2008 American Control Conference
Westin Seattle Hotel, Seattle, Washington, USA
June 11-13, 2008

Adaptive Control and the NASA X–15 Program: A Concise History, Lessons Learned, and a Provably Correct Design

Zachary T. Dydek, Anuradha M. Annaswamy, Eugene Lavretsky

Abstract— **The NASA X-15 research airplane was one of the earliest aircraft to feature an adaptive control scheme, making its first flight in 1959. The program is largely considered a success, the one exception being the fatal accident that occurred on November 15, 1967. The circumstances which led to the anomalous behavior of the X-15 during that flight are reproduced using a fully nonlinear six-degree-of-freedom aircraft model and a detailed model of the original MH-96 adaptive controller. Using the lessons learned from the X-15 program as well as decades of work in the field of stable adaptive control, a provably correct (PC) adaptive controller is designed for the X-15. When subjected to the same conditions that caused the original MH-96 adaptive controller to fail, the PC adaptive controller is able to recover and complete the maneuver successfully.**

I. INTRODUCTION

Air-breathing hypersonic vehicles are a more effective way of launching small satellites or other vehicles into low earth orbit (LEO) than expendable rockets. Over the years, the high performance missions involving X-15 and X-43A are instances of ambitious projects that have demonstrated impressive successes. With this type of aircraft, quick response and global strike capabilities for Air Force missions will be more practical. A crucial component of such projects is the controlled articulation of various inputs that allow desired maneuvers over large flight envelopes. This paper examines the role of advanced control in the earliest of such missions, the X-15, flown in the 50s and 60s.

Fig. 1. The USAF X-15 Aircraft in flight.

This work was supported by the Boeing Strategic University Initiative.

Z. Dydek (zac@mit.edu) and A. M. Annaswamy (aanna@mit.edu) are members of the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02138, USA.

E. Lavretsky (eugene.lavretsky@boeing.com) is a Technical Fellow of the Boeing Company - Phantom Works, Huntington Beach, CA 92647, USA.

Control of air-breathing hypersonic vehicles is a significant challenge that occurs precisely due to the changes in the aircraft dynamics as the maneuver takes them over large flight envelopes. The X-15 research airplane was one of the earliest aircraft to feature an adaptive control scheme, where certain parameters were adjusted in order to enforce uniform performance of the aircraft throughout flight envelope. Preliminary design work for the X-15 started in 1955 and the program recorded nearly 200 successful flights from 1959-1968. The program is largely considered a success, the one exception being the fatal accident that occurred on November 15, 1967. According to [1], events that lead to the accident can be described as follows. Shortly after the aircraft reached its peak altitude, it began a sharp descent; the aircraft had entered a Mach 5 spin. The pilot was able to recover from the spin, but the adaptive controller was unable to reduce the pitch angle and consequently the aircraft continued to dive. Encountering rapidly increasing dynamic pressures, the X-15 broke apart about 65,000 feet above sea level.

Fig. 2. The site of the 1967 crash.

The years 1970 to the present have witnessed the genesis and development of formal methodologies for adaptive control systems, focusing on the design and analysis methods for systems with parametric uncertainties. While the field began with the motivation of developing advanced controllers that can generate improved performance, the sobering lessons of tradeoffs between stability and performance of feedback control diverted the evolution of the field towards the design, analysis, and synthesis of stable adaptive systems. With the history of this field and the efforts of dedicated research over the past thirty years, we are now at a stage where several adaptive control methods are in vogue that can be used for the control of linear and nonlinear dynamic systems with parametric and dynamic uncertainties [2], [3], [4], [5], [6].

Keeping in mind the lessons learned from the X-15 and numerous other programs [7], we pose the following question: what if we were to design the flight control system for the X-15 today? To answer this question, we present a dissection of the X-15 aircraft dynamics and the Minneapolis Honeywell MH-96 adaptive controller in an effort to better understand how the sequence of events and the interplay between the controller and the aircraft dynamics led to the instability and the eventual crash. We follow it with a depiction of a provably correct (PC) adaptive control architecture that answers the question of what if the task of designing the adaptive controller for the X-15 were to be presented today and what results would accrue if some of the stable adaptive control principles outlined in the literature were to be adopted. As the readers may guess, indeed it is shown in this paper that had the X-15 controllers been implemented today, all of the 200 flights, without a single exception, would have been performed safely, without incident.

