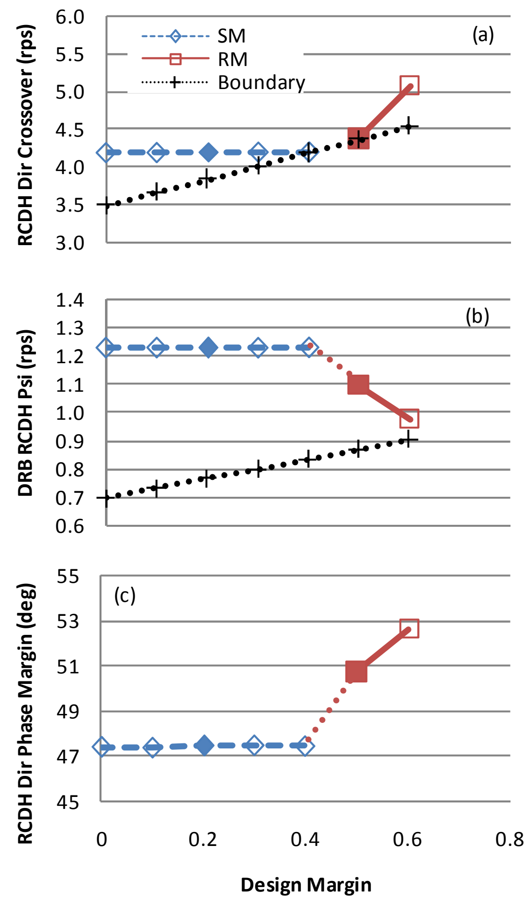

Figure 19 – Directional axis results (flight cases filled)

ensure good model following performance, it would be expected that the system response to pilot stick inputs would be similar with the two gain sets. Figure 20 compares the roll attitude responses of the system to a step cyclic input in roll with the two gain sets. As may be seen the two responses are very similar. The only difference is that with the higher Design Margin case (DM = 0.5) a slightly tighter command following is achieved, as would be expected from the higher crossover frequencies of this gain set. In contrast, Fig. 21 compares the roll attitude responses of the system to a unit roll attitude disturbance and shows that, unlike the responses to pilot inputs, the disturbance responses of the two gain sets are different. As may be seen, the disturbance response for the higher DRB/lower stability margins case (DM = 0.5) is noticeably faster and has a larger peak overshoot than the lower DRB/higher stability margins case (DM = 0.2). Both gain sets, however, were shown to have disturbance response damping ratios of better then 0.45, based on log decrement calculations for responses to pulse-type inputs, thus satisfying the ADS-33 requirements as well as the guidance of reference 15.

