

# Verification and Tuning of an Adaptive Controller for an Unmanned Air Vehicle\*

Luis G. Crespo<sup>†</sup>

*National Institute of Aerospace, Hampton, VA, 23666, USA*

Megumi Matsutani<sup>‡</sup>, and Anuradha M. Annaswamy<sup>§</sup>

*Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*

**This paper focuses on the analysis and tuning of a controller based on the “Adaptive Control Technology for Safe Flight (ACTS)” architecture. The ACTS architecture consists of a nominal, non-adaptive controller that provides satisfactory performance under nominal flying conditions, and an adaptive controller that provides robustness under off-nominal ones. A framework unifying control verification and gain tuning is used to make the controller’s ability to satisfy the closed-loop requirements more robust to uncertainty. In this paper we tune the gains of both controllers using this approach. Some advantages and drawbacks of adaptation are identified by performing a global robustness assessment of both the adaptive controller and its non-adaptive counterpart. The analyses used to determine these characteristics are based on evaluating the degradation in closed-loop performance resulting from uncertainties having increasing levels of severity. The specific adverse conditions considered can be grouped into three categories: aerodynamic uncertainties, structural damage, and actuator failures. These failures include partial and total loss of control effectiveness, locked-in-place control surface deflections, and engine out conditions. The requirements considered are the peak structural loading, the ability of the controller to track pilot commands, the ability of the controller to keep the aircraft’s state within the reliable flight envelope, and the handling/riding qualities of the aircraft. The nominal controller resulting from these tuning strategies was successfully validated using the NASA GTM Flight Test Vehicle.**

**Keywords:** Adaptive control, control tuning, control verification, safety, UAV.

## I. Introduction

An adaptive reconfigurable controller autonomously changes the controller gains to maintain satisfactory performance when unforeseen changes in the system dynamics occur. Adaptive control has the potential to improve flight safety, as the loss-of-control is one of the major causes of abnormal flight and fatal accidents. Over the past three decades, adaptive control has been developed extensively and its main performance and robustness properties have been established [2, 8–11].

In this paper we tune a controller designed for the Generic Transport Model (GTM) and evaluate its robustness characteristics. The GTM is a dynamically scaled test article of a transport aircraft for which NASA Langley Research Center has developed a high-fidelity mathematical model. Since the dynamics of this model, as those of the vehicle itself, depart considerably from the Linear Time Invariant (LTI) setting supporting the theory, it remains to be determined if the improvements in stability and performance predicted in such setting are realized in practice. This paper addresses this question by performing a comparative analysis of the robustness of the non-adaptive controller at the core of the ACTS architecture and of its adaptive augmentation. Details of these controllers structure and their implementation are available in the companion paper [6]. The studies performed herein, which are based on the developments in [4, 7], evaluate the system robustness to multiple adverse conditions solely based on the performance observed in the high-fidelity simulation. As expected, many of the assumptions and simplifications supporting the control design procedures do not hold in this setting.

---

\*This work was supported by the NRA NNX08AC62A of the IRAC project of NASA.

<sup>†</sup>Senior Research Scientist, 100 Exploration Way, AIAA Member.

<sup>‡</sup>Graduate Student, Mechanical Engineering, 77 Massachusetts Avenue, Room 3-441.

<sup>§</sup>Senior Research Scientist, Mechanical Engineering, 77 Massachusetts Avenue, Room 3-339A.

The specific adverse conditions considered can be grouped into three categories: aerodynamic uncertainties (e.g., changes in pitch stiffness, roll and yaw damping), structural damage (e.g., situations where the CG moves from its nominal location), and actuator failures (e.g., situations where symmetric and asymmetric failures in control surfaces and engines occur). These failures include partial and total loss of control effectiveness, locked-in-place control surface deflections, and engine out conditions. The requirements considered are (i) the peak structural loading, (ii) the ability of the controller to track commands, (iii) whether the aircraft's state stays within the reliable flight envelope (i.e., region in the state space where the aircraft dynamics are properly modelled and flying is safe), (iv) the handling/riding qualities, and (v) the ability of the controller to track the dynamics of the reference model.

This paper is organized as follows. Section II briefly introduces the framework and set up used for the verification of a controller's robustness. This is followed in Section III by robustness analysis of a baseline nominal controller designed using classical control concepts. The control tuning framework proposed is presented in Section IV. The sections that follow demonstrate this framework by tuning both the nominal and adaptive controllers. Finally, a few concluding remarks are made.

## II. Control Verification

The terms verification and validation have not been defined or used consistently by the scientific community in general. The notions used in this paper are in agreement with the NASA technical standards on the subject [1]. In particular, *Verification* is defined as “the process of determining that a computational model accurately represents the underlying mathematical model and its solution from the perspective of the intended uses of modeling and simulation”. Therefore, a verified controller is one that satisfies the requirements set for the closed-loop system in a mathematical setting. On the other hand, *Validation* is “the process of determining the degree to which a model or a simulation is an accurate representation of the real world from the perspective of the intended uses of the model or simulation”. Therefore, a validated controller is one that satisfies the requirements set for the closed-loop system experimentally. Note that the emphasis of verification is on the mathematics and the emphasis of validation is on the physics.

In this paper, verification refers to evaluating the robustness of a controller to various uncertainties and failures according to its ability to satisfy the closed-loop requirements. A brief introduction to the methodology used for this is presented next [4].

### A. Framework

Consider a plant whose dynamics depend on the uncertain parameter vector  $\hat{p}$ . Assume that controller  $c$  with gains  $d$  is applied to this plant. The resulting closed-loop system is deemed acceptable if a set of requirements are satisfied. These requirements can be described as the set of inequality constraints  $g < 0$ . The satisfaction of the constraints implies the satisfaction of the requirements. The overall objective of this analysis is to determine the ranges of uncertainty in  $\hat{p}$  in the set  $\hat{p}_{min} \leq \hat{p} \leq \hat{p}_{max}$  for which the controller  $c$  with gains  $d$  satisfies  $g(\hat{p}, c(d)) < 0$ . A brief introduction to the notions used to make this analysis is presented next.

Assume that the controller  $c$  satisfies all requirements at the nominal parameter point  $\bar{p}$ . The nominal parameter point is the realization of the uncertainty corresponding to nominal flying conditions. Call the region in the parameter space where at least one of the requirements is violated the *Failure Domain*. The complement set of this region, called the *Non-failure Domain*, is composed of all parameter realizations where all the requirements are satisfied. Let  $\Omega$  be a set in  $\hat{p}$ -space, called the *Reference Set*, whose geometric center is  $\bar{p}$ . The task of interest is to assign a measure of robustness to controller  $c$  based on measuring how much the reference set can be deformed before intersecting failure domain.

The *Homothetic Deformation* of  $\Omega$  with respect to  $\bar{p}$  by a factor of  $\sigma \geq 0$ , is the set of  $\hat{p}$  points satisfying  $\hat{p} = \bar{p} + \sigma(\hat{p} - \bar{p})$  such that  $\hat{p} \in \Omega$ . The factor  $\sigma$ , is called the *Similitude Ratio*. While expansions of  $\Omega$  are accomplished when  $\sigma > 1$ , contractions result when  $0 \leq \sigma < 1$ . Intuitively, one imagines that the reference set is being deformed until its boundary touches the failure domain. Any point where the deforming set touches the failure domain is a *Critical Parameter Value* (CPV). The CPV might not be unique. The deformed set, called the *Maximal Set*, is fully contained in the non-failure domain by design.

