Guaranteed Margins and Performance for an Adaptive Flight Control System and Application on the X–15 Research Airplane

by

Zachary Thompson Dydek

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2007
[June 2007]

© Massachusetts Institute of Technology 2007. All rights reserved.

Author ....
Department of Mechanical Engineering
June 11, 2007

Certified by ....
Anuradha M. Annaswamy
Senior Research Scientist
Thesis Supervisor

Accepted by ....
Lallit Anand
Chairman, Department Committee on Graduate Students

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY
JUL 1 8 2007
LIBRARIES

BARKER

Robustness Margins and High Performance for an Adaptive Flight Control System with Application to Hypersonic Vehicles

by

Zachary Thompson Dydek

Submitted to the Department of Mechanical Engineering
on June 11, 2007, in partial fulfillment of the
requirements for the degree of
Master of Science in Mechanical Engineering

Abstract

The design tools developed for use with linear controllers such as gain and phase margins do not apply to nonlinear control architectures such as adaptive control. For decades, flight control engineers have used these tools extensively to measure the robustness of their linear control systems and make guarantees on the performance of the closed-loop system in the presence of uncertainties. Stringent demands on performance for safety-critical flight systems, as in the case of hypersonic vehicles, make advanced control methods such as adaptive control increasingly attractive. The major obstacle in the widespread application of adaptive control to such applications is the lack of guarantees on performance and robustness. This thesis presents robustness margins, adaptive control analogs to the linear control notions of gain and phase margins, which can be used to make those guarantees. This paves the way for a systematic Verification and Validation (V&V) approach for adaptive controllers.

The operation of an adaptive controller can be broken down into two distinct phases: the adaptation mode, in which the adaptive parameters are varying, and the steady-state mode, in which the adaptive parameters have converged to their steady-state values. During the steady-state mode, the nonlinear adaptive controller converges to a linear time-invariant (LTI) system, and many tools exist for the calculation of the requisite margins. However, during the adaptation mode, which is arguably a more crucial mode of operation for the aircraft, no such tools exist. This thesis provides the tools for the numerical calculation of robustness margins during the adaptation mode. Robustness with respect to a range of uncertainties including parametric uncertainties, disturbances, time-delays, unmodeled dynamics, and actuator saturation is derived. The robustness of the adaptive controller is then demonstrated on a fully nonlinear model of a high-performance hypersonic aircraft.

The importance of theoretically justified adaptive controllers is illustrated using the historical example of the NASA X–15 research airplane. NASA's three X–15 aircraft together flew nearly 200 flights, acting as test beds for many bleeding-edge technologies, including the nonlinear adaptive controller implemented on the X–15–3. The application of this controller demonstrated the advantages of adaptive control including improved performance and a shorter design cycle. However, when the X–15–3 crashed in 1967, one of the severe disadvantages of this early adaptive control was highlighted: the lack of guaranteed stability and performance. Using modern adaptive control theory and the tools developed in this thesis, the control design of the X–15 is revisited and it is demonstrated 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.

Thesis Supervisor: Anuradha M. Annaswamy
Title: Senior Research Scientist

Acknowledgments

Firstly, I would like to acknowledge my advisor Anuradha Annaswamy and my colleagues and lab-mates Jinho Jang and Himani Jain. I would like to thank Eugene Lavretsky and Manu Sharma at Boeing for many insightful discussions/e-mails. I would also like to thank the Boeing Strategic University Initiative and the National Defense Science and Engineering Graduate fellowship for supporting this work. Lastly, I want to thank my fiancée Vicki Loewer for her support.

Contents

1 **Introduction** **15**

2 **Modeling the X–15** **19**

2.1 X–15 Aircraft Model 19

2.1.1 Conservation Equations 20

2.1.2 Aerodynamics 21

2.1.3 Actuators and Sensors 24

2.2 The MH–96 Adaptive Controller 24

2.2.1 Overall control architecture 25

2.2.2 Inner-loop controller 27

2.2.3 Outer-loop controller 28

2.3 Modern Adaptive Control 30

2.3.1 The Modern Inner-loop Adaptive Controller Design 31

3 **Margins** **37**

3.1 Parametric Uncertainties 37

3.2 Unstructured Uncertainties 39

3.3 Numerical Calculation of Margins 40

4 **X–15 Demonstration** **45**

4.1 The 1967 Incident 45

4.2 The 2007 Approach 52

5 **Summary** **57**

List of Figures

2-1 The 6DOF X–15 Aircraft Model. 21

2-2 Planform for the X–15 research airplane.[5] 22

2-3 Schematic of the inner loop control architecture. These block diagrams represent the pitch axis control used in **(a)** the literature and **(b)** the simulation. The control loops for roll and yaw have essentially the same diagram[18]. 26

2-4 Gain Changer Operation. These plots display the typical gain changer performance and control surface activity for **(a)** the original MH–96 adaptive controller[16] and **(b)** the controller implemented for simulation. 29

2-5 Block diagram of the modern adaptive inner-loop controller. Note that pitch, roll, and yaw axes are combined so that $\delta_c = [\delta_{c_e} \quad \delta_{c_a} \quad \delta_{c_r}]^T$ and $\delta = [\delta_e \quad \delta_a \quad \delta_r]^T$. 31

2-6 Overall control structure for the modern adaptive control system. The baseline controller (composed of the Baseline Control and Integral control blocks) is augmented by the adaptive control. 36

3-1 Example of an aircraft control surface that has sustained some type of damage, resulting in a 75% loss of effectiveness. 38

3-2 The stability region for both the adaptive and the nominal system in the presence of parametric uncertainty. The adaptation rate is normalized by the value prescribed by an empirical rule given by Equation (2.27). Note that the nominal stability region shows no dependance on the adaptation rate since the baseline controller includes no adaptation. 41

3-3 The stability region for both the adaptive and the nominal system in the presence of unmodeled dynamics in the form of aeroelastic modes. 42

3-4 The stability region for both the adaptive and the nominal system in the presence of time delay at both the input and the output to the plant. 43

4-1 Tracking performance of the X–15 (nominal case). 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. 47

4-2 Tracking performance of the X–15 (failure case). Even after the failure at $t = 80$ s, the system continues to track fairly well until the dramatic departure from the commanded values around $t = 200$ s. 48

4-3 Flight Path of the X–15 (failure case) 49

4-4 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. 50

4-5 Limit Cycles in the X–15 adaptive flight control system. **(a)** shows the input to the adaptive gain computer over time. Recall that the forward-loop gain is reduced when this input signal becomes larger than the set-point. **(b)** displays a blow up of the signal between $t = 250$ s and $t = 275$ s, showing the existence of undamped oscillations which prevent the gain changer from correctly adjusting the gain. 50

4-6 Accelerations experienced by the X–15 (failure case). 51

4-7 The commanded trajectory in Altitude/Speed space. This figure shows the commanded path for the X–15 simulations along with labels of the locations of the trim points used for controller design. 52

4-8 Tracking performance of the X–15 with modern adaptive controller (nominal case). 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. 53

4-9 Tracking performance of the X–15 with modern adaptive controller (failure case). The adaptive controller maintains high performance despite the severe failure present. 54

4-10 The control inputs for the maneuver above. Note that while the left actuator's limit is at the standard value of ±30°, the right actuator's limit is actually ±6° because of the 80% loss of control effectiveness. . 55

List of Tables

2.1 Nomenclature 20

4.1 Simulation Parameter Values. 47
4.2 Simulation Parameter Values (Adaptive). 52

Chapter 1

Introduction

The application of adaptive control to aircraft promises benefits in both safety and robustness, especially for high-performance aircraft such as air-breathing hypersonic vehicles. Early attempts at adaptive flight control used controllers with unproven stability properties, sometimes with disastrous consequences; for example the fatal crash of the NASA X–15 in November, 1967. As a result, much of the theoretical work up to the present time has been rightly focused on stability of adaptive architectures. Currently, there exists an assortment of stable adaptive control strategies, as well as techniques for preserving stability in the presence of unknown, bounded disturbances[10, 9].

Despite these advances in the theory of stable adaptive systems, theoretically verifiable Verification and Validation (V&V) techniques for adaptive systems are conspicuously absent. Current V&V techniques rely heavily on the fact that the underlying control system is linear, which makes them inadequate for adaptive flight control systems which are intentionally nonlinear. During this V&V process the linear tools of gain and phase margin are used extensively to measure the relative stability of the closed-loop system. Currently, the chief practical obstacle to transitioning adaptive flight controllers into aerospace applications is an inability to analytically assert that the closed-loop system will have adequate stability/robustness margins with respect to time-delays, unmodeled dynamics, and disturbances. This work seeks to lay the foundation for methods of calculating these margins for adaptive systems.

In an adaptive flight control system the operation of the adaptive controller can be categorized into two distinct phases: the adaptation mode and the steady-state mode. During the adaptation mode, the adaptive parameters are being adjusted constantly, hence the system is nonlinear. Its stability and robustness properties and therefore its margins are determined by the corresponding nonlinear system. By contrast, in the steady-state mode the same adaptive flight control system transitions to a linear time-invariant system (in the case of constant command signals) when the adaptive parameters converge to their steady-state values. During this mode of operation, the margins are determined by the corresponding LTI system. This thesis concentrates on the calculation of the margins during the adaptation mode for two reasons. First, the adaptation mode occurs during periods of high activity, such as during the execution of a complex maneuver. During these periods, the aircraft is more likely to encounter large disturbances or excite unmodeled dynamics, making the notion of guaranteed robustness all the more important. Second, there already exist numerous tools for the calculation of margins for an LTI system, and these tools can be similarly applied to the steady-state mode of an adaptive system.

Control of air-breathing hypersonic vehicles is a significant challenge that occurs precisely due to the significant changes in the dynamics as the maneuver takes the aircraft over its large flight envelope. The X–15 research airplane was one of the earliest aircraft to feature an adaptive control scheme, wherein certain parameters were adjusted according to the aircraft's performance. Preliminary design work for the X–15 was started in 1955 and the program recorded nearly 200 flights from 1959–1968. The program is largely considered to be one of NASA's most successful programs, the one blemish being the fatal accident that occurred on November 15, 1967. 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 and consequently the aircraft continued to dive. Encountering rapidly increasing dynamic pressures, the X–15 broke apart about 65,000 feet above sea level.

Independent of the above snafu, the years 1970 to the present have witnessed the genesis and development of the field of adaptive control theory that has addressed the problem of control in the presence of 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 [8, 13, 14, 3, 6].

With the solid understanding of how adaptive control systems ought to be designed and the tools that need to be deployed while closing the loop with adaptation algorithms behind us, this is an appropriate moment in time to revisit the events of 1967 and play a "how and a what if" scenario. A dissection of the X–15 aircraft dynamics as well as the original Minneapolis Honeywell MH–96 adaptive controller is presented in an effort to better understand how the sequence of events and the interplay between the controller and the dynamics led to the instability and the eventual crash. This is followed by a depiction of a stable 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 stable adaptive control principles outlined in the literature were to be adopted. As the readers may guess, indeed it is shown that had the X–15 controllers been implemented now, all of the 200 flights, without a single exception, would have been performed safely, without incident.

Chapter 2

Modeling the X–15

In order to evaluate the stability and performance of the adaptive controller, a fully nonlinear six-degree-of-freedom aircraft model is formulated using suitable aerodynamic data from a variety of sources. Additionally, two separate control schemes are implemented: the original 1960's adaptive controller and a modern adaptive controller. A model of the MH–96 adaptive controller was synthesized based on the descriptions in [16, 18]. The stable adaptive controller based on theory [10] 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 [8]. In this section, we describe the aircraft model of the X–15, a model of the original MH-96 adaptive controller as well as our theory-based modern adaptive controller, and the overall control loop architecture.