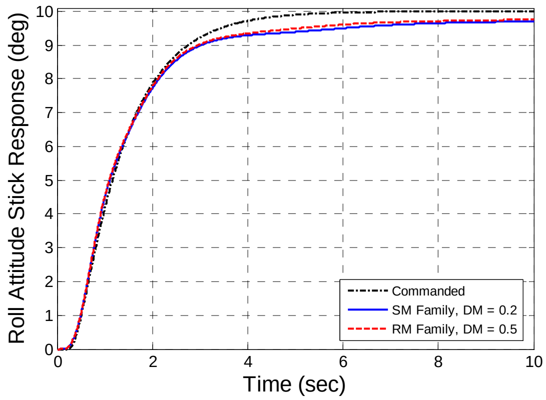

Figure 20 – Roll attitude response to step input at stick

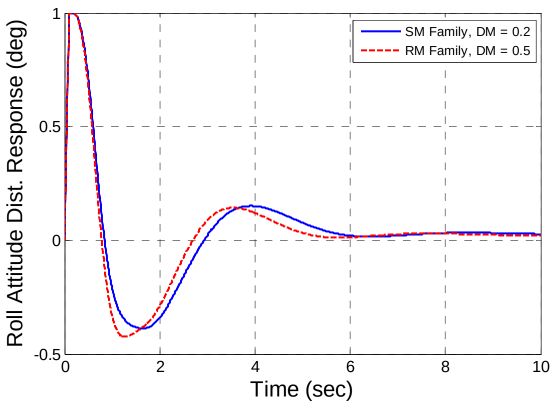

Figure 21 – Roll attitude response to unit disturbance

Flight Test Results

Five ADS-33 hover and low speed MTE's were flown by two pilots and performance measures, HQRs (using the Cooper-Harper handling qualities rating scale [14]), and pilot comments were collected. The MTE's selected were 1) Hover, 2) Hover Turn, 3) Lateral Reposition, 4) Depart/Abort, and 5) Vertical Maneuver. The Lateral Reposition and Depart/Abort are not strictly hover and low speed MTEs but they were selected based on the fact that most of the interesting dynamic portions of the maneuvers happen at low speed and near hover. Additionally, though the heave axis was not considered or evaluated as part of the trade-off study, the vertical maneuver was selected to determine if there would be noticeable differences between the two designs in terms of off-axes drift and hold functionality.

The two selected designs, along with a baseline set from earlier work aboard RASCAL, were tested and compared by flying each MTE with all three gain sets, changing the gain sets in flight between the maneuvers. The pilots were kept unaware, as much as possible, of which gain set was being flown to avoid the buildup of any bias. Also, the ordering of the gain sets was changed from MTE to MTE, and the pilots were given the option to fly the first gain set again after completing the testing of the third gain set, to avoid possible bias due to learning.

As an example of the flight test data obtained, Fig. 22 compares two Lateral Reposition results. Both cases shown were flown by the same pilot, one with the 0.2 Design Margin (SM) gain set and the other with the 0.5 Design Margin (RM) gain set. The figure shows that the pilot was able to get more aggressive with the SM gain set, achieve slightly higher lateral speed, and complete the maneuver around 2 seconds faster than with the RM case, though both cases were completed within the 18 seconds desired time for the maneuver. The figure also shows that the ground track, altitude, and heading errors are slightly less with the SM gain set. Note, however, that the excursions into the "adequate" region for ground track and altitude seen with the RM case seem to be more a result of initial starting point for the specific run than the pilot's inability to maintain "desired" ground track and altitude.

Table 5 shows a compilation of the HQRs. As may be seen, the 0.2 Design Margin case (standard margins family) was rated higher than (3 MTEs), or equal to (1 MTE), the 0.5 Design Margin case (relaxed margins family) for all but the Hover Turn MTE. One important distinction between the ratings should be noted here. In the case of the Hover, Lateral Reposition, and Depart/Abort maneuvers, the 0.2 Design Margin case was rated, on average, HQR = 3 (HQR = 3.25 for lateral reposition) while the 0.5 Design Margin case was rated, on average, HQR = 4. At first glance the difference of one HQR may not seem very significant but the difference is large owing to the fact that, as discussed by Cooper and Harper [14], HQR = 3 is Level 1, or "satisfactory without improvement," while HQR = 4 is Level 2, or "deficiencies warrant improvement." In assigning HQRs pilots are very cognizant of the boundary between Level 1 and Level 2 (HQR = 3.5) and do not cross this boundary unless they feel the system is very well behaved.

Performance data and pilot comments gathered were generally consistent with the HQRs and showed a preference for the 0.2 Design Margin gain set (standard margins). Pilot comments generally indicated that the 0.2 Design Margin case felt "more comfortable and natural" and made it "easier to get the response you want from the aircraft and freely drive aggressiveness." Even in the case of the Hover Turn MTE for which the 0.2 Design Margin case was on average assigned a worse rating than the 0.5 Design Margin case, pilot comments indicated that the system was "more predictable" but that there seemed to be some "increase in drift." In contrast, pilot comments for