The *Critical Similitude Ratio* (CSR), denoted as  $\tilde{\sigma}$ , is the similitude ratio of the deformation leading to the maximal set. While the CSR is a non-dimensional number, the Parametric Safety Margin (PSM), denoted as  $\rho$  is its dimensional equivalent. The values taken on by the CSR and the PSM are proportional to the size of the maximal set and measure the separation between the nominal parameter point  $\bar{p}$  and the failure domain. If the PSM takes on a negative value, the controller does not even satisfy the requirements for the nominal parameter point  $\bar{p}$ . If the PSM is zero, the controller

exhibits no robustness because there are arbitrarily small perturbations of  $\hat{p}$  from  $\bar{p}$  leading to a requirement violation. If the PSM is positive, the requirements are satisfied at  $\bar{p}$  and its vicinity. The larger the PSM, the larger the  $\Omega$ -shaped neighborhood of  $\bar{p}$  for which the requirements are satisfied. The calculation of the CPV requires solving a standard optimization problem. Such a problem is non-convex when the dependency of  $g$  on  $\hat{p}$  is nonlinear.

The reference set  $\Omega$  used in this paper is the hyper-rectangle

$$\mathcal{R}(\bar{p}, m) = \{\hat{p} : \bar{p} - m \leq \hat{p} \leq \bar{p} + m\}, \quad (1)$$

where  $m > 0$  is the vector of half-lengths of the sides. In this case, the CPV for the  $j$ th requirement is given by

$$\langle \tilde{p}^j, \tilde{\sigma}^j \rangle = \underset{\hat{p}, \sigma}{\operatorname{argmin}} \{ \sigma : g_j(\hat{p}, c) \geq 0, \hat{p}_{min} \leq \hat{p} \leq \hat{p}_{max}, \bar{p} - \sigma m \leq \hat{p} \leq \bar{p} + \sigma m \}. \quad (2)$$

The overall CPV  $\tilde{p}$  is the  $\tilde{p}^j$  that is the closest to the nominal parameter point  $\bar{p}$  according to the norm  $\|x\|_m^\infty = \max_i \{|x_i|/m_i\}$ . The overall CSR  $\tilde{\sigma}$  is the CSR associated with the overall CPV. Alternatively, the overall CPV and CSR can be found from

$$\langle \tilde{p}, \tilde{\sigma} \rangle = \underset{\hat{p}, \sigma}{\operatorname{argmin}} \{ \sigma : \max\{g(\hat{p}, c)\} \geq 0, \hat{p}_{min} \leq \hat{p} \leq \hat{p}_{max}, \bar{p} - \sigma m \leq \hat{p} \leq \bar{p} + \sigma m \}. \quad (3)$$

The PSM is defined as

$$\rho(c) = \tilde{\sigma} \|m\| \quad (4)$$

In the one-dimensional setting  $\hat{p} \in \mathbb{R}$ , the maximal set is the largest range of  $\hat{p}$  centered at  $\bar{p}$  for which  $g(\hat{p}, c) < 0$  for all set members. The PSM corresponding to controller  $c$  for the  $j$ th requirement in  $g$  is

$$\rho^j(c) = -\operatorname{sign}\{g_j(\bar{p}, c)\} |\tilde{p} - \bar{p}| \quad (5)$$

The corresponding CPV is given by

$$\tilde{p}^j = \min_{\hat{p}} \{ |\bar{p} - \hat{p}| : g_j(\bar{p} + (\hat{p} - \bar{p}), c) = 0, \hat{p}_{min} \leq \hat{p} \leq \hat{p}_{max} \} \quad (6)$$

## B. Application to the GTM

The components of the uncertain parameter vector  $\hat{p}$  considered in the analyses that follow are described next. The parameters  $\Lambda_{THL}$ ,  $\Lambda_{THR}$ ,  $\Lambda_{ERO}$ ,  $\Lambda_{ERI}$ ,  $\Lambda_{ELO}$ ,  $\Lambda_{ELI}$ ,  $\Lambda_{AL}$ ,  $\Lambda_{AR}$ ,  $\Lambda_{RU}$  and  $\Lambda_{RD}$  are the control effectiveness of the main ten inputs [6]. These are the throttle input to the left and right engines, the four elevators, the left and right ailerons; and the upper and lower rudders. These parameters, which are used to describe both loss in control effectiveness and the locked-in-place failures, should be interpreted as follows: if  $0 < \Lambda \leq 1$  the actuator will suffer a loss in control effectiveness, making the actual plant's input  $\Lambda R_s(u + U_0)$  instead of  $R_s(u + U_0)$  [6]. If  $-1 \leq \Lambda \leq 0$ , the actuator will be locked at the fixed value  $u = \Lambda u_{max}$ .

$\tau > 0$  is an unknown time delay,  $\mu > 0$  is a scaling factor proportional to the amplitude of the commands<sup>a</sup>,  $\Delta_x$  is a shift in the CG from its nominal location in the  $x$ -direction<sup>b</sup>, and  $\delta C_{m\alpha}$ ,  $\delta C_{lp}$  and  $\delta C_{nr}$  are aerodynamic uncertainties in pitch stiffness, roll and yaw damping. These aerodynamic uncertainties are implemented by using primary or secondary control surfaces. These surfaces are no longer driven by the flight control system, but instead by additional feedback loops aimed at artificially degrading the resulting stability derivatives. From the perspective of the flight control system this is equivalent to having different stability and control derivatives. In particular, the generation of non-zero values for  $\delta C_{m\alpha}$  and  $\delta C_{nr}$  require the loss in control effectiveness  $\Lambda_{ERI} = \Lambda_{ELI} = 0$  and  $\Lambda_{RD} = 0$ , respectively.

The flight condition corresponding to  $\hat{p} = \bar{p}$  is one in which the amplitude of the pilot commands is moderate ( $\mu = 1$ ), the actuators are fully functional (i.e., effectiveness of all actuators is one), there is no additional time delay in the processing and communication of signals, the CG remains at its nominal location and there are no aerodynamic uncertainties.

The design requirements that define a satisfactory closed-loop response are as follows:

<sup>a</sup>The desired maneuver is composed of a sequence of doublets in  $\alpha_{cmd}$ ,  $\beta_{cmd}$  and  $p_{cmd}$ . When  $\mu = 0$  the commands keep the vehicle trimmed at the point used for control design. As  $\mu$  increases the commands make the vehicle depart from this trim point. The larger the value of  $\mu$ , the larger is such a departure, and the larger are the effects of the nonlinear dynamics. Note that the vehicle is supposed to return to the trim point in steady state no matter how large  $\mu$  is. By including  $\mu$  into  $\hat{p}$  we evaluate the degradation in performance caused by inaccuracies in the LTI representation used in the underlying control theory.

<sup>b</sup>The value of  $\Delta_x$  will be specified as a percentage of the mean aerodynamic chord. Besides, its sign is consistent with the orientation of the body axes, e.g.,  $\Delta_x = -0.5$  implies an aft shift in the CG location equal to 50% of the mean aerodynamic chord.

- **Structural loading:** the requirement  $g_1 < 0$  ensures that the loading factor caused by the vehicle's acceleration does not exceed the structural stress limit.
- **Command tracking:** the requirements  $g_i < 0$  for  $i = 2, 3, 4, 5$  ensure the aircraft satisfactorily tracks the  $\alpha_{\text{cmd}}$ ,  $V_{\text{cmd}}$ ,  $\beta_{\text{cmd}}$  and  $p_{\text{cmd}}$  commands generated by the pilot.
- **Reliable Flight Envelope (RFE):** the requirement  $g_6 < 0$  ensures that the vehicle remains within a safe flight envelope, i.e., the vehicle's state stays within a given bounded set, throughout the maneuver.
- **Riding quality/high frequency oscillation:** the requirement  $g_7 < 0$  bounds the amount of energy in the high-frequency part of the power spectrum of the response. Excessively large adaptation rates induce high frequency oscillations that can lead to the violation of this requirement.
- **Reference tracking:** the requirements  $g_i < 0$  for  $i = 8, 9$  bound the offset between the plant and the reference model for the pitch rate  $q$  and the yaw rate  $r$ .