In order to evaluate the MH-96 and the PC adaptive controllers, a fully nonlinear six-degree-of-freedom aircraft model is formulated using suitable aerodynamic data from a variety of sources. A model of the MH-96 adaptive controller was synthesized based on the descriptions in [8], [9], [10], [11]. These are presented in Section II. The PC adaptive controller based on theory [12] explicitly takes into account the structure of the aircraft dynamics, and assumes that the parameters of the system are unknown, that the flight envelope encompasses multiple trim points, and that the requisite actuators can be driven into saturation by virtue of its high performance goals [2]. These are described in Section III.

II. THE MH-96 ADAPTIVE CONTROLLER AND THE 1967 INCIDENT

The Minneapolis Honeywell MH-96 "self-adaptive" controller, originally developed for the X-20 Dyna-Soar, was found to be particularly suitable for the X-15 [1]. The design intent was to achieve high performance throughout the flight envelope. Towards that end, it was observed that rapid changes in the forward-loop gain would be required, and that the controller gains needed to be near their critical values at all times, thereby necessitating a control design which constantly adjusted these gains. To accomplish this, the system output was monitored in the frequency range at which instability was expected to occur. When the signs of instability became apparent, the gains were reduced to maintain stability. When no instability was observed the gains were increased. In this manner, the gains were kept as high as possible while maintaining system stability throughout the entire flight envelope.

Three different hypersonic planes, the X-15-1, X-15-2, and X-15-3, were flown as a part of the NASA X-15 program [1]. Only X-15-3 was equipped with the MH-96 adaptive controller. Pilot ratings for the X-15-3 were generally equal to or higher than those of the X-15-1 and 2, which both employed a standard fixed gain stability augmentation system (SAS). This was especially noticeable during the re-entry phase, when changes in the flight condition were most dramatic [8]. In addition to the performance advantages, the adaptive controller required no external gain scheduling and therefore could be designed and implemented quickly and efficiently. For mainly these reasons the MH-96 was flown extensively with great success.

A. Modeling the X-15

In this section, we describe the aircraft model of the X-15, a model of the MH-96 adaptive controller, and the overall control loop architecture. Using this closed-loop system, we present sufficient conditions that may have led to the crash.

The X-15 dynamics are modeled using five subsystems: Equations of Motion, Aerodynamics, Actuator Dynamics, Actuator Saturation, and Sensor Dynamics. The overall control architecture includes inner-loop and outer-loop controllers, with the outer-loop intended for the slow states such as altitude and speed, and the inner-loop for the fast states that evolve according to the aircraft longitudinal and lateral dynamics. Each of these subsystems is described in more detail below.

1) Conservation Equations: The standard conservation equations assuming a flat earth, which can be found in any introductory aircraft text [13].

2) Aerodynamics: The aerodynamics of the X-15 were built up using a component approach which separates each force and moment into a linear combination of individual states and control surface deflections [13]. The coefficients of the linear combinations, the so-called non-dimensional derivatives, vary substantially with the angle-of-attack α as well as with the Mach number M as the X-15 moves through its flight envelope. A more detailed discussion of this approach can be found in [14]. These non-dimensional derivatives along with the standard conservation equations completely describe the open-loop dynamics of the X-15.

3) Actuators and Sensors: The X-15 had four control inputs in total: δ_{th}, which generates thrust through an XLR-99 rocket engine; δ_{f_1}, and δ_{f_2}, the combined pitch/roll control surfaces on each wing; and δ_r, a large ventral rudder. The control input deflections were actuated by irreversible hydraulic systems. The dynamics of these actuators can be modeled as second order oscillators with damping ratio $\zeta = 0.7$ and natural frequency ω_n = 90 Hz for the elevators/ailerons and ω_n = 70 Hz for the rudder. The actuator saturation limits were taken to be $\pm 30^\circ$ for both the elevator/ailerons and the rudder.

The aircraft's angular rates p, q, and r were measured by rate gyroscopes, however the dynamics of these sensors were neglected for this simulation.