2.1 X–15 Aircraft Model

The X–15 dynamics is modeled using five subsystems as shown in Figure 2-1. These include the Equations of Motion, Aerodynamics, Actuator Dynamics, Actuator Saturation, and Sensor Dynamics. The overall control architecture, also shown in Figure 2-1, includes an inner-loop and an outer-loop controller, with the outer-loop controlling the slow states such as altitude and speed, and the inner-loop controlling the fast

Table 2.1: Nomenclature

a_x	Acceleration (x-direction)	X	Aerodynamic Forces (x-direction)
a_y	Acceleration (y-direction)	Y	Aerodynamic Forces (y-direction)
a_z	Acceleration (z-direction)	Z	Aerodynamic Forces (z-direction)
N	Aircraft position (North)	L	Aerodynamic Moment (x-axis)
E	Aircraft position (East)	M	Aerodynamic Moment (y-axis)
D	Aircraft position (Down)	N	Aerodynamic Moment (z-axis)
α	Angle-of-Attack	δ_{th}	Thrust
β	Sideslip angle	δ_{f_1}	Left control surface deflection
p	Roll rate	δ_{f_2}	Right control surface deflection
q	Pitch rate	δ_r	Rudder deflection
r	Yaw rate	δ_e	Elevator deflection
V_t	True Airspeed	δ_a	Aileron deflection
ϕ	Roll angle	b_{ref}	Aircraft wingspan
θ	Pitch angle	c_{ref}	Mean aerodynamic chord
ψ	Heading angle	S	Wing surface area
u	Aircraft velocity (x-direction)	W	Aircraft gross weight
v	Aircraft velocity (y-direction)	I_{xx}	Moment of Inertia (x-axis)
w	Aircraft velocity (z-direction)	I_{yy}	Moment of Inertia (y-axis)
h	Altitude	I_{zz}	Moment of Inertia (z-axis)
M	Mach number	I_{xz}	Product of Inertia (xz-plane)

states that describe the longitudinal and lateral dynamics. Each of these subsystems is described in more detail below.

2.1.1 Conservation Equations

With the notations as given in Table 2.1, the standard conservation equations can be derived as shown in Equation (2.1).

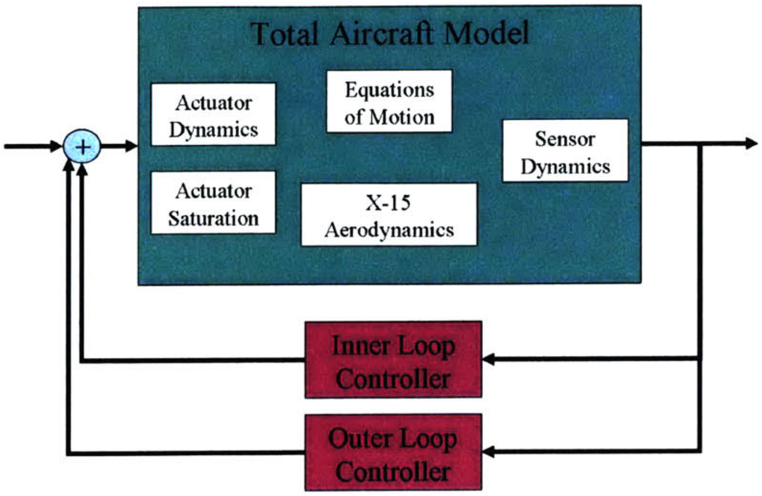
Total Aircraft Model
+
Actuator Dynamics
Equations of Motion
Sensor Dynamics
Actuator Saturation
X-15 Aerodynamics
Inner Loop Controller
Outer Loop Controller

Figure 2-1: The 6DOF X–15 Aircraft Model.

$$
\begin{aligned}
&X - mg\sin\theta = m(\dot{u} + qw - ru)\\
&Y + mg\cos\theta\sin\phi = m(\dot{v} + ru - pw)\\
&Z + mg\cos\theta\cos\phi = m(\dot{u} + qw - ru)\\[2ex]
&L = I_{xz}\dot{p} - I_{xz}\dot{r} + qr(I_{zz} - I_{yy}) - I_{xz}pq\\
&M = I_{yy}\dot{q} + rq(I_{xx} - I_{zz}) + I_{xz}(p^2 - r^2)\\
&N = -I_{xz}\dot{p} + I_{zz}\dot{r} + pq(I_{yy} - I_{xx}) + I_{xz}qr\\[2ex]
&p = \dot{\phi} - \dot{\psi}\sin\theta\\
&q = \dot{\theta}\cos\phi + \dot{\psi}\cos\theta\sin\phi\\
&r = \dot{\psi}\cos\theta\cos\phi - \dot{\theta}\sin\phi\\[2ex]
&\dot{\theta} = q\cos\phi - r\sin\phi\\
&\dot{\phi} = p + q\sin\phi\tan\theta + r\cos\phi\tan\theta\\
&\dot{\psi} = (qsin\phi + r\cos\phi)\sec\theta
\end{aligned}
\tag{2.1}
$$

2.1.2 Aerodynamics

It is well known[11] that the aerodynamic forces and moments acting on the aircraft can be expressed in terms of the nondimensional force and moment coefficients through multiplication by a dimensionalizing factor and, in the case of the forces, a

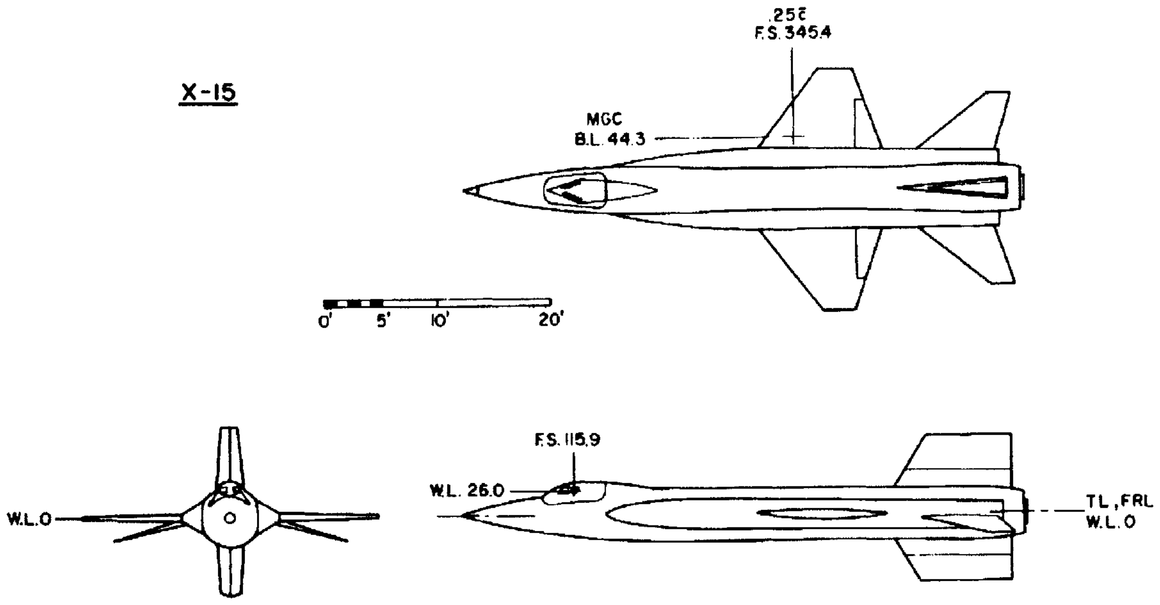

Figure 2-2: Planform for the X–15 research airplane.[5]

transformation from wind to body axes as described in Equation (2.2) below.

$$\begin{aligned}
\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} &= \bar{q}S \begin{bmatrix} \cos\alpha & 0 & -\sin\alpha \\ 0 & 1 & 0 \\ \sin\alpha & 0 & \cos\alpha \end{bmatrix} \begin{bmatrix} -C_D \\ C_Y \\ -C_L \end{bmatrix} \\
\begin{bmatrix} L \\ M \\ N \end{bmatrix} &= \bar{q}S \begin{bmatrix} b_{ref}C_l \\ c_{ref}C_m \\ b_{ref}C_n \end{bmatrix} ,
\end{aligned} \tag{2.2}$$

where C_L, C_D, and C_Y are the lift, drag, and side-force coefficient respectively and C_l, C_m, and C_n are the moment coefficients.

These force and moment coefficients in turn are functions of the aircraft states as well as the control inputs. 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 X–15 is also equipped with speed brakes which extended from the upper section of the rudder.

Due to their position on the aircraft, the speed brakes not only serve to increase drag, but they also add a positive pitching moment. Although they are included in the model, the speed brake deflection is not adjusted by the control system in this simulation.

It can be seen that the pitch/roll control surface deflections can be transformed into aileron and elevator deflections as

$$\begin{aligned} \delta_e &= \frac{\delta_{f_1} + \delta_{f_2}}{2}, \\ \delta_a &= \frac{\delta_{f_1} - \delta_{f_2}}{2}. \end{aligned} \tag{2.3}$$

For the purposes of the control discussion, it will be useful to speak in terms of these virtual elevators and ailerons as opposed to the actual combined pitch/roll control surface deflections. These control inputs are perhaps more similar to those of a standard subsonic or supersonic aircraft and unlike those of more recent hypersonic aircraft models [4, 1], where inputs such as the temperature change across the combustor, and the diffuser area ratio are used instead.

With the above control inputs, the force and moment coefficients are then broken down into their components in the following form:

$$\begin{aligned} C_L &= C_{L_{wing-body}} + C_{L_{\delta_e}}\delta_e \\ C_D &= C_{D_{wing-body}} + C_{D_{\delta_e}}\delta_e + C_{D_{\delta_{SB}}}\delta_{SB} \\ C_Y &= C_{Y_\beta}\beta + C_{Y_p}p\frac{b_{ref}}{2V_t} + (C_{Y_r} - C_{Y_{\dot\beta}})(r - \dot\beta)\frac{b_{ref}}{2V_t} + C_{Y_{\delta_a}}\delta_a + C_{Y_{\delta_r}}\delta_r \\ C_l &= C_{l_\beta}\beta + C_{l_p}p\frac{b_{ref}}{2V_t} + (C_{l_r} - C_{l_{\dot\beta}})(r - \dot\beta)\frac{b_{ref}}{2V_t} + C_{l_{\delta_a}}\delta_a + C_{l_{\delta_r}}\delta_r \\ C_m &= C_{m_{wing-body}} + (C_{m_q} + C_{m_{\dot\alpha}})(q - \dot\alpha)\frac{c_{ref}}{2V_t} + C_{m_{\delta_e}}\delta_e + C_{m_{\delta_{SB}}}\delta_{SB} \\ C_n &= C_{n_\beta}\beta + C_{n_p}p\frac{b_{ref}}{2V_t} + (C_{n_r} - C_{n_{\dot\beta}})(r - \dot\beta)\frac{b_{ref}}{2V_t} + C_{n_{\delta_a}}\delta_a + C_{n_{\delta_r}}\delta_r \end{aligned} \tag{2.4}$$

With the exception of $C_{L_{wing-body}}$ and $C_{D_{wing-body}}$, which are simply the contributions of the wing and body to the lift and drag respectively, each of the coefficients in

Equation (2.4) are non-dimensional derivatives, where C_{y_x} denotes the derivative of C_y with respect to x. It should be noted that as the X–15 moves through its flight envelope, these coefficients vary substantially with the angle-of-attack α as well as with the Mach number M. Equations (2.1)–(2.4) completely describe the open-loop dynamics of the X–15.

2.1.3 Actuators and Sensors

The control input deflections on the X–15 aircraft were actuated by irreversible hydraulic systems. The dynamics of these actuators can be modeled as second order systems with transfer functions

$$G_a(s) = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n + \omega_n^2} \tag{2.5}$$

where the damping ratio $\zeta = 0.7$ and natural frequency $\omega_n = 90$ Hz for the elevators/ailerons and $\omega_n = 70$ 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.

2.2 The MH–96 Adaptive Controller

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[7]. This was because in order to achieve high performance throughout the flight envelope, it was observed that rapid changes in the forward-loop gain would be required, and that the gains needed to be near their critical values at all times, thereby necessitating an adaptive control design which constantly adjusted these gains. To accomplish this, the system output was monitored in the frequency range at which instability occurs. 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[7]. The pilot ratings for the X–15–3, which was the only one equipped with the MH–96 adaptive controller, were generally higher than those of the X–15–1 and 2, which both employed a standard fixed gain stability augmentation system (SAS). This was especially true during re-entry, when changes in the dynamics were most dramatic[16]. 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 these reasons and more the MH–96 was flown extensively with great success.

2.2.1 Overall control architecture

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 the MH–96 adaptive controller[16, 18] and it has a more complex structure, with three individual loops for each of the pitch, roll, and yaw axes. Figure 2-3**(a)** shows the block diagram of the pitch axis of the MH–96[18] which is a slightly simplified representation of the MH–96 that was equipped on the X–15–3. Figure 2-3**(b)** presents the block diagram of the controller that was used in our studies.