Because all tracking requirements are based on  $L_2$  norms, they are a function of the transient response. Details of the functional form of the requirements can be found in [7]. Notice that the dependency of all the components of  $g$  on the parameter  $\hat{p}$  assume an unknown and implicit functional form. Further notice that evaluating  $g$  for a particular realization of  $\hat{p}$  requires simulating the closed-loop response and calculating the performance function  $g$ . In this study the values of  $\hat{p}_{\min}$  and  $\hat{p}_{\max}$  are given by  $-1 \leq \Lambda \leq 1$  for all controllable inputs, and by  $0 \leq \tau \leq 0.07$ ,  $0 \leq \mu \leq 3$ ,  $-0.5 \leq \Delta_x \leq 1$ ,  $0 \leq \delta C_{lp} \leq 1$ ,  $0 \leq \delta C_{m\alpha} \leq 1$  and  $0 \leq \delta C_{nr} \leq 1$ . In the inequalities for the aerodynamic uncertainties the zeros correspond to conditions where there is no uncertainty. The ones, on the other hand, correspond to a severe degradation in the corresponding aerodynamic coefficient well beyond open-loop neutral stability.

### III. Robustness Analysis of the Baseline Nominal Controller

The baseline version of the nominal controller will be denoted as  $c_{\text{nom, base}}$ . The gains of this controller were chosen based on classical margins of stability and performance as well as on the pilot's feedback. After several rounds of tuning and exhaustive testing in the real time simulator, this controller was rendered suitable for flight by the pilot. Since the pilot's determination of whether a controller is suitable for flight or not mostly depends on testing under nominal flying conditions, the acceptance or rejection of any given controller does not take into consideration its robustness characteristics.

$\hat{p}$	$\rho_1$	$\rho_2$	$\rho_3$	$\rho_4$	$\rho_5$	$\rho_6$	$\rho_7$	$\rho_8$	$\rho_9$
$\Lambda_{THL}$	2.00	2.00	2.00	2.00	1.87	1.55	2.00	1.92	1.47
$\Lambda_{ELO}$	2.00	2.00	2.00	2.00	1.50	2.00	2.00	2.00	1.67
$\Lambda_{AL}$	2.00	2.00	2.00	1.81	1.30	1.72	1.75	1.87	1.24
$\Lambda_{RU}$	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	1.27
$\tau$	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07	0.07
$\mu$	2.00	2.00	2.00	2.00	2.00	1.16	2.00	2.00	2.00
$\Delta_x$	0.46	0.28	0.46	0.42	0.17	0.12	0.31	0.29	0.42
$\delta C_{lp}$	1.00	0.66	1.00	0.53	0.65	0.66	0.43	1.00	0.72
$\delta C_{m\alpha}$	1.00	0.79	1.00	1.00	1.00	1.00	0.76	0.77	1.00
$\delta C_{nr}$	1.00	0.89	0.86	0.80	1.00	1.00	0.82	0.90	0.73

Table 1. One dimensional PSMs of  $c_{\text{nom, base}}$ .

The PSMs corresponding to  $c_{\text{nom, base}}$  for all requirements and several one-dimensional uncertainties are listed in Table 1. In these analyses and those that follow, all the uncertain parameters excluded from  $\hat{p}$  assume their nominal value  $\bar{p}$ . The results are fairly intuitive. The critical requirement for  $\Lambda_{THL}$ ,  $\Lambda_{RU}$ ,  $\Lambda_{AL}$  and  $\delta C_{nr}$  uncertainty is the yaw rate tracking requirement  $g_9$ . The tracking of  $V$  on the other hand, degrades the most when the elevator or the aileron are locked. Furthermore,  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  uncertainties trigger unacceptable levels of high frequency oscillation before compromising asymptotic tracking of the commands. The metrics in Table 1 will be used to compare the robustness characteristics of  $c_{\text{nom, base}}$  with those of other controllers.

Multi-dimensional robustness analyses of the baseline nominal controller led to the data in Table 2. The second and third columns show the PSM and the failure probabilities. Suitable values of  $m$  for all the components of  $\hat{p}$  were chosen when calculating PSMs. Probabilities were calculated after assuming a joint probability density function for  $\hat{p}$  and evaluating  $g$  for 200 Monte Carlo samples. The parameter  $\eta$ , which implies the combination of aerodynamic uncertainties  $\delta C_{lp} = \eta$  and  $\delta C_{m\alpha} = 0.8\eta$ , is used to evaluate the effects of the simultaneous pitch and roll uncertainties on the aircraft's performance. The metrics above enable the comparative analysis of competing control alternatives.

$\hat{p}$	$\min(\rho(c_{\text{nom, base}}))$	$P[\max(g) > 0]$
$[\tau, \mu]$	2.1	7.5e-2
$[\delta C_{lp}, \delta C_{m\alpha}]$	0.63	3.6e-1
$[\Delta_x, \Delta_y]$	0.095	7.7e-1
$[\tau, \delta C_{lp}]$	0.39	3.3e-1
$[\tau, \delta C_{m\alpha}]$	0.52	1.2e-1
$[\tau, \delta C_{nr}]$	0.67	8.5e-2
$[\delta C_{lp}, \delta C_{nr}]$	0.59	3.7e-1
$[\Lambda_{THR}, \Lambda_{RU,RL}]$	1.08	4.5e-2
$[\tau, \eta]$	0.39	3.8e-1
$[\Lambda_{ELO}, \Lambda_{ELI}]$	1.59	1.5e-2
$[\tau, \delta C_{lp}, \delta C_{m\alpha}, \delta C_{nr}]$	0.64	5.6e-1
$[\Lambda_{THR}, \dots, \Lambda_{RD}]$	1.68	2.4e-1

**Table 2. Multi-dimensional PSMs and failure probabilities of  $c_{\text{nom, base}}$ .**

With this aim we postpone a discussion on their significance to a later section. Figures 1 and 2 show the parameter space for  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$  and  $\hat{p} = [\tau, \delta C_{lp}]$ . The coloring of the figures indicate the number of the requirements being violated. Therefore, while the non-failure domain (i.e., the set where all the requirements are satisfied) is colored in green, the failure domain is colored otherwise. These figures are generated by simulating the closed-loop response at the points of an uniformly spaced grid, calculating the corresponding performance vector  $g$ , and counting its number of non-negative components. This exceedingly large number of simulations have only been made for illustrative purposes. The maximal set, which is a subset of the non-failure domain by design, is superimposed. Note that the size of the maximal set, measured by the PSM, is an indicator of the separation between acceptable and unacceptable closed-loop performances and as such it is a measure of robustness to these uncertainties.

Figure 1 shows that  $\delta C_{lp}$  uncertainty degrades the system performance faster than  $\delta C_{m\alpha}$ . Furthermore, the almost rectangular shape of the non-failure domain indicates no interaction between these two uncertainties. This observation is in agreement with the control design procedure in which the longitudinal and lateral/directional controllers were designed independently. The CPV, shown in the figure as a circle, occurs when  $\delta C_{m\alpha} = 0$ . Note that the physics of the problem on the surfaces  $\delta C_{m\alpha} = 0$  and  $\delta C_{m\alpha} = \varepsilon$ , where  $0 < \varepsilon \ll 1$ , differ significantly since in the former case the four elevators are fully functional while in the latter case the two inner elevators are totally ineffective. Figure 2 shows how the performance function  $g$  depends on the aerodynamic uncertainty  $\eta$  and the time delay  $\tau$ . In contrast to the previous figure, the compound effect of both uncertainties affect the geometry of the non-failure domain. This dependency is almost linear in the vicinity of the maximal set.