4) The Feedback Controller: The control architecture, as mentioned previously, is composed of two feedback loops, with the outer-loop for the slower states, and the inner-loop for the faster ones. The outer-loop is simply a PID controller, and ensures that the altitude and speed follow the commanded signals closely. The inner loop controller is a reconstruction of the MH-96 adaptive controller [8], [9],

[10], [11] and it has a more complex structure, with three individual loops for each of the pitch, roll, and yaw axes.

As opposed to the standard fixed gain augmentation systems used previously, the MH-96 controller augmented the pilot's inputs with an adaptive control input. This adaptive input was found by feeding the error between the measured and the desired angular rates through a variable gain and a servo loop which essentially acts as a high pass filter. In this manner, the angular rates are made to follow the desired rates. Furthermore, the variable gain is adjusted aggressively over the course of the maneuver so that it stays as high as possible while maintaining stability. This ensures high performance over the entire flight envelope.

The MH-96's model of the X-15, or the desired aircraft dynamics are given by a first order lag with time constants $\tau_m = 0.5$ in the pitch axis and $\tau_m = 0.33$ in the roll axis [8]. The error between the desired angular rates (e.g. q_m) and the measured rates (q) is fed back through a variable gain (in this case k_q). As opposed to extensive external scheduling of these gains, the MH-96 utilized an adaptive algorithm to make changes to the gains online, via a gain computer. The gain changes are initiated based on the amplitude of the control output at frequencies where system instability may occur. Since the dominant frequencies of interest were observed to be around 0.5 Hz, the controller output δ_{ad} is passed through a bandpass filter centered at that frequency. The resulting signal is then rectified and compared to a constant set-point k_{set}. The gain computer input y is thus

$$y = |G_f(s)\delta_{ad}| - k_{set} \tag{1}$$

where k_{set} is the threshold between acceptable and unacceptable oscillations. The resulting controller matched the behavior of the MH-96 as reported in the literature [14].

The outer loop controller, on the other hand, is designed to operate on a much slower time-scale than the inner loop controller. The inputs to this controller are $u = [V_t \quad h]^T$ as well as the commanded trajectory as a function of time

$$u_{cmd}(t) = [V_{cmd}(t) \quad h_{cmd}(t)]^T \tag{2}$$

The commanded trajectory was extracted from the literature as an example of a typical high altitude mission [15]. The output of the outer loop control is the reference signal $\delta_c = [\delta_{th} \quad \delta_e]^T$ which consists of thrust and elevator commands. The controller is a PID (with approximate derivative) control on the error between the actual and the commanded values. The gains for both the speed and altitude loops are selected using the Ziegler-Nichols ultimate sensitivity method [16]. This loop is meant to simulate the corrective action of the pilot which keeps the aircraft traveling along the planned trajectory.

B. The 1967 Incident

With the reconstruction of the overall X-15 dynamics and the MH-96 controller implemented in the test flights, we now investigate the accident itself and a possible explanation for how the sequence of events transpired.

While flight recorded data from the actual crash is not available, much of the information can be reconstructed from the transcript of communication between the pilot and ground control [1]. With this reconstruction, the order of the events that occurred was as follows:

(1) 85,000 ft: Electrical Disturbance slightly degrades control, pilot switches to backups.
(2) Planned wing-rocking procedure was excessive.
(3) X-15 began slow drift in heading.
(4) 266,000 ft: Peak Altitude — drift in heading pauses with airplane yawed 15° to the right.
(5) Drift continues, plane begins descending at right angles to the flight path.
(6) X-15 enters a Mach 5 spin.
(7) 118,000 ft: Pilot recovers from the spin, enters inverted Mach 4.7 dive.
(8) MH-96 begins limit cycle oscillation, prevents any further recovery techniques.
(9) X-15 experiences 15-g vertically, 8-g laterally — aircraft breaks apart.

The X-15 open-loop dynamics, the multi-loop control architecture described in Section II-A.4, together with MH-96 adaptive controller were simulated to represent the overall flight control system and the flight dynamics. In addition to these parameters, the X-15 aerodynamics block utilizes a lookup table for each of the non-dimensional coefficients with axes in Mach number, M and angle-of-attack α. The data for these tables was extracted from various sources: flight recorded data [17], [18], [19], wind tunnel measurements [19], [20], [21], and theoretical calculations [19], [21].