One of the simplest ways of controlling this aircraft through feedback is that which was implemented on the X–15–1 and 2, the stability augmentation system (SAS). The pilot's control inputs were augmented by the SAS signals, the angular rates p, q, and r fed back through a series of fixed gains. This essentially added damping to the X–15, increasing its stability and improving its handling properties. The MH–96 controller, on the other hand, augmented the pilot's inputs with a more complex 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

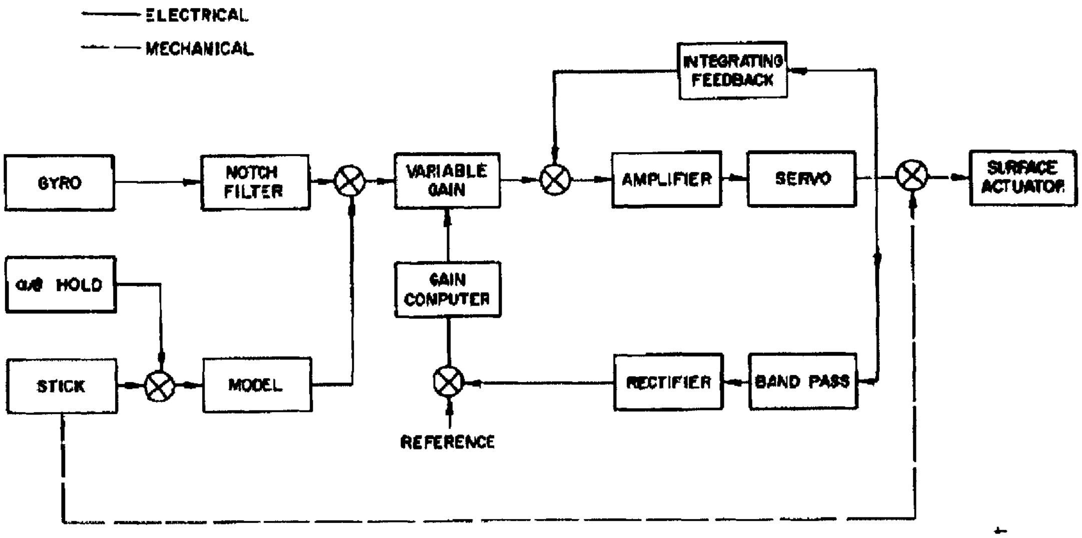
ELECTRICAL
MECHANICAL
INTEGRATING FEEDBACK
GYRO
NOTCH FILTER
VARIABLE GAIN
AMPLIFIER
SERVO
SURFACE ACTUATOR
α/θ HOLD
GAIN COMPUTER
STICK
MODEL
RECTIFIER
BAND PASS
REFERENCE

(a) Literature

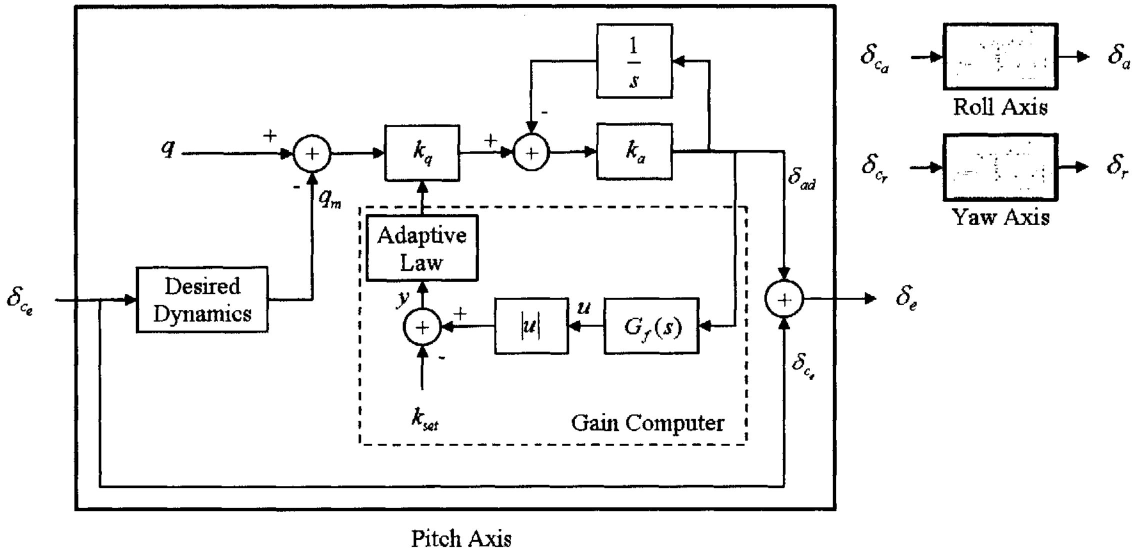
$\frac{1}{s}$
q
+
-
q_m
k_q
+
k_a
δ_{ad}
Adaptive Law
Desired Dynamics
δ_{c_e}
y
+
-
$|u|$
u
$G_f(s)$
k_{sat}
Gain Computer
δ_e
δ_{c_e}
Pitch Axis
δ_{c_a}
δ_a
Roll Axis
δ_{c_r}
δ_r
Yaw Axis

(b) Simulation

Figure 2-3: Schematic of the inner loop control architecture. These block diagrams represent the pitch axis control used in (a) the literature and (b) the simulation. The control loops for roll and yaw have essentially the same diagram[18].

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.

2.2.2 Inner-loop controller

The MH–96's model of the X–15, or the desired dynamics of the X–15 in the eyes of the controller, is given by

$$G_m(s) = \frac{1}{\tau_m s + 1} \tag{2.6}$$

where $\tau_m = 0.5$ in the pitch axis and $\tau_m = 0.33$ in the roll axis[16]. For the yaw axis, $G_m(s)$ is taken to be 0, that is, the desired yaw rate $r_m \equiv 0$. 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 designed as:

$$G_f(s) = \frac{s}{s^2 + 0.1s + \pi^2} \tag{2.7}$$

The resulting signal is then rectified and compared to a constant set-point k_{set}. The gain computer input y is thus

$$y = |G_f(s)\delta_{ad}| - k_{set} \tag{2.8}$$

where k_{set} represents the threshold between acceptable and unacceptable oscillation amplitude.

The algorithm for adjusting k_q was determined using the following guidelines: (i) the change in k_q had to be smooth, (ii) the amplitude of k_q had to lie within certain bounds k_{q1} and k_{q2} which ensure that structural feedback is minimized and are specified in the literature [16, 18], and (iii) the rate of change of k_q must be within certain bounds k_{qd1} and k_{qd2} which are chosen such that k_q can be rapidly reduced from large gain values that may trigger instabilities and more slowly increased so that the gains stay near critical values for longer periods of time. The above considerations led the following adaptive law for k_q:

$$\dot{k}_{q_{int}} = \begin{cases} k_{d_1} & \text{if } k_{q_0}y \leq k_{qd_1}, \\ k_{qd_2} & \text{if } k_{q_0}y \geq k_{d_2}, \\ k_{q_0}y & \text{otherwise.} \end{cases} \quad \text{and } k_q = \begin{cases} k_{q_1} & \text{if } k_{q_{int}} \leq k_{q_1}, \\ k_{q_2} & \text{if } k_{q_{int}} \geq k_{q_2}, \\ k_{q_{int}} & \text{otherwise.} \end{cases} \tag{2.9}$$

Equation (2.9) ensures that if the band-passed signal y is smaller than the set-point, the gain computer increases the forward-loop gain. Conversely, when the signal becomes large, signalling the onset of instability, the forward-loop gain is decreased. Typical time profiles of these variable gains are displayed in Figures 2-4(**a**) and 2-4(**b**), which show the original time profiles of the MH–96 adaptive controller[16] and of the simulated controller, respectively.

It should be noted that the notch filter used in the X–15 is not included in our studies as its original purpose was to reduce structural modes, which are not present in our rigid body model. The autopilot action represented by the α/θ hold block in Figure 2-3(**a**) was also not modeled in our simulation. Lastly, the conversion from electrical to mechanical signals provided by the amplifier and servo in Figure 2-3(**a**) is modeled simply as a gain k_a in Figure 2-3(**b**).

2.2.3 Outer-loop controller

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

$$u_{cmd}(t) = [V_{cmd}(t) \quad h_{cmd}(t)]^T \tag{2.10}$$

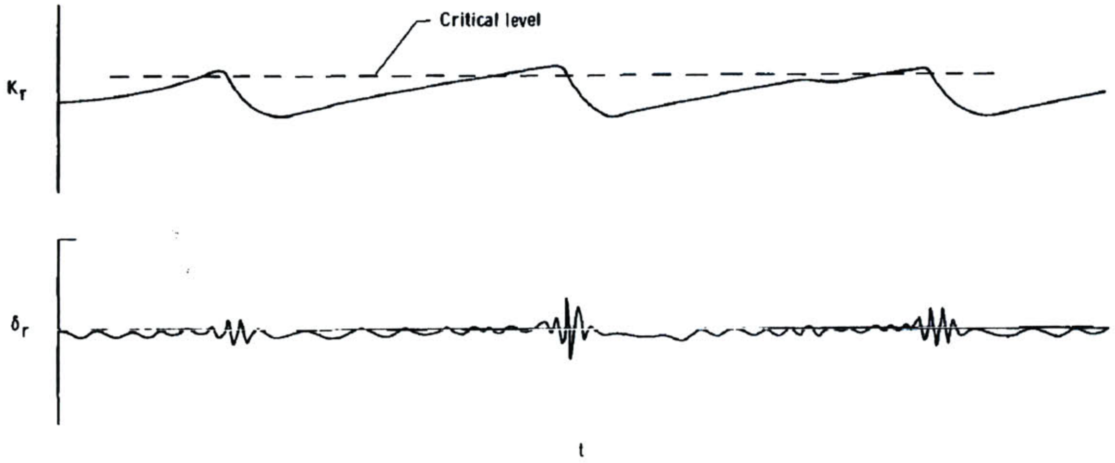

(a) Literature

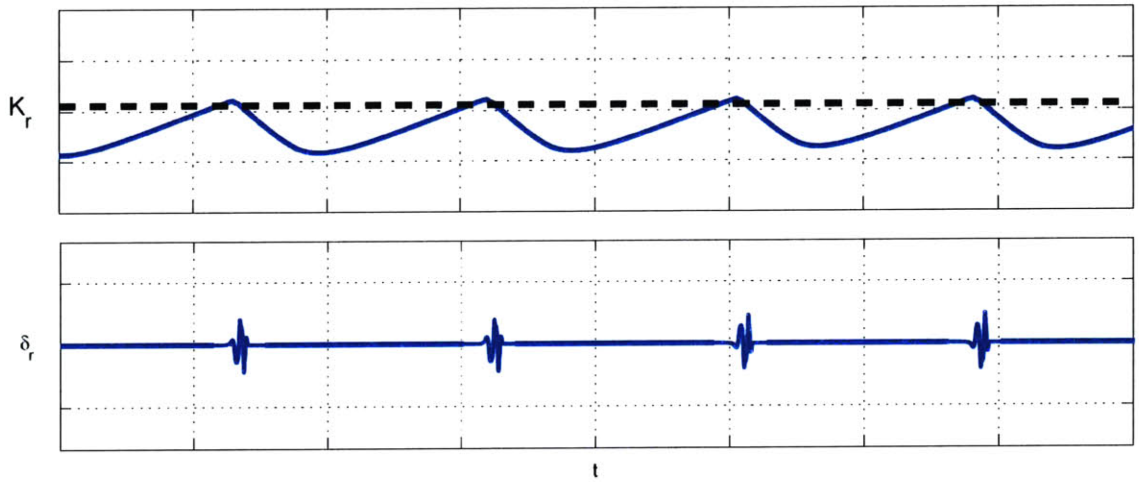

(b) Simulation

Figure 2-4: Gain Changer Operation. These plots display the typical gain changer performance and control surface activity for **(a)** the original MH–96 adaptive controller[16] and **(b)** the controller implemented for simulation.

