AIAA Guidance, Navigation, and Control Conference
2 - 5 August 2010, Toronto, Ontario Canada

AIAA 2010-8404

Retrospective Cost Adaptive Control of the NASA GTM Model

Benjamin C. Coffer[*], Jesse B. Hoagg[†] and Dennis S. Bernstein[‡]

Department of Aerospace Engineering, The University of Michigan, Ann Arbor, MI 48109-2140

A retrospective cost adaptive control algorithm is used to control the NASA Generic Transport Model (GTM) under various operational environments and flight scenarios. In particular, the adaptive control algorithm is used to follow commanded flight trajectories and reject undesired flight disturbances (e.g., wind gusts) under nominal flight scenarios as well as damaged flight scenarios (e.g., limited control surface effectiveness). Retrospective cost adaptive control is effective for multi-input, multi-output systems that are either minimum phase or nonminimum phase. The adaptive control algorithm requires limited model information, specifically, the first nonzero Markov parameter and the nonminimum-phase transmission zeros of the transfer function from the control signal to the performance, which can be estimated from a finite number of Markov parameters. Furthermore, the adaptive control algorithm is effective for stabilization as well as command following and disturbance rejection, where the command and disturbance spectrum are unknown.

I. Introduction

Nonminimum-phase zeros present a major challenge to direct adaptive control, and thus many direct adaptive control methodologies invoke a minimum-phase assumption.[1–5] In fact, sampling may give rise to nonminimum-phase zeros even if the continuous-time system is minimum phase,[6] and these nonminimum-phase zeros must ultimately be accounted for by any adaptive control algorithm implemented digitally.

Various direct adaptive control algorithms exist for discrete-time systems.[5,7–12] However, the majority of these algorithms are restricted to minimum-phase systems. In Refs. 5, 7, discrete-time adaptive control laws are presented for stabilization and command following of minimum-phase systems. An extension is given in Ref. 8, which addresses the combined stabilization, command following, and disturbance rejection problem. Note that the results of Refs. 5, 7, 8 are restricted to minimum-phase systems.

Discrete-time adaptive controllers using a retrospective cost are known to be effective for stabilization, command following, and disturbance rejection for systems that are either minimum phase or nonminimum phase provided that knowledge of the nonminimum-phase zeros is available.[13–15] Furthermore, retrospective cost adaptive controllers are effective for command following and disturbance rejection where the spectra of the commands and disturbances are unknown and the disturbance is unmeasured. Proof of stability and convergence in the minimum-phase case is given in Ref. 8, while extensions to the nonminimum-phase case are described in Refs. 14, 15. Retrospective cost adaptive control uses a retrospective performance measure, in which the performance measurement is modified based on the difference between the actual past control inputs and the recomputed past control inputs, assuming that the current controller had been used in the past. This technique does not require matching assumptions on either the plant uncertainty or the disturbances, and no prior uncertainty set is required.

The adaptive laws of Refs. 13, 14 are derived by minimizing an instantaneous retrospective cost, which is a function of the retrospective performance at the current time step, whereas the adaptive laws of Ref. 15 are derived by minimizing a cumulative retrospective cost function, which is a function of the retrospective performance at the current time step and all previous time steps. The instantaneous retrospective cost adaptive controllers of Ref. 13,14 have been demonstrated on various experiments and applications, including

[*]Graduate student.
[†]Postdoctoral Research Fellow, AIAA Member.
[‡]Professor, AIAA Member.

the Air Force's deployable optical telescope testbed in Ref. 16, the NASA Generic Transport Model (GTM) in Ref. 17, and flow control problems in Ref. 18.