The tracking performance obtained under nominal conditions is satisfactory–the altitude error is generally less than 1% of the maximum altitude achieved. The speed error is generally less than 3% of the maximum speed.

We now introduce the effect of an electrical disturbance, described in Event (1), as an 80% loss of controller effectiveness on the right elevator/aileron at $t = 80$ s. The first indication of an instability is a steady drift in heading angle ψ, as in Event (3), which can be attributed to the asymmetry of the control failure. In Figure 3**(a)**, we can see this drift is initiated at the onset of the disturbance, and that the drift oscillates around 15° for a period between $t = 120$ s and $t = 200$ s as described in Event (4). Examination of the aircraft states reveals additional similarities between the simulation and Events (1)-(9). In Figure 4, the X-15 makes a dramatic departure from the commanded trajectory as it nears its peak altitude as described by Event (5). The simulation is stopped at $t = 385$ s when the altitude reaches 0 ft; however, the accuracy of the model most likely breaks down earlier, perhaps soon after the dive is initiated at a time between $t = 250$ s and $t = 300$ s. Figure 3**(b)** shows that this dive is indeed a spin, corresponding to Event (6). We can also observe limit cycle behavior in the adaptive gain system, Event (8), which does indeed prevent the adaptive controller from reducing its pitch.

It is important to note that some of the events reported in [1] were not reproducible. For example, the planned

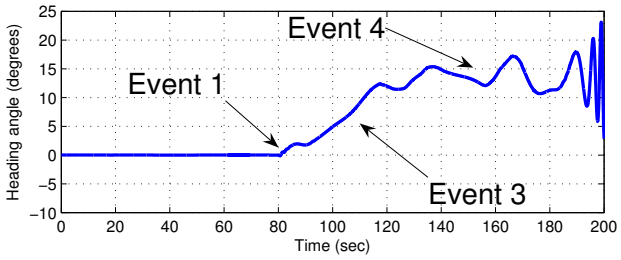

(a) Heading Angle

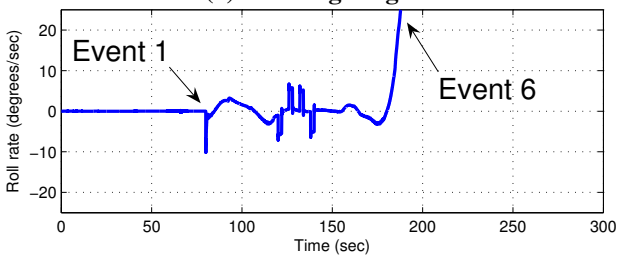

(b) Roll Rate

Fig. 3. X-15 heading angle ψ and roll rate r. In **(a)**, we can see the slow drift in heading, which briefly halts at around 15° as the X-15 reaches it's peak altitude. In **(b)**, we can see that the roll rate becomes very large as the X-15 enters the dive, corresponding to a rapid spin.

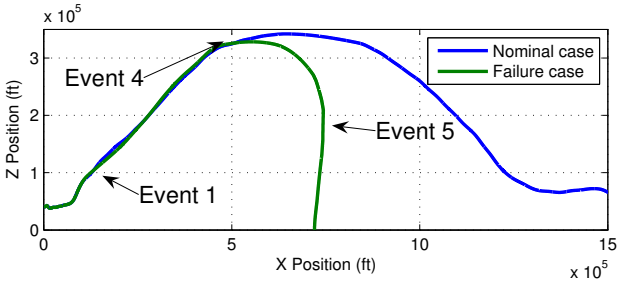

Fig. 4. Flight Path of the X-15 and MH-96 adaptive controller.

wing-rocking maneuver in Event (2) was not excessive in our simulation. In Figure 3**(b)**, it can be seen that the aircraft is unable to recover from the spin as described in Event (7). This is most likely due to the fact that the pilot employed spin-recovery techniques which are outside the scope and capability of our controller. With the exception of these differences, the remaining events of the crash were all observed in our studies as well.