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 $\delta_c = [\delta_{th} \quad \delta_e]^T$. The controller is a PID (with approximate derivative) control on the error between the actual and the commanded values, with a transfer function given by

$$G_c(s) = K_P + \frac{K_I}{s} + \frac{K_D s}{\frac{1}{N}s + 1}, \tag{2.11}$$

with $K_P = \begin{bmatrix} K_{P_V} & 0 \\ 0 & K_{P_h} \end{bmatrix}$, $K_I = \begin{bmatrix} K_{I_V} & 0 \\ 0 & K_{I_h} \end{bmatrix}$, and $K_D = \begin{bmatrix} K_{D_V} & 0 \\ 0 & K_{D_h} \end{bmatrix}$,

The gains for both the speed and altitude loops are selected using the Ziegler-Nichols ultimate sensitivity method[23].

2.3 Modern 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. The overall block diagram of the modern adaptive controller is shown in Figure 2-6, 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. The inner-loop controller described in Section 2.3.1 below is offered as a modern replacement to the MH–96 controller discussed in Section 2.2.2.

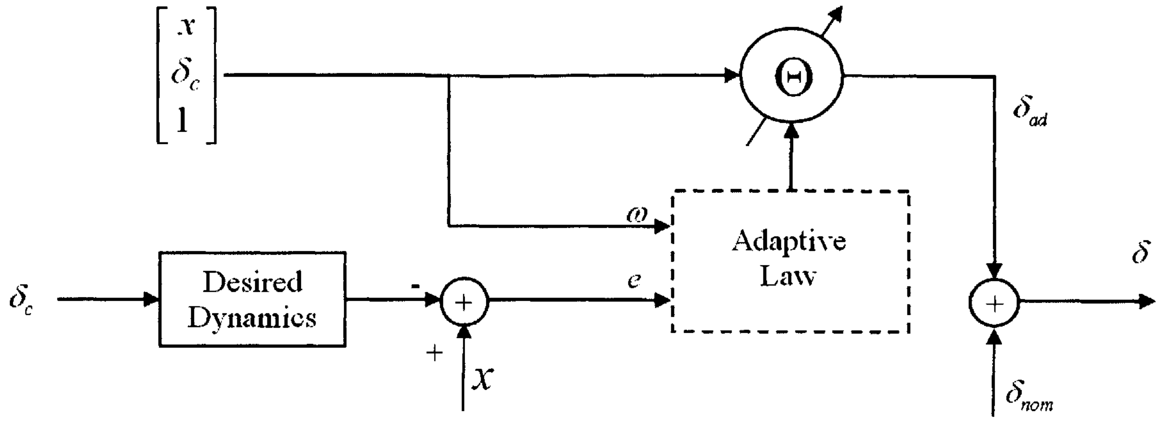

Figure 2-5: Block diagram of the modern adaptive inner-loop controller. Note that pitch, roll, and yaw axes are combined so that $\delta_c = [\delta_{c_e} \quad \delta_{c_a} \quad \delta_{c_r}]^T$ and $\delta = [\delta_e \quad \delta_a \quad \delta_r]^T$.

2.3.1 The Modern Inner-loop Adaptive Controller Design

In contrast to the MH–96 inner-loop controller shown in Figure 2-3, the modern 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.

Baseline Controller

One of the key elements of our proposed controller is that a good baseline controller is augmented with adaptive control. 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 given in Equations (2.1)–(2.4), which can be written as

$$\dot{X}_t = f_p(X_t, U_t), \tag{2.12}$$

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

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

$$\dot{x}_p = A_{p_i} x_p + B_{p_i} \delta, \tag{2.13}$$

where

$$A_{p_i} = \left. \frac{\partial \dot{X}}{\partial X} \right|_{X_{0_i},\ U_{0_i}} \qquad B_{p_i} = \left. \frac{\partial \dot{X}}{\partial U} \right|_{X_{0_i},\ U_{0_i}}, \tag{2.14}$$

where $x_p = X - X_{0_i} \in \Re^n$, $\delta = U - U_{0_i} \in \Re^m$. The baseline controller is then designed straightforwardly using the parameters given by Equation (2.14).

To overcome the drift in the lateral dynamics due to the trim disturbance, an integral controller was added in the roll rate p and the combined yaw rate/sideslip angle term $r - \beta$,

$$x_c = [p_I \quad (r - \beta)_I]. \tag{2.15}$$

The decision to combine r and β was made to reduce the number of states by exploiting their strong coupling. Note that integral states for the longitudinal terms α and q are not included since deviations from the trim values of these states are used heavily by the outer loop controller to control altitude and speed. We can write the dynamics of these integral controller states as

$$\dot{x}_c = A_c x_c + B_c x_p. \tag{2.16}$$

The nominal baseline LQ controller is then designed in the standard feedback form as

$$\delta_{nom} = K_{x_i} x, \tag{2.17}$$

where $x = [x_p^T \quad x_c^T]^T$, and K_{x_i} denotes the nominal feedback gain matrix designed for the dynamics given by Equations (2.13) and (2.15) around the i^{th} trim point, and minimizes the cost function

$$J = \int_0^\infty (x^T Q_{lqr} x + \delta_{nom}^T R_{lqr} \delta_{nom}) dt. \tag{2.18}$$

A schedule of ten nominal LQ gain matrices was thus constructed for the baseline controller.

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, one that includes a parametric uncertainty matrix Λ, which is discussed in more detail in Section 3.1, and another that includes a saturation nonlinearity in the actuator. Both these effects are incorporated in the linearized dynamics (2.13) as

$$\dot{x}_p = A_{p_i} x_p + B_{p_i} \Lambda \text{sat}(\delta) + d_{p_i}, \tag{2.19}$$

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

$$\text{sat}(\delta) = \begin{cases} \delta & \text{if } |\delta| \leq \delta_{max} \\ \delta_{max} \text{sgn}(\delta) & \text{if } |\delta| > \delta_{max} \end{cases}, \tag{2.20}$$

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

$$\begin{bmatrix} \dot{x}_p \\ \dot{x}_c \end{bmatrix} = \begin{bmatrix} A_{p_i} & 0 \\ B_c & A_c \end{bmatrix} \begin{bmatrix} x_p \\ x_c \end{bmatrix} + \begin{bmatrix} B_{p_i} \\ 0 \end{bmatrix} \Lambda \text{sat}(\delta) + \begin{bmatrix} d_{p_i} \\ 0 \end{bmatrix}, \tag{2.21}$$

or equivalently

$$\dot{x} = A_i x + B_i \Lambda u + d_i, \tag{2.22}$$

where $u = \text{sat}(\delta)$. The overall dynamics given by Equation (2.22) is used for the adaptive control design.

In order to ensure safe adaptation, a target dynamics is specified for the adaptive controller using a reference model. This is designed using the baseline controller and the plant dynamics in the presence of no actuator uncertainties. That is, the desired goal for the adaptive controller is to generate the same performance that would have been obtained had there been no actuator uncertainties. This is given by

$$\dot{x}_{ref} = (A_i + B_i K_{x_i})x_{ref} = A_{ref_i} x_{ref}, \tag{2.23}$$

where all the matrices A_{ref_i} are assumed to be Hurwitz.

Using Equations (2.22) and (2.23), an adaptive control input is added to the baseline controller as

$$\delta_{ad} = \theta_x^T(t)x + \theta_\delta(t)\delta_c + \theta_d(t) = \Theta^T \omega, \tag{2.24}$$

where $\Theta^T = [\theta_x^T \quad \theta_\delta^T \quad \theta_d^T]$ are adaptive parameters that will be adjusted in the adaptive law given in Equation (2.25) below and the linear regressor $\omega = [x \quad \delta_c \quad 1]^T$.

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 Reference [8].

The adaptive law given by Reference [9],

$$\dot{\theta} = -\Gamma \omega e^T P B_p sign(\Lambda) - \theta \left(1 - \frac{||\theta||}{\theta^*_{max}}\right)^2 f(\theta) \tag{2.25}$$

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

$$f(\theta) = \begin{cases} 1 & \text{if } ||\theta|| \geq \theta^*_{max} \\ 0 & \text{otherwise,} \end{cases} \tag{2.26}$$

ensures gauranteed stability while providing an upper bound on the adaptive parameters Θ.

The adaptive gains at the i^{th} trim point Γ_i are selected according to the following empirical formula, which arises from inspection of the structure of the adaptive laws[3]:

$$\Gamma_i = |\frac{\text{diag}(\vartheta_i)}{\tau_{min} p_i \delta_{c_{max}}}| + \Gamma_0 \tag{2.27}$$

where

i. $\vartheta_i \in \Re^n$ is vector given by the sum of the columns of $\bar{\Theta}_i^*$ where $\bar{\Theta}_i^*$ corresponds to the uncertainty $\Lambda = \bar{\Lambda}$ for which the plant has the most unstable eigenvalues. The components of $\bar{\Theta}_i^*$ are given by

$$\begin{aligned} \bar{\theta}_{x_i}^* &= -K_{x_i}\bar{\Lambda}^{-1}, \\ \bar{\theta}_{\delta_i}^* &= \bar{\Lambda}^{-1}, \\ \bar{\theta}_{d_i}^* &= [1 \quad 1 \quad 1]^T. \end{aligned} \tag{2.28}$$

ii. τ_{min} is the smallest time constant of the reference model,

iii. p_i is the norm of $B_{ref_i}^T P_i$,

iv. $\delta_{c_{max}}$ is the maximum amplitude of the reference input signal,

v. Γ_0 is a small positive definite diagonal matrix which ensures that Γ is positive definite.

Thus the full inner-loop control input becomes

$$\delta = \delta_{nom} + \delta_{ad}, \tag{2.29}$$

and the overall control architecture can be seen in Figure 2-6.

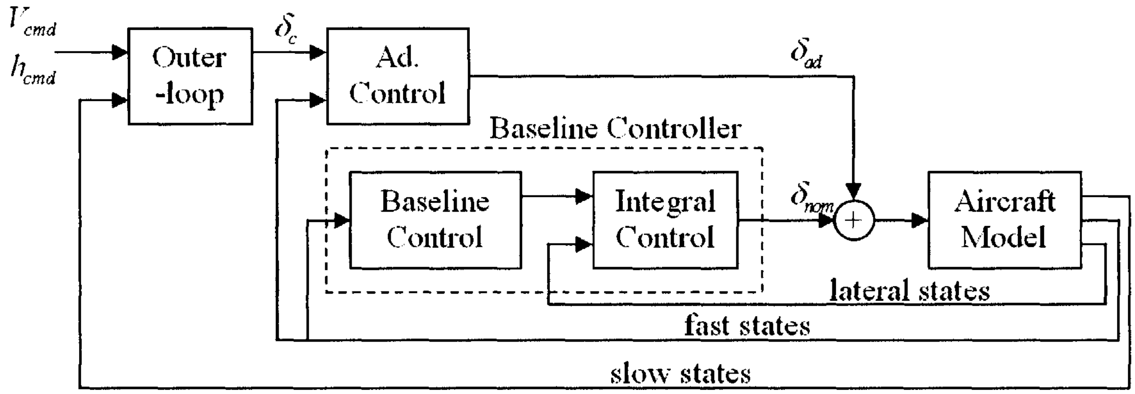

V_{cmd}
h_{cmd}
Outer -loop
δ_c
Ad. Control
δ_{ad}
Baseline Controller
Baseline Control
Integral Control
δ_{nom}
+
Aircraft Model
lateral states
fast states
slow states

Figure 2-6: Overall control structure for the modern adaptive control system. The baseline controller (composed of the Baseline Control and Integral control blocks) is augmented by the adaptive control.

Chapter 3

Margins

Traditional control design for aircraft typically relies heavily on the application of classical linear design tools, especially the notions of gain and phase margin. Since these techniques are not applicable to adaptive controllers, a new set of tools must be developed in order to certify margins for adaptive controllers. Setting aside the question of what happened to the X–15 in 1967 for a moment, we can now examine the robustness margins of the modern adaptive controller at a single operating point. In order to calculate these margins, we simulate the modern adaptive controller as described in Section 2.3 using the aircraft model discussed in Section 2.1 in the presence of several types of uncertainties. Through examination of the stability of this closed-loop system, we compute the margins numerically.