## IV. Control Tuning

Ultimately, a good controller should satisfy the stability and performance requirements imposed upon it with sufficiently large robustness margins. The PSM and CSR introduced above are examples of such margins. A computational framework that enables searching for the set of controller gains that maximize such margins is introduced next. Further details are available in [5].

Recall that the controller  $c$  depends on the control parameters  $d$ . This dependency, which is inconsequential for the analyses above, is exploited in the developments that follow. Note that the constraint function  $g < 0$  and therefore, the non-failure domain and all the robustness metrics used to evaluate it, are a function of  $d$ . A control designer wants to determine the value of the design point  $d$  for which the closed-loop system exhibits satisfactory performance and

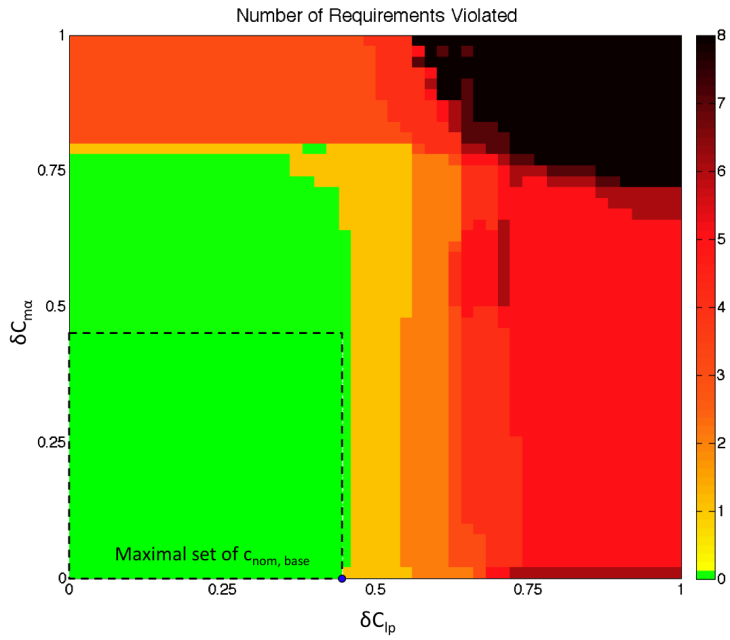


Figure 1. Maximal set, failure and non-failure domains of  $c_{nom, base}$  for  $\hat{p} = [\delta C_{ip}, \delta C_{m\alpha}]$ .

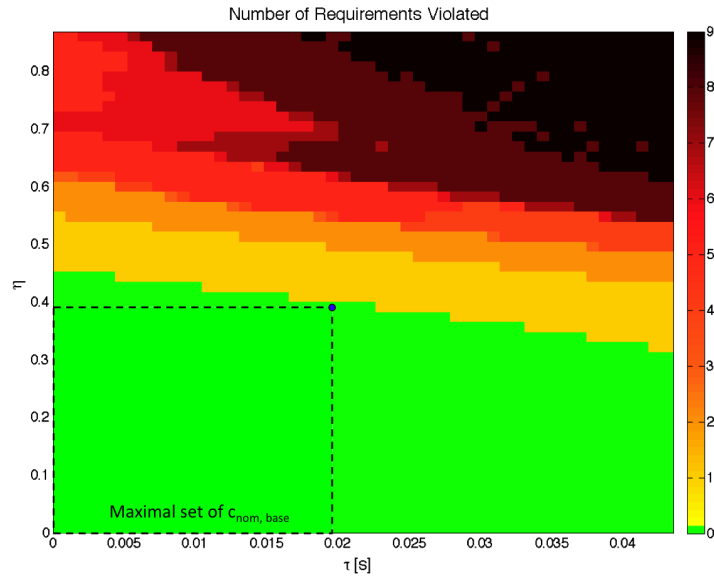


Figure 2. Maximal set, failure and non-failure domains of  $c_{nom, base}$  for  $\hat{p} = [\tau, \eta]$ .

robustness characteristics. The design points targeted herein realize the largest PSM the control structure allows for. Since these controllers are able to tolerate the largest perturbations from the nominal parameter point before violating any of the requirements, they will be called *Robustly Optimal*. The formulations that follow evaluate the robustness characteristics of any design point considered during the search for the robustly optimal one by using the exact maximal set, or by using a multi-point approximation to the maximal set. Only the former strategy always leads to the intended designs while the other strategy may not, due to the approximate nature of the formulation. However, the relaxed computational demands of the latter strategy make it attractive in spite of their potential drawbacks. Because the framework in Reference [4] allows for a detailed robustness analysis, the inaccuracies resulting from such drawbacks can be detected and mended a posteriori. To expedite the search for a robustly optimal controller, we use the design point resulting from the approximate formulation as starting point for the exact one.

### A. Enlarging the Maximal Set

The problem of interest is given by

$$d^* = \underset{d}{\operatorname{argmax}} \{ \tilde{\sigma}(d) : h(\hat{p}, d) > 0 \}, \quad (7)$$

where  $d^*$  is the set of optimal controller gains, and the CSR  $\tilde{\sigma}$  is given by Equation (3). The constraint  $h(\bar{p}, d) > 0$  is used to enforce additional characteristics, e.g.,  $h(\bar{p}, d) = \tau - \tau_{\min} > 0$  ensures that the time delay for the nominal plant is larger than  $\tau_{\min}$ . Recall that determining the CSR  $\tilde{\eta}$  corresponding to any given design point entails solving an optimization problem. Therefore, this formulation has an optimization in the outer loop and another one in the inner loop. While the outer loop searches for the robustly optimal gains  $d^*$ , the inner one searches for the CSR corresponding to the design point under consideration. Nested optimizations, such as Equation (7) and those resulting from worst-case-based design policies, impose stringent computational demands. Such demands can be substantially mitigated by using multi-point approximations. The multi-point approach in the next section eliminates the inner optimization loop at the expense of possibly finding suboptimal design points.

### B. Enlarging an Approximation to the Maximal Set

An approximation to the solution of Equation (7) is given by

$$\langle d^*, \xi^* \rangle = \underset{d, \xi}{\operatorname{argmax}} \left\{ \xi : \max_{j,i} \{ g_j(\bar{p} + \xi(p^i - \bar{p}), d) \} < 0, \xi \geq 0, h(\hat{p}, d) > 0 \right\}, \quad (8)$$

where  $p^i$  for  $i = 1, \dots, n$  are parameter-points on the surface of  $\mathcal{R}(\bar{p}, m)$ . These points, which only have to be computed once, should be uniformly distributed over the surface. Note that this formulation replaces the inner optimization loop in Equation (7) by a multi-point constraint over parameter points lying on the surface of the approximate maximal set. Further notice that  $\tilde{\sigma}(d^*)$  is approximated by  $\xi^*$ . Since the satisfaction of the multi-point constraints does not guarantee that the approximate maximal set is a subset of the non-failure domain, this formulation may lead to suboptimal designs. Besides, the larger the value of  $\xi$ , the smaller the density of points over the surface of the approximated maximal set, and the greater the chance to converge to an overly large approximation of the true maximal set. Procedures to generate the  $p^i$  points are available in reference [3].