The summary of the above study is this: a sudden change in the actuator characteristics (which could have been caused by the electrical disturbance) simulated as an 80% loss in control effectiveness causes a significant change in the aircraft dynamics, this in turn causes the dynamics to depart significantly from those represented in the model and therefore in the control design. As a result the control gain choices, despite the flexibility provided by the adaptive feature, are inadequate, causing the overall control system to be unable to recover from the onset of instability leading up to the crash.

III. PROVABLY CORRECT ADAPTIVE CONTROL

The last four decades have been a witness not only to the evolution of adaptive control theory but also the field of control theory as a whole. Notions of state, controllability, observability, stability, robustness, and uncertainty management have been studied in depth and breadth and applied to the control of complex dynamic systems in several problems. Several methods of control synthesis are currently available that are capable of accommodating the specific nature of the dynamics in a given problem. In this section, we utilize appropriate tools of control synthesis in general and adaptive control in particular to control the X-15 dynamics. Towards this end, the MH-96 inner-loop controller is replaced with an adaptive controller. The overall block diagram of the PC adaptive controller is shown in Figure **??**, which includes an integral controller, an outer-loop controller to compensate for the slow states of the Mach number and altitude, and an inner-loop adaptive controller to compensate for the fast states. In Section III-A, the latter discussed in more detail.

A. The Modern Inner-loop Adaptive Controller Design

In contrast to the inner-loop controller discussed in Section II-A.4, the PC adaptive controller is based on theory, accommodates coupling between different states, actuator saturation, and multiple parametric uncertainties, includes the measurement of more longitudinal and lateral states, integral action, baseline control action, and on-line adjustment of several parameters. The procedure for the design of both baseline control components and adaptive control components are discussed in this section. The crucial advantages that accrue from this procedure, which include improved performance and the ability to guarantee stable behavior, are also discussed here.

1) Baseline Controller: The baseline controller used in our study is an LQ controller which includes integral action on the fast aircraft states. A schedule of LQ gain matrices is designed by linearizing the flight dynamics

$$\dot{X}_t = f_p(X_t, U_t), \tag{3}$$

at multiple trim points $(X_{t_{0_i}}, U_{t_{0_i}})$ selected so that they sample the flight regime.

The fast states $X = [\alpha \quad \beta \quad p \quad q \quad r]^T$ and the corresponding control inputs $U = [\delta_e \quad \delta_a \quad \delta_r]^T$ can be extracted from the full state vector X_t and the full control vector U_t. This leads to the linearized flight dynamics

$$\dot{x}_p = A_{p_i} x_p + B_{p_i} \delta + d_{p_i}, \tag{4}$$

where $x_p = X - X_{0_i} \in \Re^n$, $\delta = U - U_{0_i} \in \Re^m$, and $d_{p_i} \in \Re$ is a constant trim disturbance. The baseline controller is then designed using the state space matrices given by (4). To overcome the drift in the lateral dynamics due to the trim disturbance, augmented integral states were added in the roll rate p and the combined yaw rate/sideslip angle term $r - \beta$.

The nominal baseline LQ controller is then designed with K_{x_i}, the nominal feedback gain matrix designed for the dynamics given by (4) and (**??**) around the i^{th} trim point,

which minimizes a quadratic cost function. A schedule of nominal LQ gain matrices is thus constructed for the baseline controller.

2) Adaptive Controller: The main problem that needs to be addressed is the accommodation of uncertainties that occur due to actuator anomalies. These uncertainties are represented by a combination of two features: a parametric uncertainty matrix Λ that represents loss of control effectiveness, and a saturation nonlinearity in the actuator. Both these effects are incorporated in the linearized dynamics (4) as

$$\dot{x}_p = A_{p_i} x_p + B_{p_i} \Lambda \text{sat}(\delta) + d_{p_i}, \tag{5}$$

where the saturation function $\text{sat}(\delta)$ is defined as

$$\text{sat}(\delta) = \begin{cases} \delta & \text{if } |\delta| \leq \delta_{max} \\ \delta_{max} \text{sgn}(\delta) & \text{if } |\delta| > \delta_{max} \end{cases}, \tag{6}$$