3.1 Parametric Uncertainties

The first type of uncertainty comes in the form of a control failure. In the case of a control failure, the nonlinear dynamics (2.12) can be represented as

$$\dot{X} = F(X, \Lambda U), \tag{3.1}$$

where $\Lambda^{m \times m}$ is an unknown diagonal matrix with nonzero diagonal entries λ_i which represents loss of control effectiveness or control reversal. This control failure parameterization differs in several ways from that proposed by References [2, 17]. First, as opposed to assuming the aircraft is equipped with many redundant actuators which either fail completely (corresponding to some $\lambda_i = 0$) or do not fail at all ($\lambda_i = 1$), we allow λ_i to take on any value between -1 and 1 and typically only require $\lambda_i \neq 0$ for controllability. This corresponds to the scenario where a portion of the control surface is damaged or separated so that the overall effectiveness of the control surface is reduced to a fraction of its nominal value, and that fraction is in fact λ_i. Second, we do not consider the effect of control surfaces becoming "stuck" in a certain position. It is assumed that the failed portion of the control surface does not produce any aerodynamic forces or moments from the time when the failure occurs onward. An example of this type of loss of controller effectiveness is shown in Figure 3-1 below.

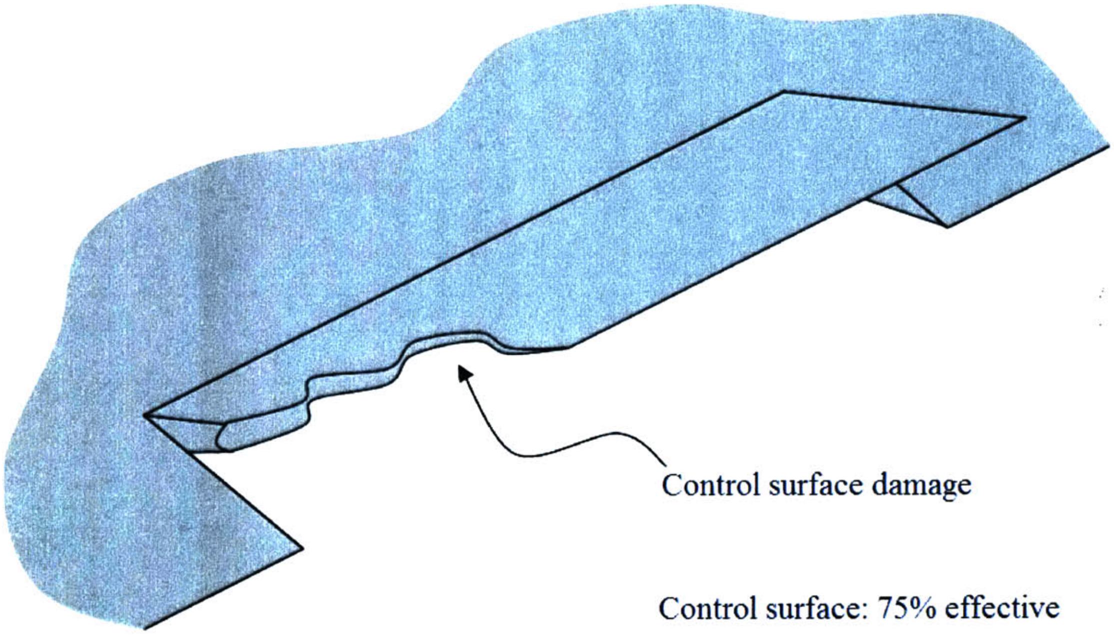

Figure 3-1: Example of an aircraft control surface that has sustained some type of damage, resulting in a 75% loss of effectiveness.

If the nonlinear plant in (3.1) is linearized about a single trim point (X_0, U_0), we obtain a linearized dynamics

$$\dot{x}_p = A_{p\lambda} x_p + B_{p\lambda} \delta + d_p, \tag{3.2}$$

$$A_{p\lambda} = \left.\frac{\partial F(X, \Lambda U)}{\partial X}\right|_{X=X_0,\ U=U_0}, \quad B_{p\lambda} = \left.\frac{\partial F(X, \Lambda U)}{\partial U}\right|_{X=X_0,\ U=U_0}, \quad d_p = F(X_0, \Lambda U_0) \tag{3.3}$$

where $A_{p\lambda}$, $B_{p\lambda}$, and d_p are unknown, and d_p is a constant and bounded disturbance due to the fact that the aircraft is no longer trimmed properly. Assuming that $F(X, U)$ is affine in U, we obtain that $B_{p\lambda} = B_p\Lambda$ and $d_p = B_p\Lambda d$. Consequently, the linearized dynamics in (3.2) can be written as

$$\dot{x}_p = A_{p\lambda}x_p + B_p\Lambda(\delta + d). \tag{3.4}$$

We assume that $(A_{p\lambda}, B_p\Lambda)$ is controllable. Equation (3.4) therefore represents the aircraft dynamics in the presence of control failure and constant, bounded disturbances. Note that for the purpose of the margins calculation, actuator nonlinearity due to saturation is not included.

3.2 Unstructured Uncertainties

There are numerous unstructured uncertain dynamics which are not captured in the linearized dynamics (3.4), the first and most obvious of which are the nonlinearities which are ignored in the linearized model. The residues due to these nonlinearities can sometimes significantly affect the dynamics, especially when the system is driven at high frequencies. Second, structural modes such as the aeroelastic mode are not captured in many aircraft models, which assume rigid body dynamics. And finally, there will always be time delays as a result of CPU time at each step of the control process where calculations take place. With these uncertainties present, we can express the full dynamics as

$$\dot{x} = A_{p\lambda}x_p + B_p\Lambda(\delta + d) + g(x_p, \delta, \mu). \tag{3.5}$$

where $g(x_p, \delta, \mu)$ represents uncertainties due to nonlinearities, structural modes, or time delay that may be present in the system and $\mu \geq 0$ represents a measure of the severity of the particular uncertainty present, or in the presence of multiple uncertainties, a vector of these values.

3.3 Numerical Calculation of Margins

Using the high-performance aircraft model described in Equations (2.1)–(2.4) with the adaptive controller described in Equations (2.24) and (2.25), we examine the robustness of the overall closed-loop system. The simulation studies begin with the aircraft trimmed at an altitude of 60,000 ft and Mach number of 2.6. This flight condition corresponds to the maximum dynamic pressure, $\bar{q}$, typically experienced by this particular aircraft. An amount of uncertainty is then introduced and the aircraft is commanded to perform an aggressive maneuver simultaneously at $t = 20$ sec. The simulation is then allowed to run until either it completes successfully or one of two failure conditions occurs: 1) the simulation is unable to continue without reducing the step-size below .001 sec, 2) the aircraft altitude is reduced below 55,000 ft. The first condition represents a failure that is normally the result of high frequency oscillations in the adaptive controller which are not physically realizable on aircraft hardware. The second condition represents a dramatic departure from the trimmed condition which constitutes a performance objective failure. If neither of the first two conditions occur after 100 sec, the maneuver is deemed successful.

In addition to the nonlinearities inherently present in our full aircraft model, three types of uncertainty are considered: actuator failure, unmodeled dynamics, and time delay. For each of these types of uncertainty, the stability boundary of the adaptive controller is determined using the simulation environment described above over several values of the adaptive gain Γ. It is assumed that μ can be chosen such that $g(x_p, \delta, 0) = 0$. Intuition tells us that the system will remain stable for small values of μ and instability will first occur at some critical value μ^*. This critical μ^* can be considered as a margin of the adaptive system with respect to the particular uncertainty present. For each choice of Γ, the parameter μ is increased from 0 until the simulation fails. Backstepping is then used to determine the critical value μ^* to a specified tolerance. This approach computes the margins in a time-efficient manner.

First, parametric uncertainty is added in the form of Λ, the actuator failure matrix as in Equation (3.1). Figure 3-2 below shows the stability region for both the baseline and the adaptive controllers. For this particular example, the parameter μ is chosen as $\mu = 1 - \lambda_3$, corresponding to a failure of the aircraft's right control surface. The adaptive system is able to withstand far greater levels of parametric uncertainty than the nominal system alone. Consequently, the stability region for the adaptive controller is much larger than that of the nominal system. Note that in the case of a linear plant, the adaptive controller is proven to be stable for any $0 < \lambda_i \leq 1$, for any choice of Γ. It is clear from Figure 3-2 that the full nonlinear aircraft does not retain this quality. The overall closed loop system robustness to actuator failures is in fact a measure of the severity of the nonlinear residue present after linearization.

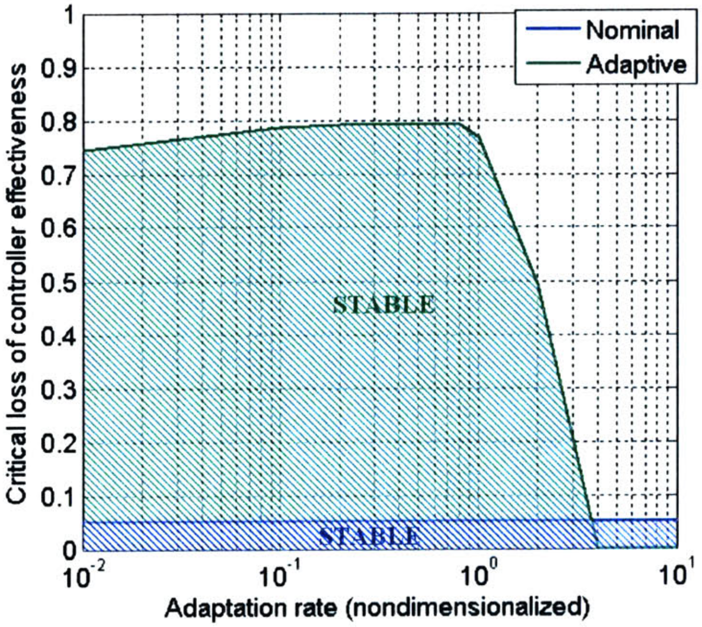

Figure 3-2: The stability region for both the adaptive and the nominal system in the presence of parametric uncertainty. The adaptation rate is normalized by the value prescribed by an empirical rule given by Equation (2.27). Note that the nominal stability region shows no dependance on the adaptation rate since the baseline controller includes no adaptation.

Second, we examine the robustness to unmodeled dynamics in the form of aeroelastic modes. Figure 3-3 shows the respective stability regions for the nominal baseline and adaptive controllers. Despite the adaptive system's lack of knowledge of the structure of these uncertainties, the adaptive controller outperforms the baseline controller alone for most choices of Γ. The aircraft model discussed thus far is a rigid-body model, but flexible aircraft modes can be added using the methods discussed in References [4] and [1]. Essentially the aircraft is modeled as two cantilever beams extending fore and aft from the center of mass. In this configuration, first mode vibrations effect the pitch rate q. Here μ represents the excitability of the aeroelastic mode which in practice will depend on numerous physical parameters of the aircraft. Note that the baseline controller is already quite robust to this type of disturbance since it is designed using LQ techniques.

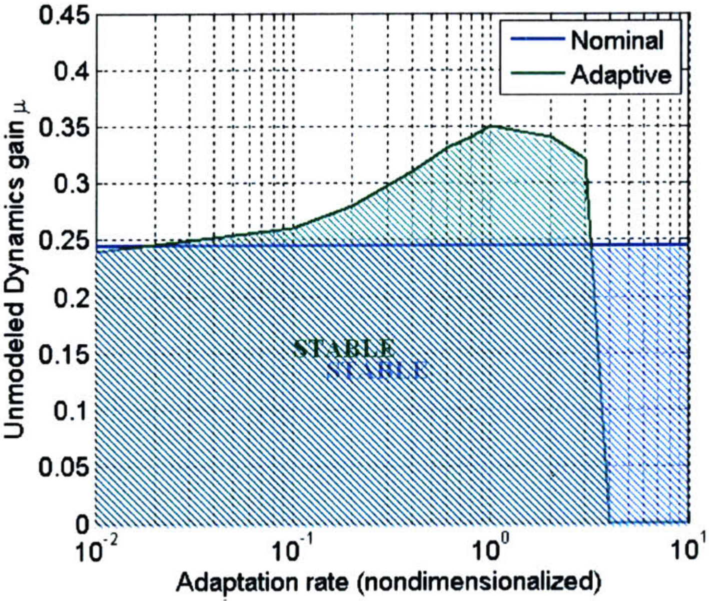

Figure 3-3: The stability region for both the adaptive and the nominal system in the presence of unmodeled dynamics in the form of aeroelastic modes.