## V. Robustness Analysis of the Tuned Nominal Controller

The methods of Section IV are first applied to the nominal controller (see the companion paper [6]). In particular, we search for the set of gains of the lateral/directional controller that maximize the CSR for  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$ . The constraint  $h(\hat{p}, d) > 0$  was used to enforce frequency domain and stability margin considerations. The same reference set used in the previous  $\delta C_{lp}$ - $\delta C_{m\alpha}$  deformation is used. We will denote the resulting controller  $c_{\text{nom, tuned}}$ . One- and multi-dimensional robustness analyses of this controller are presented next.

Table 3 presents the change in the PSMs attained by  $c_{\text{nom, tuned}}$  relative to  $c_{\text{nom, base}}$  for the one-dimensional cases in Table 1. This change is given by


$$\kappa_i(c_{\text{nom, tuned}}, c_{\text{nom, base}}) = \left( \frac{\rho_i(c_{\text{nom, tuned}})}{\rho_i(c_{\text{nom, base}})} - 1 \right) \times 100\% \quad (9)$$

where  $i = 1, \dots, 9$ . While positive values indicate a robustness improvement, negative values indicate a loss. Changes in the PSMs enable determining the relative advantages and drawbacks of using one controller over another according to particular requirements and uncertainties.

$\hat{p}$	$\kappa_1$	$\kappa_2$	$\kappa_3$	$\kappa_4$	$\kappa_5$	$\kappa_6$	$\kappa_7$	$\kappa_8$	$\kappa_9$
$\Lambda_{THL}$	0	0	0	0	0	0	0	0	0
$\Lambda_{ELO}$	0	0	0	0	0	0	0	0	-16
$\Lambda_{AL}$	0	0	0	-3	-2	-2	3	-3	-2
$\Lambda_{RU}$	0	0	0	0	0	0	0	0	2
$\tau$	0	0	0	0	0	0	0	0	0
$\mu$	0	0	0	0	0	0	0	0	0
$\Delta_x$	0	0	0	0	0	0	7	0	0
$\delta C_{lp}$	0	45	0	48	35	45	59	0	34
$\delta C_{m\alpha}$	0	0	0	-8	0	0	0	0	0
$\delta C_{nr}$	0	2	2	2	0	0	2	2	0

**Table 3.** Change in the PSMs of  $c_{nom, tuned}$  relative to  $c_{nom, base}$ .

Note that even though  $c_{nom, base}$  and  $c_{nom, tuned}$  have the same longitudinal controller, their robustness to failure in the left-outer elevator differ. This is because the CPV for  $\Lambda_{ELO}$  is negative, and therefore, the locked elevator is coupling the longitudinal and lateral/directional dynamics. The PSMs remain practically unchanged for most of the uncertainties,  $\delta C_{lp}$  being a notable exception. Figure 3 shows  $g(\hat{p}, c_{nom, base})$  and  $g(\hat{p}, c_{nom, tuned})$  as a function of  $\hat{p} = \delta C_{lp}$ . Note that the performance functions  $g$  for  $c_{nom, tuned}$  cross the threshold of satisfactory performance  $g = 0$  at values of  $\delta C_{lp}$  that are much larger than those corresponding to  $c_{nom, base}$ . Overall, the tuned controller is much more robust to aerodynamic uncertainties in the roll axis than the baseline. This improvement is attained without a significant reduction in the robustness to other uncertainties. As with  $c_{nom, base}$ ,  $c_{nom, tuned}$  was evaluated extensively by the pilot in a real time simulator and it was regarded suitable for flight. During its evaluation, and consistently with the results above, piloted simulations showed that  $c_{nom, tuned}$  tolerated 25% more uncertainty in  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  than  $c_{nom, base}$ .



**Figure 3.** Dependence of  $g$  on  $\delta C_{lp}$  for  $c_{nom, base}$  and  $c_{nom, tuned}$ . Line conventions are as follows:  $g_1$ (solid),  $g_2$ (squares),  $g_3$ (diamond),  $g_4$ (asterisk),  $g_5$ (triangle),  $g_6$ (dotted),  $g_7$ (dashed),  $g_8$ (circle),  $g_9$ (x)

Table 4 presents the change in the PSM and in the failure probability attained by  $c_{nom, tuned}$  relative to  $c_{nom, base}$  for the multi-dimensional cases in Table 2. Note that changes in the PSM and in the failure probability have con-



$\hat{p}$	$\left(\frac{\rho(c_{\text{nom, tuned}})}{\rho(c_{\text{nom, base}})} - 1\right) \times 100\%$	$P[g(c_{\text{nom, base}}) > 0] - P[g(c_{\text{nom, tuned}}) > 0]$
$[\tau, \mu]$	-43	0.015
$[\delta C_{lp}, \delta C_{m\alpha}]$	51	0.344
$[\Delta_x, \Delta_y]$	-9	-0.05
$[\tau, \delta C_{lp}]$	49	0.230
$[\tau, \delta C_{m\alpha}]$	0	0.005
$[\tau, \delta C_{nr}]$	0	0.03
$[\delta C_{lp}, \delta C_{nr}]$	52	0.25
$[\Lambda_{THR}, \Lambda_{RU,RL}]$	1	-0.02
$[\tau, \eta]$	49	0.1
$[\Lambda_{ELO}, \Lambda_{ELI}]$	0	-0.005
$[\tau, \delta C_{lp}, \delta C_{m\alpha}, \delta C_{nr}]$	33	0.265
$[\Lambda_{THR}, \dots, \Lambda_{RD}]$	-24	-0.17

**Table 4. Changes in PSMs and failure probabilities of  $c_{\text{nom, tuned}}$  relative to  $c_{\text{nom, base}}$ .**

sistent trends most of the time. In other words, an increase(decrease) of the PSM commonly implies a proportional decrease(increase) in the failure probability. This is the case in all analyses with the exception of  $\hat{p} = [\tau, \mu]$  for which a reduction of 43% in the PSM does not yield a sizable reduction in the failure probability<sup>c</sup>. This situation arises because the failure domain has a long and thin spike intruding deeply into the non-failure domain. This spike prevents the realization of a larger homothetic deformation. This is a situation where the PSM analysis is able to find critical failure modes (since they are the closest to the nominal operating conditions associated with  $\bar{p}$ ) that cannot be found otherwise. In this particular case the spike is caused by a violation of the roll rate tracking requirement  $g_4$ . Since the rate of change of  $g_4$  within the spike is very low, the effects of this violation can be easily overcome by the pilot. Note, however, that the failure probability and the PSM are qualitatively different metrics whose trends do not have to be in agreement.

Figure 4 shows the number of requirements violated by  $c_{\text{nom, tuned}}$  as a function of  $\delta C_{lp}$  and  $\delta C_{m\alpha}$ . The same color conventions used earlier apply. The maximal sets corresponding to  $c_{\text{nom, base}}$  and  $c_{\text{nom, tuned}}$  are superimposed. Note that the tuned controller is able to tolerate 52% larger  $\delta C_{lp}$  uncertainties than the pilot-approved baseline controller before reaching the failure domain. As before, Figure 5 shows the dependency of  $g$  on the aerodynamic uncertainty  $\eta$  and the time delay  $\tau$ . The comparison of the maximal sets and of the non-failure domains shows that the improvements in robustness to  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  do not reduce the robustness to time delay. With respect to the baseline controller, the tuned controller exhibits a substantial improvement to  $\delta C_{lp}$  uncertainty at the expense of a moderate reduction in the robustness to actuator failure. After considering the risk/benefit ratio of using  $c_{\text{nom, tuned}}$  illustrated by the above metrics, we decided to build the ACTS architecture around this controller.