and δ_{max} is the known input saturation limit. This leads to an augmented plant dynamics given by

$$\begin{bmatrix} \dot{x}_p \\ \dot{x}_c \end{bmatrix} = \begin{bmatrix} A_{p_i} & 0 \\ B_c & A_c \end{bmatrix} \begin{bmatrix} x_p \\ x_c \end{bmatrix} + \begin{bmatrix} B_{p_i} \\ 0 \end{bmatrix} \Lambda \text{sat}(\delta) + \begin{bmatrix} d_{p_i} \\ 0 \end{bmatrix}, \tag{7}$$

or equivalently

$$\dot{x} = A_i x + B_i \Lambda u + d_i, \tag{8}$$

where $u = \text{sat}(\delta)$. The overall dynamics given by (8) are used for the adaptive control design.

In order to ensure safe adaptation, target dynamics are specified for the adaptive controller using a reference model. This model is selected using the baseline controller and the plant dynamics without actuator uncertainties. Thus, the main goal of the adaptive controller is to recover/maintain baseline closed-loop system performance, in the presence of uncertainties. The reference model is defined as:

$$\dot{x}_{ref} = (A_i + B_i K_{x_i}) x_{ref} = A_{ref_i} x_{ref}, \tag{9}$$

where the matrices A_{ref_i} are Hurwitz.

Using (8) and (9), an adaptive control input is added to the baseline controller as

$$\delta_{ad} = \theta_x^T(t) x + \theta_d(t) = \Theta^T \omega, \tag{10}$$

where $\Theta^T = [\theta_x^T \quad \theta_d^T]$ are adaptive parameters that will be adjusted in the adaptive law given in (11) below and the linear regressor $\omega = [x \quad 1]^T$.

The error e between the state of the plant and that of the reference model may be the result of a number of factors including both parametric uncertainties and the effects of actuator saturation. However, by exploiting our explicit knowledge of the actuator saturation limits, we can calculate the error due to saturation e_Δ and instead only adapt to the augmented error $e_u = e - e_\Delta$. This approach provides guaranteed stability in the presence of actuator saturation limits, as shown in [2].

The adaptive law given by [22],

$$\dot{\theta} = -\Gamma \omega e^T P B_p sign(\Lambda) - \theta \left(1 - \frac{\|\theta\|}{\theta^*_{max}}\right)^2 f(\theta) \tag{11}$$

where θ^*_{max} is a known constant and

$$f(\theta) = \begin{cases} 1 & \text{if } \|\theta\| \geq \theta^*_{max} \\ 0 & \text{otherwise,} \end{cases} \tag{12}$$

ensures guaranteed stability while providing an upper bound on the adaptive parameters Θ. The stability analysis for this controller is carried out by linearizing (3) about a desired command trajectory [23].

The adaptive gains at the i^{th} trim point Γ_i are selected according to an empirical formula, which arises from inspection of the structure of the adaptive laws [5]. The full inner-loop control input becomes

$$\delta = \delta_{nom} + \delta_{ad}. \tag{13}$$

B. Simulation Results

The first step in the simulation study is to select multiple trim points (X_{0_i}, U_{0_i}) so that they span the commanded trajectory, which is a path in the space with coordinates of altitude and speed. Instead of attempting to span the entire flight envelope, trim points were selected along the commanded trajectory. It was found that 10 such trim points provided good nominal control performance during the maneuver.

The next step is to simulate the PC adaptive inner-loop controller described above in (10)-(12) with the X-15 model and PID outer loop controller. The initial conditions for the simulation are given by (X_{0_1}, U_{0_1}), that is, the aircraft begins trimmed at the 1^{st} trim point.