Finally, time delay is added to the system at both the input and the output of the plant. Figure 3-4 shows the stability regions for both controllers in the presence of time delay. Time delay poses quite a challenge for adaptive control. The adaptive system attempts to counteract the effects of the time-delay by introducing successively higher frequencies into the system. However, it can not possibly match the frequency content of the delay, and unfortunately these high frequencies typically result in the simulation reaching failure condition 1. For this reason, high adaptation rates should be avoided in order to maximize robustness with respect to time delay. Here $\mu = \tau_1 = \tau_2$, the time delay in ms added to both the input and output of the plant.

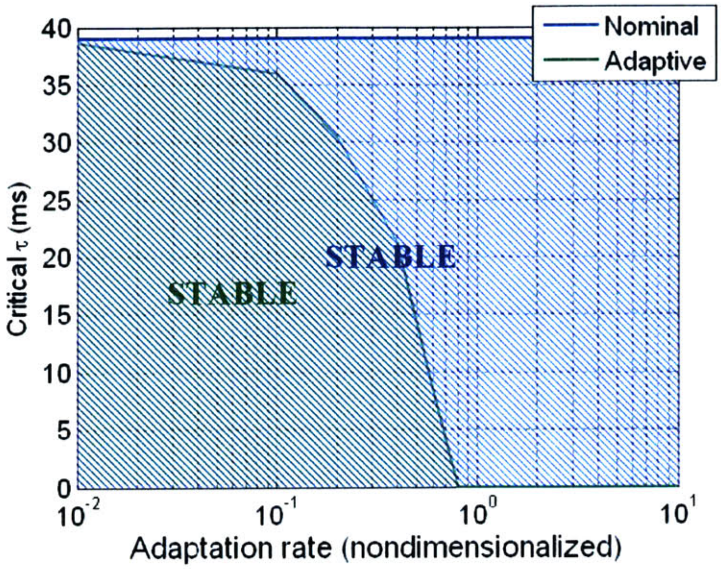

Figure 3-4: The stability region for both the adaptive and the nominal system in the presence of time delay at both the input and the output to the plant.

The critical μ^* provides a compact method of expressing the robustness of the adaptive controller, similar to the notions of gain and phase margin for linear systems. The computation of these margins can be carried out in a time-efficient manner for a number of uncertainties, including but not limited to those discussed above. This allows a complete picture of the robustness of the adaptive controller to be built in a relatively short period of time.

Chapter 4

X–15 Demonstration

With the reconstruction of the original X–15 aircraft, including both the aircraft dynamics as well as the MH–96 adaptive controller that was implemented in the flights, we now investigate the 1967 accident itself and a possible explanation for how the sequence of events transpired. We then show that, under the same conditions that led to the anomalous behavior and eventual crash of the original X–15, the modified X–15 (which includes the same aircraft dynamics coupled with a modern adaptive controller) is able to recover and complete the maneuver successfully.

4.1 The 1967 Incident

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[7]. 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.

Equations (2.1)–(2.4), that described the X–15 dynamics, the multi-loop control architecture described in Section 2.2.1, together with the adaptive algorithm in Equation (2.25) were simulated to represent the overall flight control system and the flight dynamics. Table 4.1 lists all of the parameter values that were used. In addition to these parameters, the X–15 aerodynamics block utilizes a lookup table for each of the nondimensional coefficients in Equation (2.4) with axes in M and α. The data for these tables was extracted from various sources: flight recorded data[22, 21, 20], wind tunnel measurements[20, 12, 19], and theoretical calculations[20, 19]. The pilot is assumed to engage the speed brakes during the final stages of the descent, that is

$$\delta_{SB} = \begin{cases} 1 & \text{if } 350 \leq t \leq 400, \\ 0 & \text{otherwise.} \end{cases} \tag{4.1}$$

The tracking performance obtained under nominal conditions is shown in Figure 4-1**(a)**–**(d)**, which shows that the altitude and speed track the commanded signals fairly closely. Figure 4-1 represents the nominal behavior of the X–15, corresponding to the nearly 200 successful flights in the program.

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 control surface at $t = 80$ s. As can be seen in Figure 4-2, the tracking performance begins to degrade soon after $t = 80$ s. While the extent of the degradation may not be immediately evident in Figure 4-2, Figures 4-3, 4-4, 4-5, and 4-6 illustrate the anomalous behavior of the X–15 more clearly. In Figure 4-3 we clearly see that the X–15 has made a dramatic departure from the commanded trajectory. The simulation is stopped at

Table 4.1: Simulation Parameter Values.

S	200 ft^2	k_a	1	k_{r_0}	-0.05
b_{ref}	22.36 ft	k_{d_1}	-0.5	k_{r_1}	0.1
c_{ref}	10.27 ft	k_{d_2}	0.2	k_{r_2}	1.0
W	15560 lbs	k_{p_0}	0.05	K_{P_V}	40.2
I_{xx}	3650 slug-ft^2	k_{p_1}	-0.1	K_{P_h}	720
I_{yy}	80000 slug-ft^2	k_{p_2}	-1.0	K_{I_V}	10
I_{zz}	82000 slug-ft^2	k_{q_0}	-0.1	K_{I_h}	2.18
I_{xz}	490 slug-ft^2	k_{q_1}	0.1	K_{D_V}	1030
k_{set}	0.2	k_{q_2}	1.0	K_{D_h}	2970

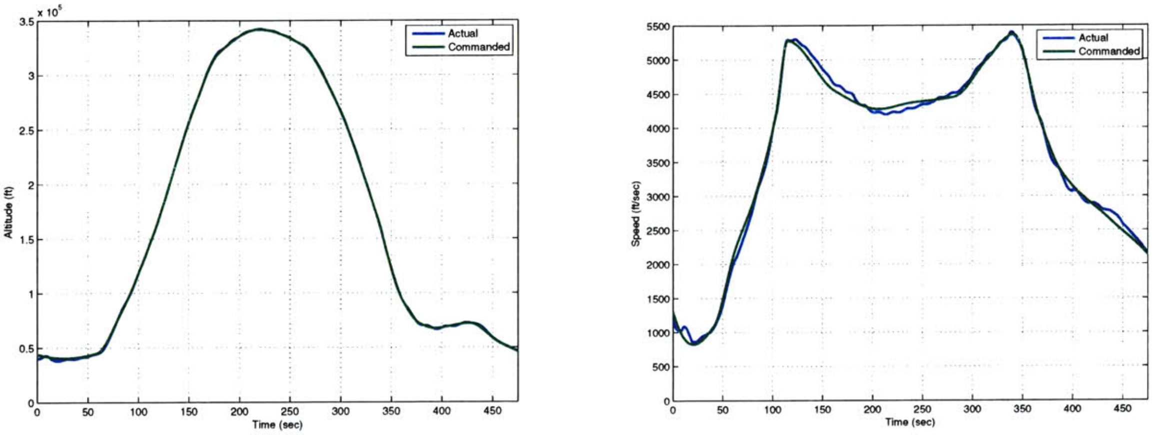

(a) Altitude **(b) Speed**

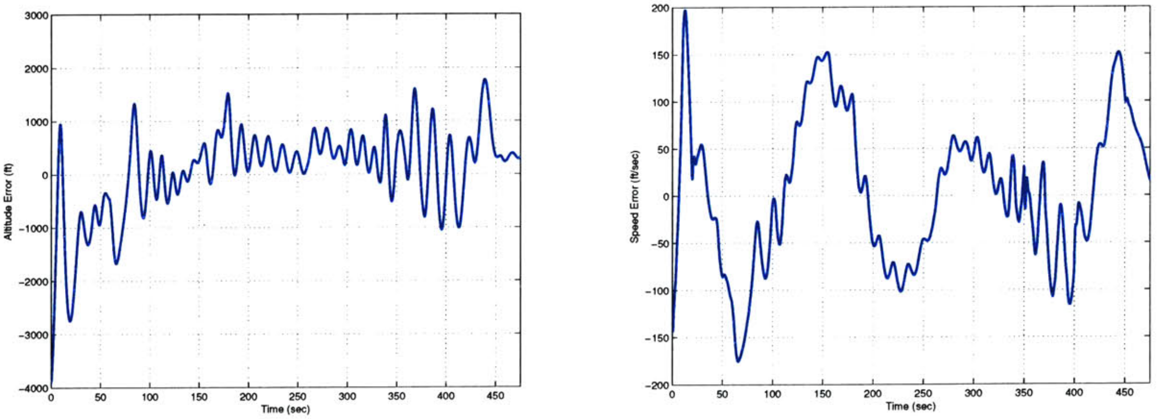

(c) Altitude Error **(d) Speed Error**

Figure 4-1: Tracking performance of the X–15 (nominal case). 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.

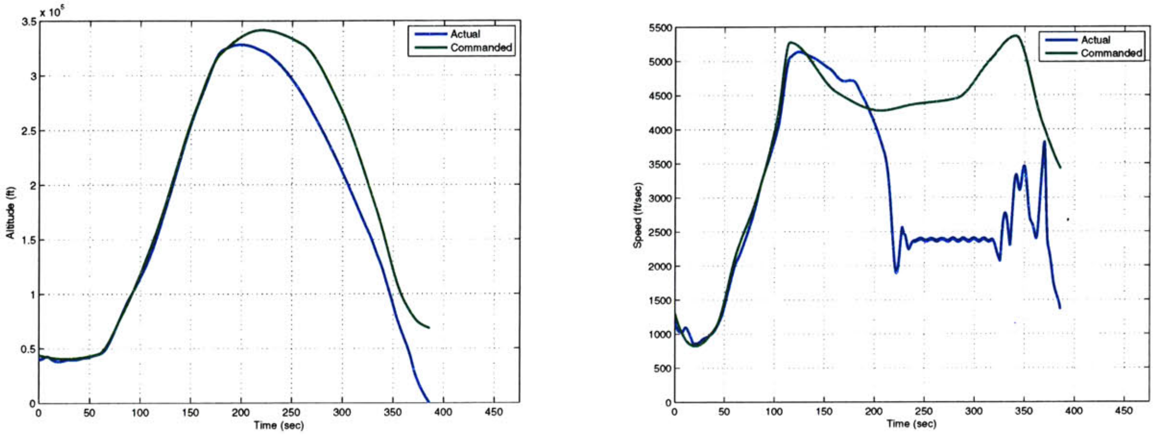

(a) Altitude **(b) Speed**

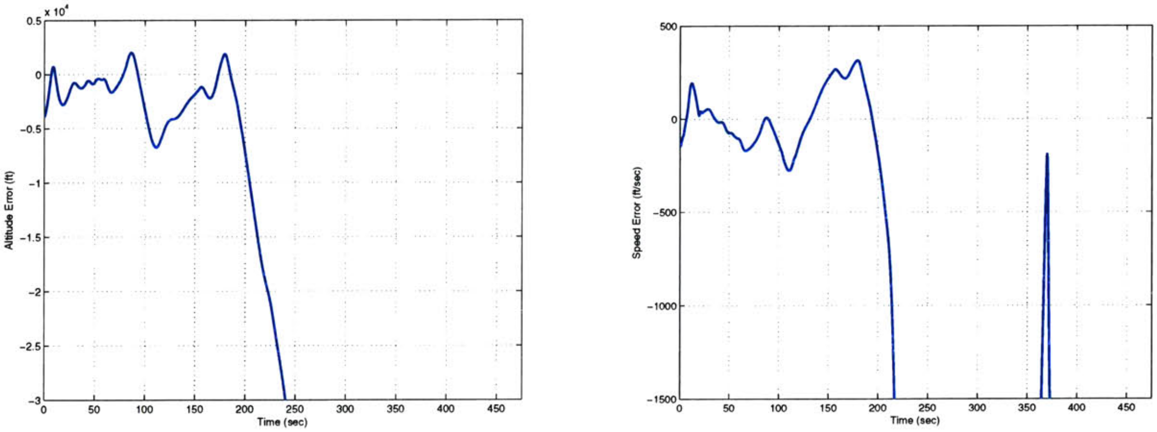

(c) Altitude Error **(d Speed Error)**

Figure 4-2: Tracking performance of the X–15 (failure case). Even after the failure at $t = 80$ s, the system continues to track fairly well until the dramatic departure from the commanded values around $t = 200$ s.