## A. Flight Validation

The controller  $c_{\text{nom, tuned}}$  was successfully validated using the NASA AirSTAR Flight Test Vehicle. In the set of flight experiments conducted, nominal and off-nominal flying conditions were evaluated. Some of the off nominal conditions included a latency test and a set of failures where the aerodynamic characteristics of the vehicle were degraded progressively beyond open-loop neutral stability. These degradations are consistent with the manner in which  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  uncertainties are enforced in this paper. None of the other controllers flown, including a DF MRAC and an L1 adaptive controllers, were able to outperform this non-adaptive controller.

## VI. Robustness Analysis of the Tuned Adaptive Controller

The adaptive control parameters to be determined are  $Q$ ,  $\theta_{\text{max}}$ ,  $\Gamma_1$ ,  $\Gamma_2$ , and  $\Gamma_\lambda$  (see the companion paper [6]).  $Q$  prescribes the rate at which the Lyapunov function decreases in time,  $\theta_{\text{max}}$  prescribes the range of adaptation and  $\Gamma_1$ ,

<sup>c</sup>Even though the failure probability estimates have the numerical error caused by calculating probabilities with a finite number of samples, moderate failure probability differences are much larger than this error.




Figure 4. Maximal set, failure and non-failure domains of  $c_{nom, tuned}$  for  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$ .

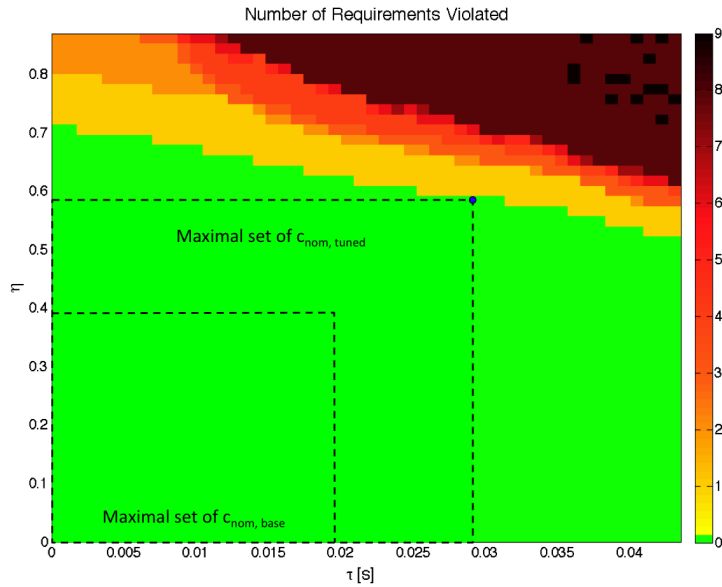


Figure 5. Maximal set, failure and non-failure domains of  $c_{nom, tuned}$  for  $\hat{p} = [\tau, \eta]$ .

$\Gamma_2$  and  $\Gamma_\lambda$  determine the rate of adaptation.

The methods of Section IV are now applied to determine some of these control parameters. In particular, we search for the set of adaptive control parameters  $\Gamma_1 > 0$  and  $Q$  that maximize the CSR for  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$ . The parameters kept were found to be the most consequential. As before, a lower bound on the time delay is enforced via the constraint function  $h(\hat{p}, d) > 0$ . We will denote the resulting controller  $c_{ada, tuned}$ . One may think of a baseline adaptive controller where  $c_{ada, base} = c_{base, tuned}$  by assuming a very small  $\Gamma_1$ . One- and multi-dimensional robustness analyses of the resulting  $c_{ada, tuned}$  are presented next.

Table 5 presents the change in the PSMs attained by  $c_{ada, tuned}$  relative to  $c_{nom, tuned}$  for the one-dimensional cases considered earlier. In this case, Equation (9) becomes  $\kappa_i(c_{ada, tuned}, c_{nom, base})$ . According to the data, the effects of adaptation are detrimental. Even in the  $\delta C_{m\alpha}$  and  $\delta C_{lp}$  directions only one PSM increases. However, the control tuning procedure indicates that the two dimensional maximal set in the  $\delta C_{m\alpha}$ - $\delta C_{lp}$  space corresponding to  $c_{ada, tuned}$  is slightly larger than that of  $c_{nom, tuned}$ . In the one-dimensional setting this is reflected on the 1% increase of the PSM of the critical requirement  $g_7$  for  $\delta C_{lp}$ . Therefore the 1% increase in the PSM of the critical requirement that enables having a larger maximal set causes a sizable reduction in the PSM of some non-critical requirements.

Figure 6 shows the number of requirements violated by  $c_{ada, tuned}$  as a function of  $\delta C_{lp}$  and  $\delta C_{m\alpha}$ . The maximal sets corresponding to  $c_{nom, tuned}$  and  $c_{ada, tuned}$  as well as the estimate of the boundary of the non-failure domain of  $c_{nom, tuned}$  are superimposed. Even though the maximal set for  $c_{ada, tuned}$  is very close to that of  $c_{nom, tuned}$ , the corresponding control parameters enlarge considerably the non-failure domain. This improvement, which is practically unnoticeable from the one-dimensional analyses, illustrates high sensitivity of the controller robustness to the joint effect of both uncertainties.

$\hat{p}$	$\kappa_1$	$\kappa_2$	$\kappa_3$	$\kappa_4$	$\kappa_5$	$\kappa_6$	$\kappa_7$	$\kappa_8$	$\kappa_9$
$\Lambda_{THL}$	0	0	0	0	0	0	0	0	-2
$\Lambda_{ELO}$	0	0	0	0	0	0	0	0	-19
$\Lambda_{AL}$	-17	-17	-72	-78	-4	-3	-11	-14	-47
$\Lambda_{RU}$	0	0	0	0	0	0	0	0	-2
$\tau$	0	0	0	-13	0	0	-9	0	-15
$\mu$	0	0	0	0	0	-7	0	0	-54
$\Delta_x$	0	-7	0	-5	0	0	-6	0	-78
$\delta C_{lp}$	0	-20	-30	-11	-13	-22	1	-19	-27
$\delta C_{m\alpha}$	0	0	0	-9	0	0	0	0	0
$\delta C_{nr}$	0	-18	-22	-19	-12	-23	-18	-17	-21

**Table 5. Changes in PSMs of  $c_{ada, tuned}$  relative to  $c_{nom, tuned}$ .**

Even though the robustness to  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  uncertainty was increased as intended, sizable reductions in robustness to  $\Lambda_{AL}$  and  $\tau$  also occur (see Table 5). Figure 7 shows the performance functions  $g(\Lambda_{AL}, c_{nom, tuned})$  and  $g(\Lambda_{AL}, c_{ada, tuned})$ . Note that at  $\Lambda_{AL} = -0.62$ , value where the left aileron is stuck at 62% of its maximum deflection, the adaptive controller becomes unstable, making all components of  $g$  to blow up. This “finite escape time”-like phenomenon, which is solely caused by adaptation, occurs abruptly without any lead time for pilot reaction. This illustrates a serious drawback of adaptation in which the controller itself causes instability. The elimination of this behavior requires smaller adaptive rates or smaller ranges of adaptation. Unfortunately, with these reductions the advantages and disadvantages of adaptation are both mitigated.

Table 6 shows the changes in the PSM and in the failure probability attained by  $c_{ada, tuned}$  relative to  $c_{nom, tuned}$  for the multi-dimensional settings considered above. This data indicates that the cost of the robustness improvements for  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  is a reduction in robustness to time delay, CG location and actuator failure. The reduction of the failure probability for  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$  is an incidental result from attempting to increase the corresponding PSM during control tuning. The increase in the PSM corresponding to  $\hat{p} = [\tau, \mu]$  results from the elimination of the spike mentioned above. The data in Table 6 is consistent with that for other optimization-based strategies where the optimal controller exhibits high sensitivity to those uncertainties which were left out of the problem formulation.