In the nominal case where no failures are present ($\Lambda = I^{3x3}$), the adaptive controller tracks the commanded trajectory given by (2) with a level of performance similar to that of the MH-96 adaptive controller. The performance of the two controllers in the failure case ($\Lambda = \text{diag}([1 \quad 1 \quad 0.2 \quad 1])$) however, is quite different. Not only does the adaptive controller maintain stability, as opposed to the MH-96 controller which loses stability, the performance in the failure case is comparable to that of the nominal case. That is, the desired performance of the PC adaptive controller is retained despite the severe parametric uncertainty. The trajectories for each of these cases are shown in Figure 5. The adaptive controller accomplishes this while at or near

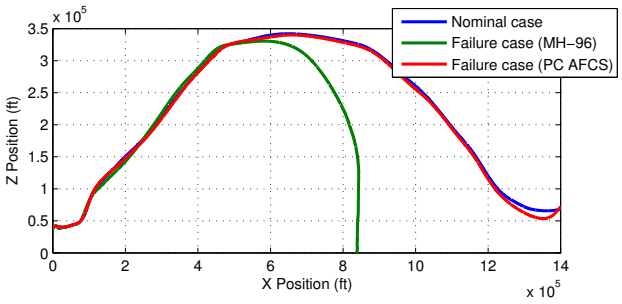

Fig. 5. Comparison of the flight path of the X-15 under nominal conditions and in the failure case with both MH-96 and PC adaptive controllers.

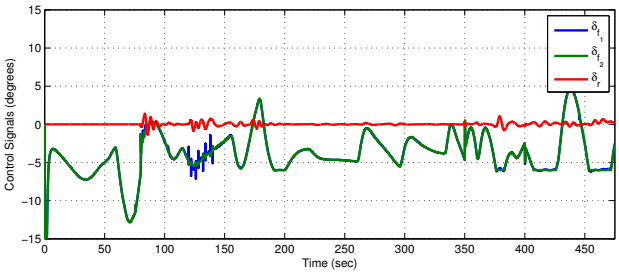

Fig. 6. The control inputs commanded by the PC adaptive system for the maneuver above (failure case). Note: the right actuator's limit is $\pm 6^\circ$ because of the 80% loss of control effectiveness.

the actuator limits for significant periods of time as can be seen in Figure 6 below.

The PC adaptive controller succeeds in stabilizing the X-15 where the MH-96 controller failed for several reasons. First, the PC adaptive controller includes the measurement of more longitudinal and lateral states and accommodates coupling between the states. Second, the adaptive controller is designed to specifically to accommodate actuator non-linearities in the form of saturation, as well as parametric uncertainties corresponding to a loss of actuator effectiveness. Lastly, the adaptive controller is built on top of a well-designed LQ baseline controller which includes integral action in the lateral states. In addition to these advantages, the PC adaptive controller is based on theory, and a proof of it's stability can be found in [23].

IV. SUMMARY

It has been shown that an accurate model of the X-15 aircraft and the original controller performs satisfactorily under nominal conditions. However, when subjected to a severe disturbance, the system fails, displaying much of the anomalous behavior observed during the crash of the actual X-15 in 1967. When the original MH-96 controller is replaced by the PC adaptive controller, the X-15 not only achieves high performance in the nominal case, but also exhibits increased robustness to uncertainties. Indeed, when subjected to the same failure as the X-15 equipped with the MH-96, the PC adaptive controller maintains both stability and much of its performance, completing the maneuver safely.

The original MH-96 adaptive flight control system accomplished its performance goal of providing a nearly invariant response across all flight conditions. Furthermore, it showed that a satisfactory adaptive control system could be designed without necessarily having accurate a priori information about the aircraft aerodynamics, and consequently, aircraft configuration changes could be easily accounted for [8]. However, the MH-96 lacked an analytically based proof of stability, which was highlighted by the fatal crash in 1967. After four decades, the theoretical ground work for the application of stable adaptive controllers has corrected this deficiency in the adaptive controllers of the 1960's.

REFERENCES

[1] D. R. Jenkins, "Hypersonics before the shuttle: A concise history of the X-15 research airplane," NASA, Special Publication SP-2000-4518, June 2000, Monographs in Aerospace History: Number 18.

[2] S. P. Karason and A. M. Annaswamy, "Adaptive control in the presence of input constraints," *IEEE Transaction on Automatic Control*, vol. 46, no. 11, pp. 2325–2330, November 1994.

[3] M. Schwager, A. M. Annaswamy, and E. Lavretsky, "Adaptation-based reconfiguration in the presence of actuator failures and saturation," in *Proc. American Control Conference (ACC'05)*, vol. 4, Portland, Oregon, June 2005, pp. 2640–2645.