The goal of the present paper is to use the cumulative retrospective cost adaptive controller, presented in Ref. 15, to control the NASA GTM.[19,20] This paper is thus an extension of Ref. 17, where an instantaneous retrospective cost adaptive controller is used to control GTM. We show that using a cumulative retrospective cost function, which is minimized by a recursive least-squares algorithm, results in improved transient performance as compared to an instantaneous retrospective cost adaptive control. In addition, the present paper extends the results of Ref. 17 by considering a wider range of operational environments and flight scenarios. In particular, the present paper addresses nominal flight scenarios (i.e., undamaged flight) and damaged flight scenarios, including reduced rudder, aileron, and elevator effectiveness as well as rudder lock. Finally, the present paper adopts an extension of retrospective cost adaptive control that is effective in the presence of amplitude and rate saturation (described in Section V). More specifically, the adaptive controller that we examine assumes that the saturation nonlinearity is known; the value of the saturated input is then used within the retrospective cost update algorithm. Consequently, when the adaptive controller converges, the resulting fixed gain controller is a nonlinear compensator.

II. Problem Formulation

Consider the multi-input, multi-output discrete-time system

$$x(k+1) = Ax(k) + Bu(k) + D_1w(k), \tag{1}$$

$$y(k) = Cx(k) + Du(k) + D_2w(k), \tag{2}$$

$$z(k) = E_1x(k) + E_2u(k) + E_0w(k), \tag{3}$$

where $x(k) \in \mathbb{R}^n$, $y(k) \in \mathbb{R}^{l_y}$, $z(k) \in \mathbb{R}^{l_z}$, $u(k) \in \mathbb{R}^{l_u}$, $w(k) \in \mathbb{R}^{l_w}$, and $k \geq 0$. Our goal is to develop an adaptive controller that generates a control signal u that minimizes the performance z in the presence of the exogenous signal w. We assume that measurements of the output y and performance z are available for feedback; however, we do not assume that a direct measurement of the exogenous signal w is available; however, a measurement of w can be used if it is available.

Note that w can represent either a command signal to be followed, an external disturbance to be rejected, or both. For example, if $D_1 = 0$, $E_2 = 0$, and $E_0 \neq 0$, then the objective is to have the output E_1x follow the command signal $-E_0w$. On the other hand, if $D_1 \neq 0$, $E_2 = 0$, and $E_0 = 0$, then the objective is to reject the disturbance w from the performance measurement E_1x. The combined command following and disturbance rejection problem is addressed when D_1 and E_0 are block matrices. Lastly, if D_1 and E_0 are empty matrices, then the objective is output stabilization, that is, convergence of z to zero.

We represent (1) and (3) as the time-series model from u and w to z given by

$$z(k) = \sum_{i=1}^{n} -\alpha_i z(k-i) + \sum_{i=d}^{n} \beta_i u(k-i) + \sum_{i=0}^{n} \gamma_i w(k-i), \tag{4}$$

where $\alpha_1, \ldots, \alpha_n \in \mathbb{R}$, $\beta_d, \ldots, \beta_n \in \mathbb{R}^{l_z \times l_u}$, $\gamma_0, \ldots, \gamma_n \in \mathbb{R}^{l_z \times l_w}$, and the relative degree d is the smallest non-negative integer i such that the ith Markov parameter, either $H_0 \triangleq E_2$ if $i = 0$ or $H_i \triangleq E_1A^{i-1}B$ if $i > 0$, is nonzero. Note that $\beta_d = H_d$.

III. Review of the Cumulative Retrospective Cost Adaptive Controller

In this section, we review the cumulative retrospective cost adaptive controller presented in Ref. 15. Consider a strictly proper time-series controller of order n_c, such that the control $u(k)$ is given by

$$u(k) = \sum_{i=1}^{n_\mathrm{c}} M_i(k)u(k-i) + \sum_{i=1}^{n_\mathrm{c}} N_i(k)y(k-i), \tag{5}$$

where, for all $i = 1, \ldots, n_\mathrm{c}$, $M_i : \mathbb{N} \to \mathbb{R}^{l_u \times l_u}$ and $N_i : \mathbb{N} \to \mathbb{R}^{l_u \times l_y}$ are determined by the adaptive control law presented below. The control (5) can be expressed as

$$u(k) = \theta(k)\phi(k), \tag{6}$$

where

$$\theta(k) \triangleq \begin{bmatrix} N_1(k) & \cdots & N_{n_c}(k) & M_1(k) & \cdots & M_{n_c}(k) \end{bmatrix},$$

and

$$\phi(k) \triangleq \begin{bmatrix} y(k-1) \\ \vdots \\ y(k-n_c) \\ u(k-1) \\ \vdots \\ u(k-n_c) \end{bmatrix} \in \mathbb{R}^{n_c(l_u+l_y)}. \tag{7}$$