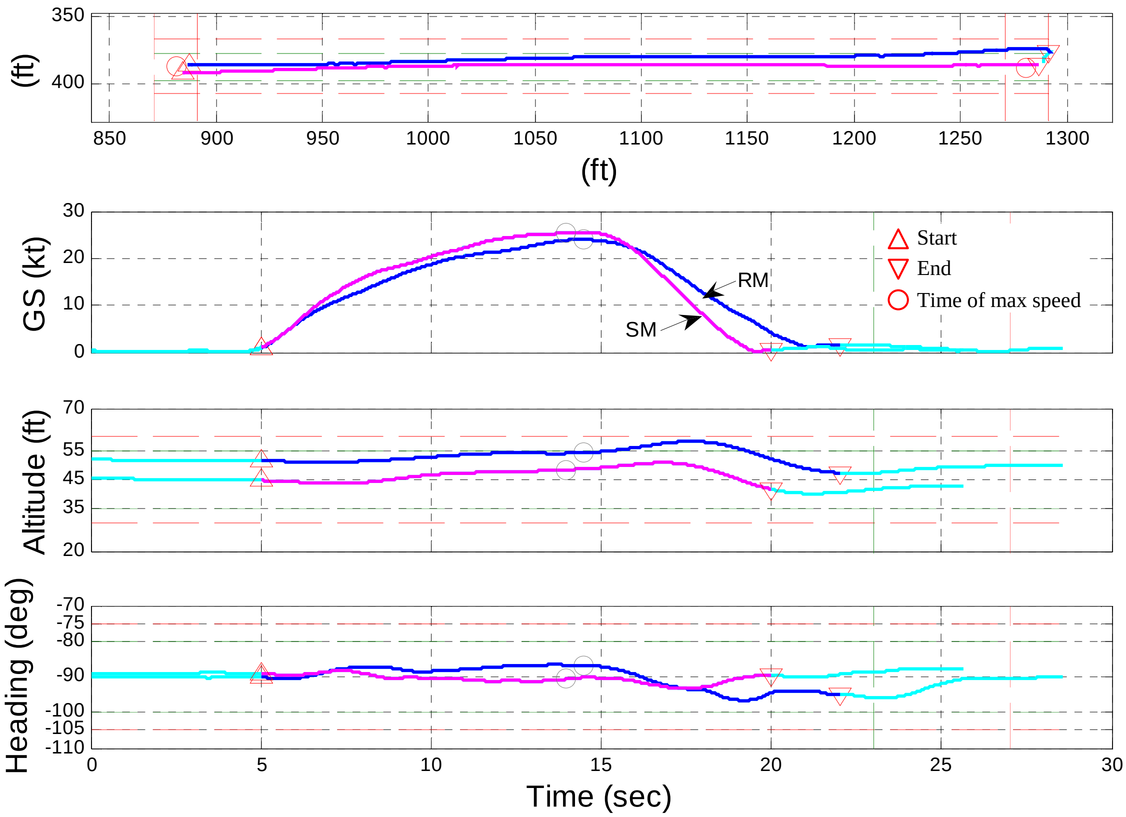

Figure 22 – Sample Lateral Reposition results for both cases

Table 5 – Compilation of Handling Qualities Ratings

	Average HQR	
	SM Case DM=0.2	RM Case DM=0.5
Precision Hover	3	4
Hover Turn	3	2.5
Lateral Reposition	3.25	4
Depart/Abort	3	4
Vertical Maneuver	3	3

the 0.5 Design Margin case generally indicated a "tendency to over control" and that with this gain set the aircraft "felt more nervous." Finally, the result for the Vertical Maneuver MTE did not show any noticeable difference between the two gain sets.

Summary

A systematic look at the effect of relaxing the flight control stability margins requirements of the system in order to achieve increased disturbance rejection performance was carried from analysis to flight on AFDD's RASCAL JUH-60A Black Hawk helicopter.

- Two families of results, one with standard stability margins (6 dB, 45 deg) and reduced disturbance rejection bandwidth (DRB) and one with relaxed stability margins (4 dB, 35 deg) and increased DRB, were generated using Design Margin Optimization in CONDUIT®.
- From these two families, one design from each was selected and flight tested in parallel using ADS-33 hover and low speed MTEs.
- Handling qualities ratings and pilot comments indicated a preference for the standard stability margins design based on a more natural feel of the system and the ability to easily get the desired response from the aircraft.