[4] M. Schwager and A. M. Annaswamy, "Adaptive control of multi-input systems with magnitude saturation constraints," in *Proc. Conference on Decision and Control (CDC'05)*, Seville, Spain, December 2005.

[5] Z. T. Dydek, H. Jain, J. Jang, A. M. Annaswamy, and E. Lavretsky, "Theoretically verifiable stability margins for an adaptive controller," in *Proc. AIAA Conference on Guidance, Navigation, and Control (GNC'06)*, Keystone, Colorado, August 2006.

[6] J. Jang, A. M. Annaswamy, and E. Lavretsky, "Towards verifiable adaptive flight control in the presence of actuator anomalies," in *Proc. Conference on Decision and Control (CDC'06)*, San Diego, California, December 2006.

[7] W. H. Dana, "The X-15 lessons learned," NASA Dryden Research Facility, Tech. Rep. AIAA-93-0309, 1993.

[8] Staff of the Flight Research Center, "Experience with the X-15 adaptive flight control system," NASA, Flight Research Center, Edwards, CA, Technical Note D-6208, March 1971.

[9] J. Taylor, Lawrence W. and E. J. Adkins, "Adaptive flight control systems - pro and con," NASA, Flight Research Center, Edwards, CA, Technical Memorandum X-56008, April 1964.

[10] K. J. Åström, "Adaptive control around 1960," in *Proc. 34th IEEE Conference on Decision and Control (CDC'95)*, New Orleans, Louisiana, December 1995, pp. 2784–2789.

[11] B. Boskovich and R. Kaufmann, "Evolution of the honeywell first-generation adaptive autopilot and its applications to F-94, F-101, X-15, and X-20 vehicles," *Journal of Aircraft*, vol. 3, no. 4, pp. 296–304, 1966.

[12] K. S. Narendra and A. M. Annaswamy, *Stable Adaptive Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1989.

[13] R. C. Nelson, *Flight Stability and Automatic Control*. McGraw-Hill, 1989.

[14] Z. T. Dydek, "Guaranteed margins and performance for an adaptive flight control system and application on the x–15 research airplane," Master's thesis, MIT, Cambridge, MA, June 2007.

[15] Staff of the Flight Research Center, "Progress of the X-15 research airplane program," NASA, Flight Research Center, Edwards, CA, Special Publication SP-90, October 1965, p.74.

[16] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," *ASME Transactions*, vol. 65, pp. 433–444, 1942.

[17] R. B. Yancey, H. A. Rediess, and G. H. Robinson, "Aerodynamic-derivative characteristics of the X-15 research airplane as determined from flight tests for Mach numbers from 0.6 to 3.4," NASA, Flight Research Center, Edwards, CA, Technical Note D-1060, January 1962.

[18] R. B. Yancey, "Flight measurements of stability and control derivatives of the X-15 research airplane to a Mach number of 6.02 and an angle of attack of 25°," NASA, Flight Research Center, Edwards, CA, Technical Note D-2532, November 1964.

[19] H. J. Walker and C. H. Wolowicz, "Stability and control derivative characteristics of the X-15 airplane," NASA, Flight Research Center, Edwards, CA, Technical Memorandum X-714, March 1962.

[20] E. J. Saltzmm and D. J. Garringer, "Summary of full-scale lift and drag characteristics of the X-15 airplane," NASA, Flight Research Center, Edwards, CA, Technical Note D-3343, March 1966.

[21] H. J. Walker and C. H. Wolowicz, "Theoretical stability derivatives for the X-15 research airplane at supersonic and hypersonic speeds including a comparison with wind tunnel results," NASA, Flight Research Center, Edwards, CA, Technical Memorandum X-287, August 1960.

[22] G. Kreisselmeier and K. S. Narendra, "Stable model reference adaptive control in the presence of bounded disturbances," *IEEE Transactions on Automatic Control*, vol. 27, no. 6, pp. 1169–1175, December 1982.

[23] J. Jang, A. M. Annaswamy, and E. Lavretsky, "Adaptive control of time-varying systems with gain-scheduling," in *Proc. American Control Conference (ACC'08) Submitted.*, Seattle, Washington, June 2008.