Next, we define the retrospective performance

$$\hat{z}(\hat{\theta},k) \triangleq z(k) + \sum_{i=d}^{\nu} \bar{\beta}_i \left[\hat{\theta} - \theta(k-i)\right] \phi(k-i), \tag{8}$$

where $\nu \geq d$, $\hat{\theta} \in \mathbb{R}^{l_u \times (n_c(l_y+l_u))}$ is an optimization variable used to derive the adaptive law, and $\bar{\beta}_d, \ldots, \bar{\beta}_\nu \in \mathbb{R}^{l_z \times l_u}$. The choice of ν and $\bar{\beta}_d, \ldots, \bar{\beta}_\nu$ is discussed in Section IV.

Next, defining $\hat{\Theta} \triangleq \text{vec}\,\hat{\theta} \in \mathbb{R}^{n_c l_u (l_y+l_u)}$ and $\Theta(k) \triangleq \text{vec}\,\theta(k) \in \mathbb{R}^{n_c l_u(l_y+l_u)}$, it follows that

$$\hat{z}(\hat{\Theta},k) = z(k) + \sum_{i=d}^{\nu} \Phi_i^{\mathrm{T}}(k) \left[\hat{\Theta} - \Theta(k-i)\right] = z(k) - \sum_{i=d}^{\nu} \Phi_i^{\mathrm{T}}(k)\Theta(k-i) + \Psi^{\mathrm{T}}(k)\hat{\Theta}, \tag{9}$$

where, for $i = d, \ldots, \nu$,

$$\Phi_i(k) \triangleq \phi(k-i) \otimes \bar{\beta}_i^{\mathrm{T}} \in \mathbb{R}^{(n_c l_u(l_y+l_u)) \times l_z},$$

where $\otimes$ represents the Kronecker product, and

$$\Psi(k) \triangleq \sum_{i=d}^{\nu} \Phi_i(k).$$

Now, define the cumulative retrospective cost function

$$J(\hat{\Theta},k) \triangleq \sum_{i=0}^{k} \lambda^{k-i} \hat{z}^{\mathrm{T}}(\hat{\Theta},i) R \hat{z}(\hat{\Theta},i) + \lambda^k \left(\hat{\Theta} - \Theta(0)\right)^{\mathrm{T}} Q \left(\hat{\Theta} - \Theta(0)\right), \tag{10}$$

where $\lambda \in (0,1]$, and $R \in \mathbb{R}^{l_z \times l_z}$ and $Q \in \mathbb{R}^{(n_c l_u(l_y+l_u)) \times (n_c l_u(l_y+l_u))}$ are positive definite. Note that λ serves as a forgetting factor, which allows more recent data to be weighted more heavily than earlier data.

The cumulative retrospective cost function (10) is minimized by a recursive least-squares (RLS) algorithm with a forgetting factor.[2,4,5] Therefore, $J(\hat{\Theta},k)$ is minimized by the adaptive law

$$\Theta(k+1) = \Theta(k) - P(k)\Psi(k) \left[\lambda R^{-1} + \Psi^{\mathrm{T}}(k) P(k) \Psi(k)\right]^{-1} z_{\mathrm{R}}(k), \tag{11}$$

$$P(k+1) = \frac{1}{\lambda} P(k) - \frac{1}{\lambda} P(k)\Psi(k) \left[\lambda R^{-1} + \Psi^{\mathrm{T}}(k) P(k) \Psi(k)\right]^{-1} \Psi^{\mathrm{T}}(k) P(k), \tag{12}$$