Figure 8 shows the dependency of  $g$  on the aerodynamic uncertainty  $\eta$  and the time delay  $\tau$ . The comparison of Figures 8 and 5 shows that the improvements in robustness to aerodynamic uncertainty results in a slight increase in the time delay margin in the vicinity of the maximal set. Note however, that the region where all requirements are

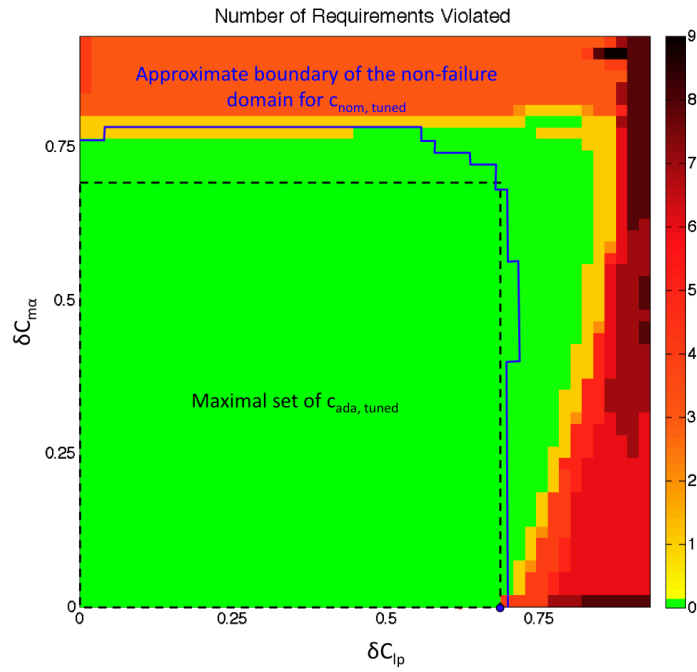


Figure 6. Maximal set, failure and non-failure domains of  $c_{ada, tuned}$  for  $\hat{p} = [\delta C_{ip}, \delta C_{m\alpha}]$ .


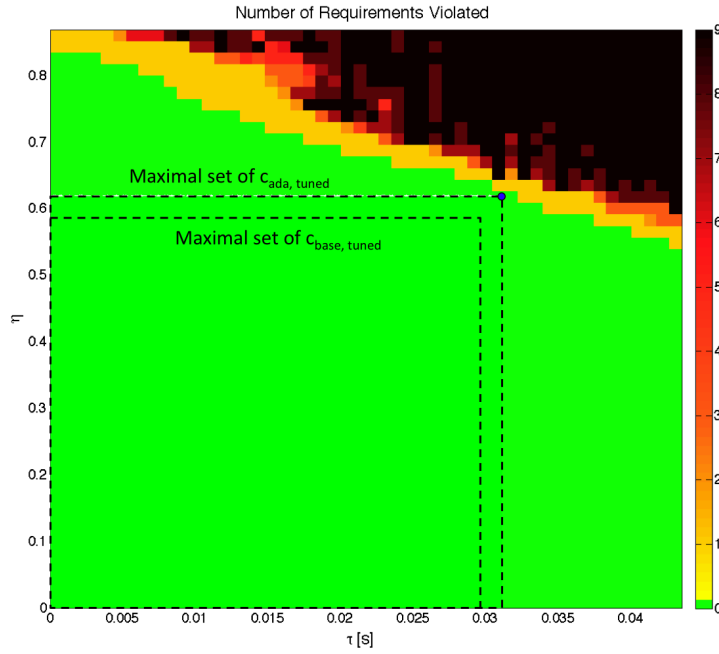


Figure 7. Dependence of  $g$  on  $\Lambda_{AL}$  for  $c_{nom, tuned}$  and  $c_{ada, tuned}$ .

$\hat{p}$	$\left(\frac{\rho(c_{\text{ada, tuned}})}{\rho(c_{\text{nom, tuned}})} - 1\right) \times 100\%$	$P[g(c_{\text{ada, tuned}}) > 0] - P[g(c_{\text{nom, tuned}}) > 0]$
$[\tau, \mu]$	57	1e-1
$[\delta C_{lp}, \delta C_{m\alpha}]$	1	-8e-2
$[\Delta_x, \Delta_y]$	-51	7e-2
$[\tau, \delta C_{lp}]$	-15	1e-1
$[\tau, \delta C_{m\alpha}]$	-2	3e-3
$[\tau, \delta C_{nr}]$	-32	1.7e-1
$[\delta C_{lp}, \delta C_{nr}]$	-22	1.6e-1
$[\Lambda_{THR}, \Lambda_{RU, RL}]$	0	-3e-3
$[\tau, \eta]$	7	-9e-3
$[\Lambda_{ELO}, \Lambda_{ELI}]$	0	3.5e-2
$[\tau, \delta C_{lp}, \delta C_{m\alpha}, \delta C_{nr}]$	-8	1e-1
$[\Lambda_{THR}, \dots, \Lambda_{RD}]$	-59	2.4e-1

**Table 6.** Changes in PSMs and failure probabilities of  $c_{\text{ada, tuned}}$  relative to  $c_{\text{nom, tuned}}$ .

violated is dangerously close to the non-failure domain.



**Figure 8.** Maximal set, failure and non-failure domains of  $c_{\text{ada, tuned}}$  for  $\hat{p} = [\tau, \eta]$ .

Figure 9 shows estimates of the failure and non-failure domain for  $c_{\text{nom, tuned}}$  corresponding to  $\hat{p} = [\tau, \delta C_{lp}]$ . The maximal set of both  $c_{\text{nom, tuned}}$  and  $c_{\text{ada, tuned}}$  are superimposed. The reduction in the size of the non-failure domain illustrates detrimental effects caused by adaptation.

## VII. Real Time Simulation

The nominal and adaptive controllers  $c_{\text{nom, tuned}}$  and  $c_{\text{ada, tuned}}$  were evaluated in the real time simulation. In this environment, and in contrast to the analyses above, pilot commands are generated in real time. Furthermore, there is a more accurate aerodynamic model, a surface dead-band in all control surfaces, sensor noise and moderate turbulence. The Figures 10 show closed-loop responses corresponding to degradations in  $\delta C_{lp}$  and  $\delta C_{m\alpha}$  that make the open-loop

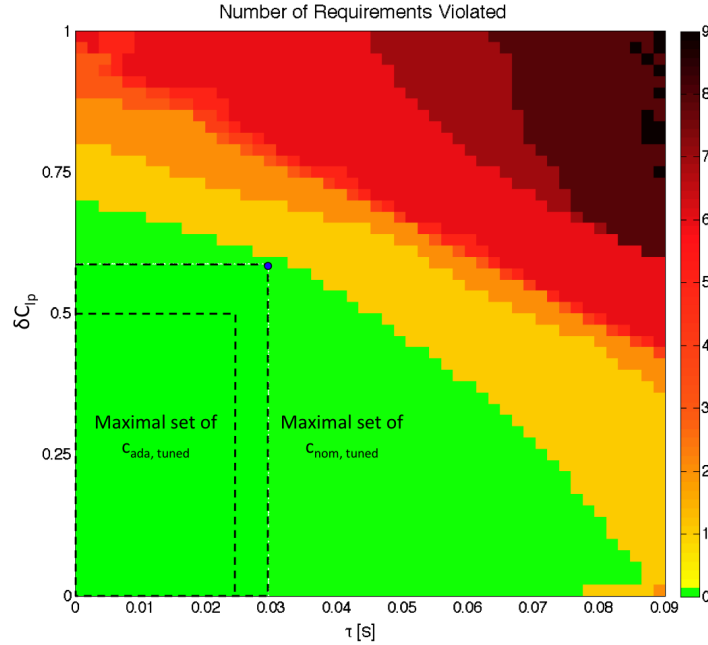


Figure 9. Maximal set, failure and non-failure domains of  $c_{nom, tuned}$  for  $\hat{p} = [\tau, \delta C_{lp}]$ .

unstable. The improvement in the transient roll rate response attained by adaptation is apparent. Unfortunately, the liabilities caused by adaptation when flying outside the region where the reference model accurately approximates the aircraft dynamics make this controller unsuitable for flight. Strategies to overcome this difficulty include expanding the range of accuracy of the reference model via gain scheduling and applying adaptive dead-zones. Efforts in these directions are underway.