References

[1] Anon., "Flight Control Systems - Design, Installation, and Test of Piloted Aircraft, General Specification for," U.S. Air Force, MIL-F-9490D, 6 June 1975.

[2] Anon., "Aerospace - Flight Control Systems - Design, Installation and Test of Piloted Military Aircraft, General Specification for," SAE Aerospace Standard, AS94900, July 2007.

[3] Blanken, C. L., Hoh, R. H., Mitchell, D. G., Key, D. L., "Test Guide for ADS-33E-PRF," U.S. Army RDECOM Special Report AMR-AF-08-07, July 2008.

[4] Harding, J. W., Mansur, M. H., Tischler, M. B., Moody, S. J., McCann, R. C., "Optimization and Piloted Simulation Results of the AH-64D Modern Control Laws," presented at American Helicopter Society 63rd Annual Forum, Virginia Beach, VA, May 1-3, 2007.

[5] Downs, J., Prentice, R., Dalzell, S., Busachio, A., Ivler, C. M., Tischler, M. B., Mansur, M. H., "Control System Development and Flight Test Experience with the MQ-8B Fire Scout Vertical Take-Off Unmanned Aerial Vehicle (VTUAV)," presented at American Helicopter Society 63rd Annual Forum, Virginia Beach, VA, May 1-3, 2007.

[6] Christensen, K. T., Campbell, K. G., Griffith, C. D., Ivler, C. M., Tischler, M. B., Harding, J. W., "Flight Control Development for the ARH-70 Armed Reconnaissance Helicopter Program," presented at American Helicopter Society 63rd Annual Forum, Virginia Beach, VA, May 1-3, 2007.

[7] Einthoven, P., et., al., "Development of Control Laws for the Chinook Digital AFCS Program," presented at American Helcopter Society 62nd Annual Forum, Phoenix, AZ, May 2006.

[8] Moralez III, E., et., al., "Flight Research Qualilfication of the Army/NASA RASACAL Variable Stability Helicopter," presneted at American Helicopter Society 58th Annual Forum, Montreal, Canada, June 11-13, 2002.

[9] Tischler, M. B., Colbourne, J. D., Morel, M. R., Biezad, D. J., Cheung, K. K., Levine, W. S., Veronica, M., "A Multidisciplinary Flight Control Development Environment and Its Application to a Helicopter," IEEE Control System Magazine, Vol. 19, No. 4, pp 22-33, August 1999.

[10] Tischler, M. B., Remple, R. K., "Aircraft and Rotorcraft System Identification: Engineering Methods and Flight Test Examples," AIAA, August 2006.

[11] Anon., "Handling Qualities Requirements for Military Rotorcraft," Aeronautical Design Standard-33 (ADS-33E-PRF), U.S. Army Aviation and Missile Command, March 21, 2000.

[12] Tischler, M. B., Ivler, C. M., Mansur, M. H., Cheung, K. K., Berger, T., Berrios, M., "Handling Qualities Optimization and Trade-offs in Rotorcraft Flight Control Design," presented at AHS Specialists Meeting on Rotorcraft Handling-Qualities, Liverpool, U. K. 4-6 November 2008.

[13] Duda, H., "Prediction of Pilot-in-the-Loop Oscillations due to Rate Saturation," AIAA Journal of Guidance, Navigation, and Control, Vol. 20, No. 3, May-June 1997.

[14] Cooper, G. E., Harper, R. P., "The Use of Pilot Rating in the Evaluation of Aircraft Handing Qualities," NASA TN D-5153, April 1969.

[15] Blanken, C. L., et. al., "An Investigation of Rotorcraft Stability-Phase Margin Requirements in Hover," presented at American Helicopter Society 65th Annual Forum, Grapevine, TX, May 27-29, 2009.