$t = 385$ 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. Closer examination of the aircraft states reveals additional similarities between the simulation and Events (1)–(9). The first of these is Event (3), a steady drift in heading angle ψ which can be attributed to the asymmetry of the control failure. In Figure 4-4(**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). This is followed by a rapid downward spiral (see Figure 4-3, $t = 200$ s and $t = 300$ s, Figure 4-4(**b**), $t = 180$ s and onward). 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 (see Figure 4-5). Lastly, we can see large accelerations in both the lateral and vertical directions (see Figure 4-6) that ultimately would have caused the X–15 to break apart, corresponding

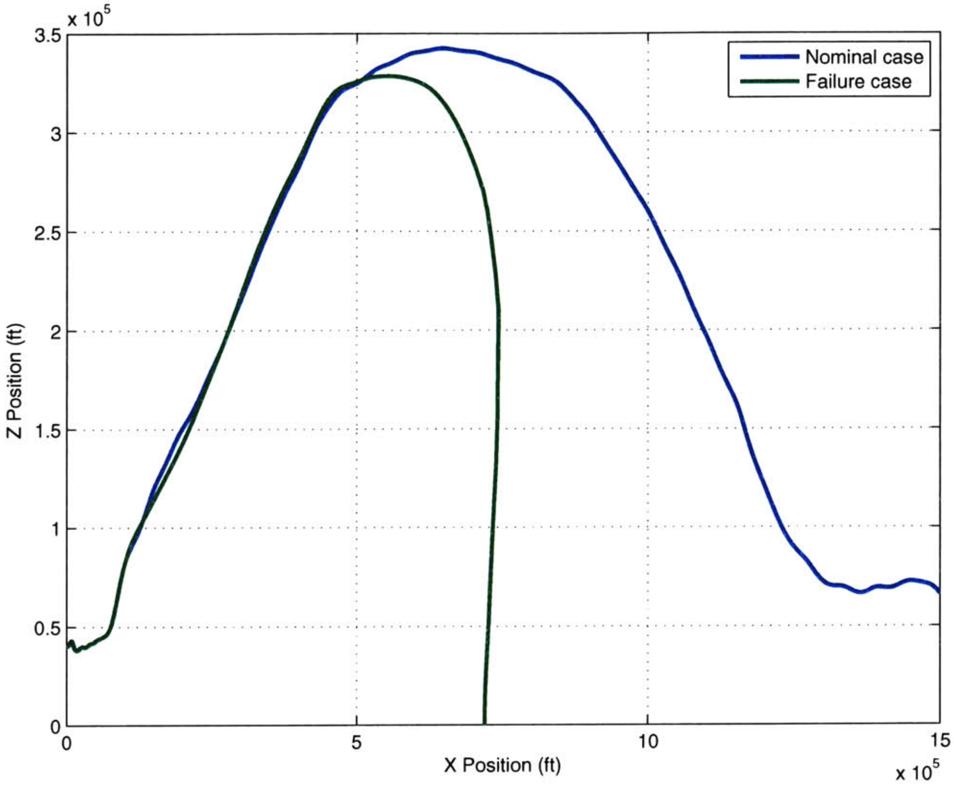

Figure 4-3: Flight Path of the X–15 (failure case)

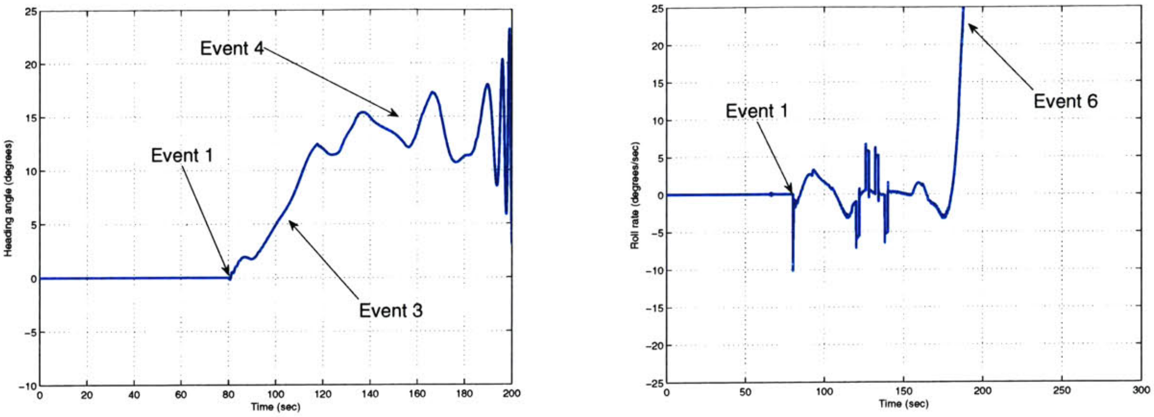

(a) Heading Angle **(b) Roll Rate**

Figure 4-4: 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.

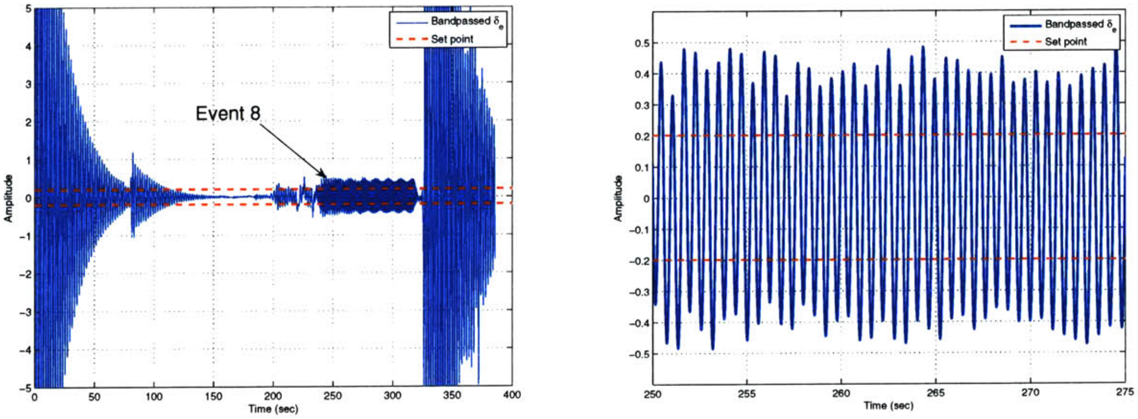

(a) Gain Computer Input **(b) Gain Computer Input (Zoom)**

Figure 4-5: Limit Cycles in the X–15 adaptive flight control system. **(a)** shows the input to the adaptive gain computer over time. Recall that the forward-loop gain is reduced when this input signal becomes larger than the set-point. **(b)** displays a blow up of the signal between $t = 250$ s and $t = 275$ s, showing the existence of undamped oscillations which prevent the gain changer from correctly adjusting the gain.

to Event (9).

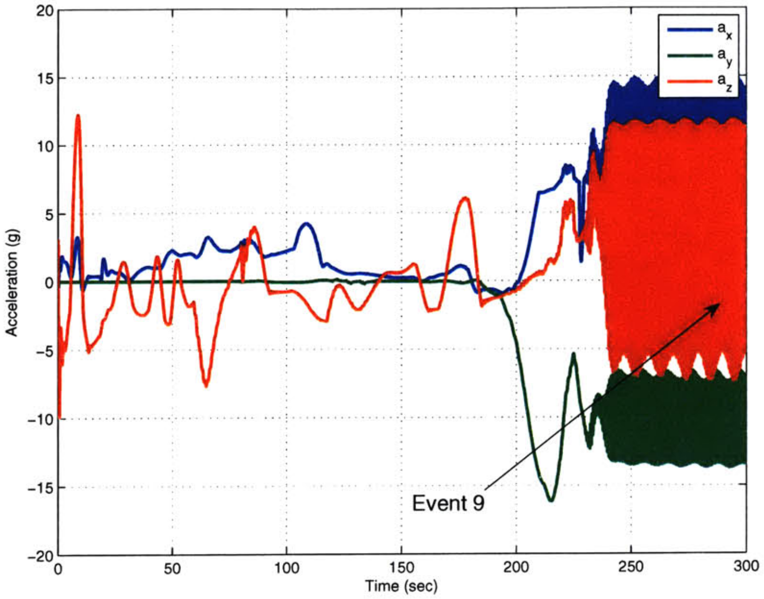

Figure 4-6: Accelerations experienced by the X–15 (failure case).

It is important to note that some of the events reported in [7] were not reproducible. For example, the planned wing-rocking maneuver in Event (2) was not excessive in our simulation. Additionally, in Figures 4-3 and 4-4, 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.

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.

4.2 The 2007 Approach

In this section we examine in detail the implementation of the modern adaptive controller described in Section 2.3 and its robustness to the same disturbance. The first step in this process 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. The trim points were distributed uniformly across this path as shown in Figure 4-7 below.

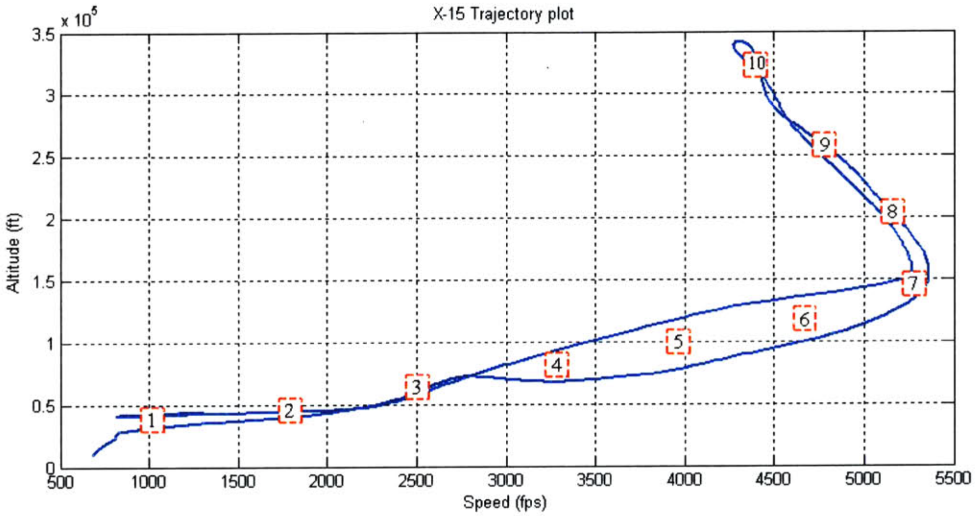

Figure 4-7: The commanded trajectory in Altitude/Speed space. This figure shows the commanded path for the X–15 simulations along with labels of the locations of the trim points used for controller design.

Table 4.2: Simulation Parameter Values (Adaptive).

Q_{lqr}	diag([10 10 100 5000 100 10 10])	Q	$10I^{7\text{x}7}$
R_{lqr}	$I^{3\text{x}3}$	$\bar{\Lambda}$	$.1I^{3\text{x}3}$
τ_{min}	0.33	Γ_λ	$I^{3\text{x}3}$
$\delta_{c_{max}}$	15°		

The next step is to simulate the modern adaptive inner-loop controller described in Equations (2.24)–(2.27) with the X–15 model and PID outer loop controller in Equation (2.11). Once again the speed brakes are engaged as described in Equation (4.1). The additional parameter values used for these simulations can be found in Table 4.2. 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 Equation (2.10) fairly well. Figure 4-8**(a)**–**(b)** show that $h \to h_{cmd}$ and $V \to V_{cmd}$. Figure 4-8**(c)**–**(d)** show that the errors are small, generally less than 1% of the maximum in the case of altitude and less than 3% of the maximum in the case of speed. This level of performance is similar to that of the MH–96 adaptive controller in the case where no failures are present.