## VIII. Rate of Adaptation

Figure 11 shows the failure and non-failure domains associated with  $\hat{p} = [\delta C_{lp}, \delta C_{m\alpha}]$  for four adaptive controllers. These controllers only differ in their adaptive rate  $\Gamma_1$  whose value is printed in the top left of each figure. Note that as the controller becomes more aggressive, spikes of the failure domain start intruding into the non-failure domain. These spikes cause a reduction in the PSM to the extent that for the most aggressive controller, the maximal set is empty (i.e., the controller does not satisfy the requirements for the nominal plant). Note that while the PSM decreases when the adaptive rate is increased from  $4\Gamma$  to  $8\Gamma$ , the size of the non-failure domain slightly increases.

Further notice that a probabilistic assessment of the most aggressive controller, such as one based on calculating failure probabilities via Monte Carlo, does not properly characterise its most prominent robustness characteristics (e.g., the fact that the controller does not even satisfy the requirements under nominal flying conditions). Besides the reduction in the size of the non-failure domain associated with this controller, note the small separation between such a domain and the region where all requirements are violated (i.e., region in black where the controller becomes unstable). The small separation between these two regions pose a grave safety risk since the transition to instability is not preceded by an observable degradation in closed-loop performance and as such, it will be very difficult to predict or avert. Recall that this mode of instability is caused by adaptation. Figure 11 illustrates the risk of over-tuning the control parameters according to particular uncertainty realizations and highlights the need for determining the robustness characteristics of the controller from a global perspective.

## IX. Conclusions

This paper concentrates on the evaluation and tuning of a controller for an unmanned air vehicle. The control architecture assumed augments a non-adaptive controller with an adaptive component. The gains of both these controllers are tuned using a framework that aims at increasing robustness to aerodynamic uncertainty. Global robustness

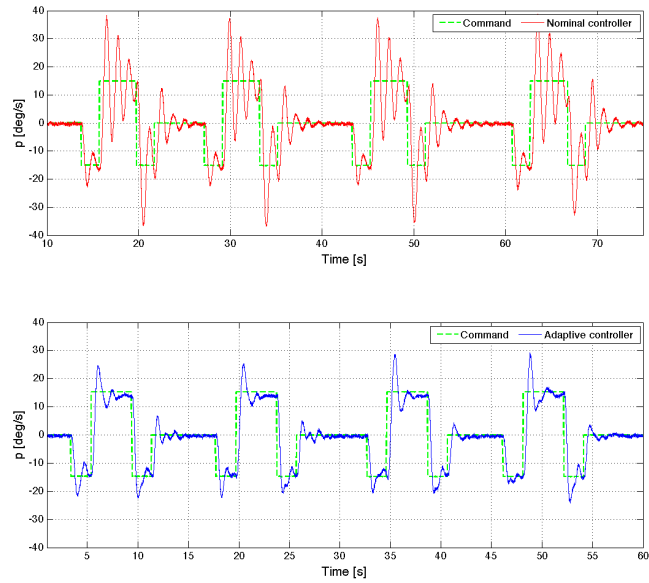


Figure 10. Closed-loop responses corresponding to the nominal and adaptive controllers




Figure 11. Failure and non-failure domains for increasingly faster adaptive controllers.

assessments of the resulting controllers are used to determine advantages and disadvantages of adaptation. While increments in the rate of adaptation have the potential of improving the robustness to aerodynamic uncertainties, they usually cause a reduction in the time delay margin. More importantly, suitable adaptation rates for some uncertainty realizations (i.e., rates for which the adaptive controller well-outperforms the nominal controller and improves safety) may turn out to be excessively high for other uncertainties (i.e., situations where the controller, and thereby the aircraft, becomes unstable). This instability, which may not be preceded by a noticeable degradation in closed-loop performance, occurs abruptly without any lead time for a pilot reaction. Failure to properly balance these trends may lead to controllers that compromise safety by adapting either too rapidly or too slowly. The nominal controller tuned by the methods proposed was successfully validated using the NASA AirSTAR Flight Test Vehicle. This flight experiment, where both the longitudinal and lateral aerodynamic characteristics of the aircraft were degraded beyond neutral stability, targets uncertainties that perfectly match the mathematical framework supporting the theory of adaptive control. None of the other controllers flown, including a DF MRAC and an L1 adaptive controllers, were able to outperform this non-adaptive controller. Efforts to make the adaptive controller suitable for flight are underway.

## References

- <sup>1</sup>National Aeronautics and Space Administration. Nasa technical standard. In *NASA-STD-7009*, Washington, DC, USA, 2008.
- <sup>2</sup>Karl J. Astrom and Bjorn Wittenmark. *Adaptive Control*. Addison-Wesley, Reading, MA, 2nd edition, 1994.
- <sup>3</sup>L. G. Crespo, D. P. Giesy, and S. P. Kenny. A verification-driven approach to control analysis and tuning. In *AIAA Guidance Navigation and Control Conference, Honolulu, Hawaii, AIAA-2008-6340*, August 2008.
- <sup>4</sup>L. G. Crespo, S. P. Kenny, and D. P. Giesy. A computational framework to control verification and robustness analysis. In *NASA TP 2010-216189*, pages 1–38, NASA Langley Research Center, Hampton, VA, 2010.
- <sup>5</sup>L. G. Crespo, S. P. Kenny, and D. P. Giesy. A computational framework to control tuning. In *NASA TC 2011-TBD*, page In review, NASA Langley Research Center, Hampton, VA, 2011.
- <sup>6</sup>L. G. Crespo, M. Matsutani, and A. Annaswamy. Design of a model reference adaptive controller for an unmanned air vehicle. In *AIAA Guidance Navigation and Control Conference, Toronto, Canada*, August 2010.
- <sup>7</sup>Luis G. Crespo, M. Matsutani, and Anuradha M. Annaswamy. Design and verification of an adaptive controller for the generic transport model. In *Proceedings of AIAA Guidance, Navigation, and Control Conference and Exhibit*, Chicago, IL, USA, 2009. AIAA 2009-5618.
- <sup>8</sup>Petros A. Ioannou and Jing Sun. *Robust Adaptive Control*. Prentice Hall, Upper Saddle River, NJ, 1996.
- <sup>9</sup>Y. D. Landau. *Adaptive Control: The model reference approach*. Marcel Dekker, New York: NY, 1979.
- <sup>10</sup>Kumpati S. Narendra and Anuradha M. Annaswamy. *Stable Adaptive Systems*. Prentice Hall, Englewoods Cliffs, NJ, 1989.
- <sup>11</sup>Shankar S. Sastry and Marc Bodson. *Adaptive Control: Stability, Convergence, and Robustness*. Prentice Hall, Englewood Cliffs, NJ, 1989.