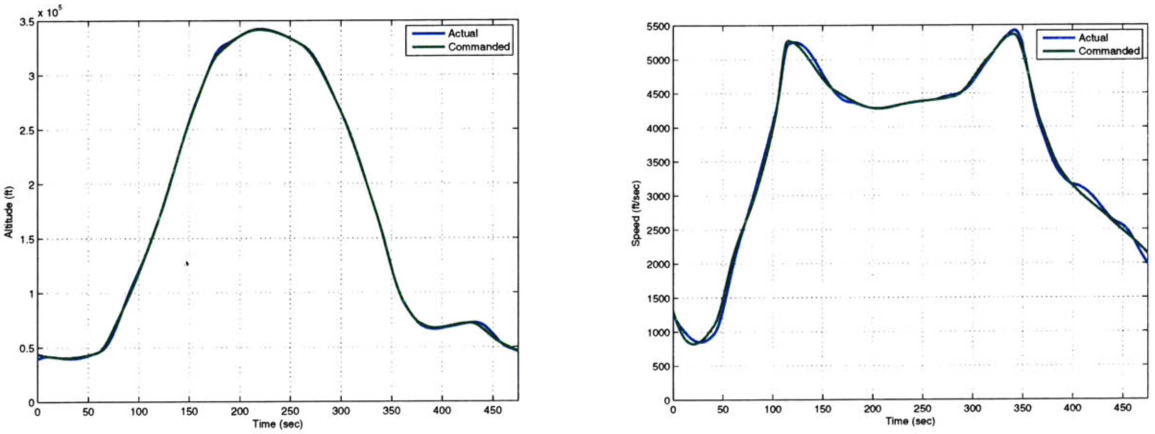

(a) Altitude **(b) Speed**

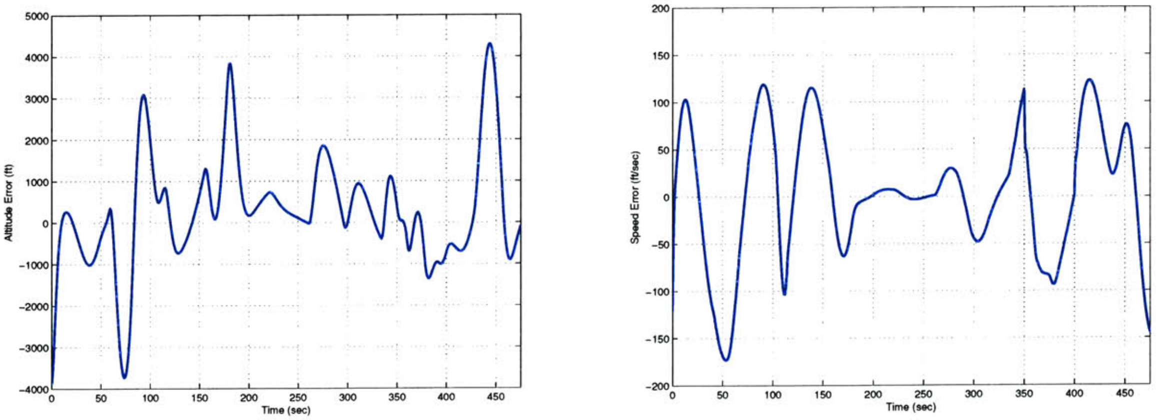

(c) Altitude Error **(d) Speed Error**

Figure 4-8: Tracking performance of the X–15 with modern adaptive controller (nominal case). 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.

The tracking performance in failure case ($\Lambda = \text{diag}([1 \quad 1 \quad 0.2 \quad 1])$) is shown in Figure 4-9 and the corresponding control inputs are shown in Figure 4-10. Not only does the adaptive controller maintain stability, as opposed to the MH–96 controller which loses stability resulting in an eventual crash as seen in Figure 4-3, the performance in the failure case is comparable to that of the nominal case! That is, the desired performance of the modern adaptive controller is retained despite severe parametric uncertainty. The adaptive controller accomplishes this while at or near the actuator limits for significant periods of time as can be seen in Figure 4-10 below.

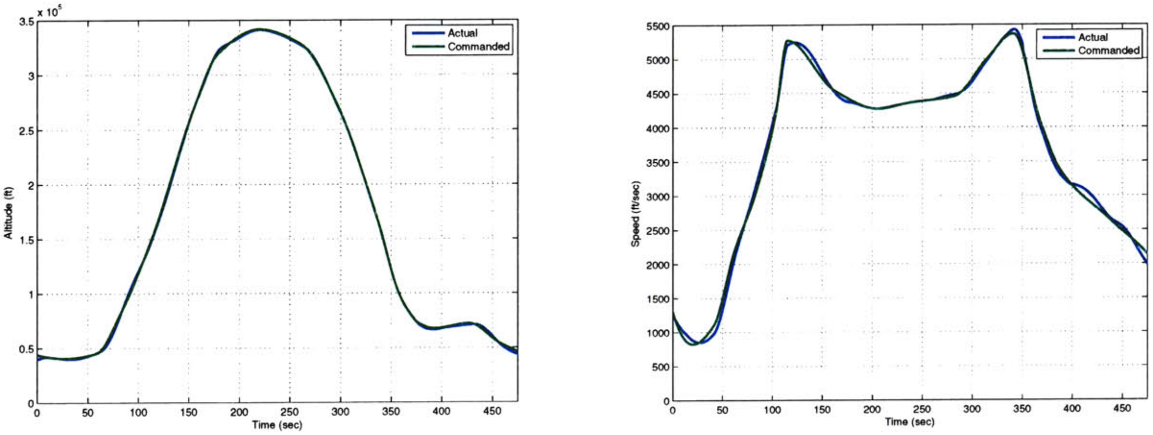

(a) Altitude

(b) Speed

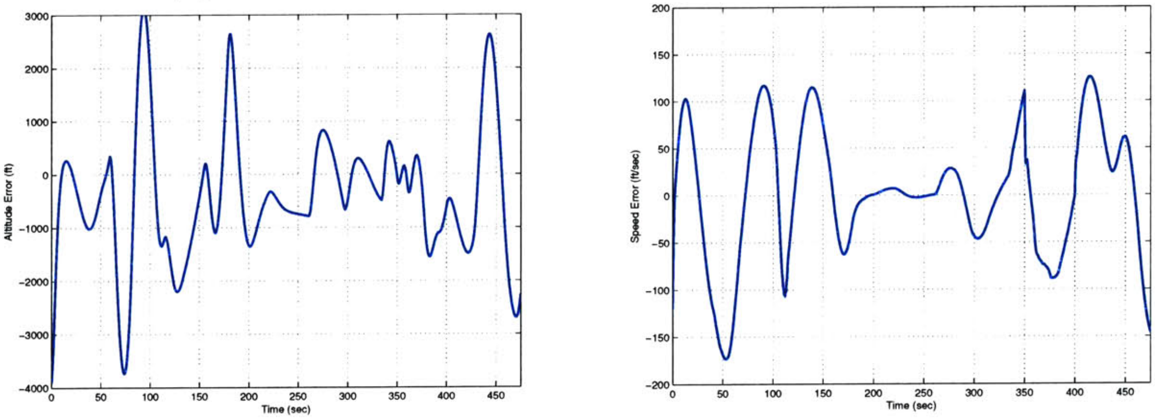

(c) Altitude Error

(d) Speed Error

Figure 4-9: Tracking performance of the X–15 with modern adaptive controller (failure case). The adaptive controller maintains high performance despite the severe failure present.

The modern adaptive controller succeeds in stabilizing the X–15 where the MH–96

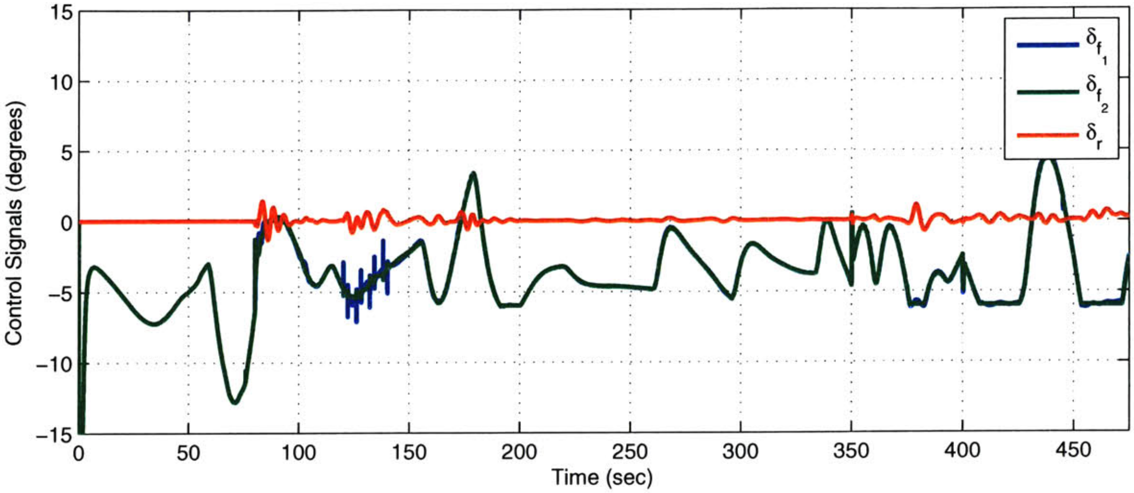

Figure 4-10: The control inputs for the maneuver above. Note that while the left actuator's limit is at the standard value of ±30°, the right actuator's limit is actually ±6° because of the 80% loss of control effectiveness.

controller failed for several reasons. First, the modern 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 nonlinearities 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 modern adaptive controller is based on theory, and has the advantage of proven stability.

Chapter 5

Summary

Guaranteeing certain levels of performance and robustness for adaptive systems requires the development of metrics which can be used to gauge the ability of nonlinear controllers to reject parametric and dynamic disturbances. The critical value of the parameter μ^* for adaptive systems is analogous to the linear control design concepts of gain and phase margin in that it concisely describes the overall robustness of the closed-loop system, and in that it can be used to assist in the selection of free parameters, such as the adaptive gain Γ. The process of the numerical calculation of these margins is demonstrated for the X–15 aircraft coupled with a modern adaptive controller.

The importance of having guaranteed robustness is further highlighted in the historical example of the X–15. 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 a modern 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 modern 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[16]. 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.

The advantages of adaptive control for aircraft, which include increased performance, increased robustness to uncertainties, more expedient design cycle, and robustness to minor changes in the aircraft design, can now be coupled with proven stability and guarantees on performance and robustness. This opens the door for widespread application of adaptive control on aerospace applications ranging from unmanned aerial vehicles and guided missiles to high-speed, high-performance manned aircraft.

Bibliography

[1] Michael A. Bolender and David B. Doman. A non-linear model for the longitudinal dynamics of a hypersonic air-breathing vehicle. *AIAA Guidance Navigation and Control Conference*, 2005.

[2] Jovan D. Bošković, Ssu-Hsin Yu, and Raman K. Mehra. A stable scheme for automatic control reconfiguration in the presence of actuator failures. Proceedings of the American Control Conference, June 1998.

[3] Zachary T. Dydek, Himani Jain, Jinho Jang, Anuradha M. Annaswamy, and Eugene Lavretsky. Theoretically verifiable stability margins for an adaptive controller. *AIAA Guidance Navigation and Control Conference*, 2006.

[4] Kevin P. Groves, David O. Sigthorssony, Andrea Serraniz, Stephen Yurkovich, Michael A. Bolender, and David B. Doman. Reference command tracking for a linearized model of an air-breathing hypersonic vehicle. *AIAA Guidance Navigation and Control Conference*, 2005.

[5] Robert K. Heffley and Wayne F. Jewell. Aircraft handling qualities data. Contractor's Report CR-2144, NASA, Systems Technology, Inc. Hawthorne, CA 90250, December 1972.

[6] Jinho Jang, A. M. Annaswamy, and E. Lavretsky. Towards verifiable adaptive flight control in the presence of actuator anomalies. San Diego, CA, 2006. Conference on Decision and Control.

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

[8] S. P. Karason and A. M. Annaswamy. Adaptive control in the presence of input constraints. *IEEE Transaction on Automatic Control*, 46(11):2325–2330, November 1994.

[9] G. Kreisselmeier and K. S. Narendra. Stable model reference adaptive control in the presence of bounded disturbances. *IEEE Transactions on Automatic Control*, 27(6):1169–1175, December 1982.

[10] Kumpati S. Narendra and Anuradha M. Annaswamy. *Stable Adaptive Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1989.

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

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

[13] Mac Schwager and A. M. Annaswamy. Adaptation-based reconfiguration in the presence of actuator failures and saturation. Portland, OR, 2005. American Control Conference.

[14] Mac Schwager and A. M. Annaswamy. Adaptive control of multi-input systems with magnitude saturation constraints. Seville, Spain, 2005. Conference on Decision and Control.

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

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

[17] Gang Tao and Suresh M. Joshi. Adaptive output feedback compensation of variant actuator failures. American Control Conference, June 2005.

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

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

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

[21] Roxanah 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°. Technical Note D-2532, NASA, Flight Research Center, Edwards, CA, November 1964.

[22] Roxanah B. Yancey, Herman A. Rediess, and Glenn 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. Technical Note D-1060, NASA, Flight Research Center, Edwards, CA, January 1962.

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