where $P(0) = Q^{-1}$, $\Theta(0) \in \mathbb{R}^{n_c l_u(l_y+l_u)}$, and the retrospective performance measure $z_{\mathrm{R}}(k) \triangleq \hat{z}(\Theta(k),k)$. Note that $z_{\mathrm{R}}(k)$ is computable from (9) using measured signals z, y, u, θ and the matrix coefficients $\bar{\beta}_d, \ldots, \bar{\beta}_\nu$. The cumulative retrospective cost adaptive control law is thus given by (11), (12), and

$$u(k) = \theta(k)\phi(k) = \text{vec}^{-1}(\Theta(k))\phi(k). \tag{13}$$

The key feature of the retrospective cost adaptive control (RCAC) algorithm is the use of the retrospective performance (9), which modifies the performance variable $z(k)$ based on the difference between the actual past control inputs $u(k-d),\ldots,u(k-\nu)$ and the recomputed past control inputs $\hat{u}(\hat{\Theta},k-d) \stackrel{\triangle}{=} \mathrm{vec}^{-1}(\hat{\Theta})\phi(k-d),\ldots,\hat{u}(\hat{\Theta},k-\nu) \stackrel{\triangle}{=} \mathrm{vec}^{-1}(\hat{\Theta})\phi(k-\nu)$, assuming that the current controller $\hat{\Theta}$ had been used in the past.

Note that the cumulative retrospective cost adaptive controller (11)-(13) requires the matrix coefficients $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$. More specifically, these matrix coefficients are needed to compute the regressor $\Psi(k)$ and the retrospective performance measurement $z_{\mathrm{R}}(k)$. In the next section, we discuss how to select $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$.

IV. Choice of the Parameters $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$

In this section, we consider the case where the transfer function from u to z is potentially nonminimum phase, that is, the invariant zeros of (A,B,E_1,E_2) are not all contained inside of the unit circle. We present three constructions for the parameters $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$.

IV.A. Controller Construction Using Numerator Coefficients

First, consider the case where $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$ are the coefficients of the numerator polynomial matrix of the transfer function from u to z, that is, $\nu = n$ and, for $i = d,\ldots,n$, $\bar{\beta}_i = \beta_i$. In this case, the controller uses knowledge of all transmission zeros from u to z.

IV.B. Controller Construction Using Nonminimum-Phase Transmission Zeros

The results of Ref. 8 for the minimum-phase case suggests that RCAC requires knowledge of only the first nonzero Markov parameter and the nonminimum-phase transmission zeros of the transfer function from u to z. In this section, we choose $\bar{\beta}_d,\ldots,\bar{\beta}_\nu$ to capture the information contained in the first nonzero Markov parameter and the portion of the numerator polynomial matrix that includes the nonminimum-phase transmission zeros. Consider the matrix transfer function from u to z given by

$$G_{zu}(\mathbf{z}) \stackrel{\triangle}{=} \frac{1}{\alpha(\mathbf{z})}\beta(\mathbf{z}),$$

where $\alpha(\mathbf{z}) \stackrel{\triangle}{=} \mathbf{z}^n + \alpha_1\mathbf{z}^{n-1} + \cdots + \alpha_{n-1}\mathbf{z} + \alpha_n$ and $\beta(\mathbf{z}) \stackrel{\triangle}{=} \mathbf{z}^{n-d}\beta_d + \mathbf{z}^{n-d-1}\beta_{d+1} + \cdots + \mathbf{z}\beta_{n-1} + \beta_n$. Next, let $\beta(\mathbf{z})$ have the polynomial matrix factorization

$$\beta(\mathbf{z}) = \beta_{\mathrm{U}}(\mathbf{z})\beta_{\mathrm{S}}(\mathbf{z}),$$

where $\beta_{\mathrm{U}}(\mathbf{z})$ is an $l_z \times l_u$ polynomial matrix of degree $n_{\mathrm{U}} \geq 0$ whose leading matrix coefficient is β_d, $\beta_{\mathrm{S}}(\mathbf{z})$ is a monic $l_u \times l_u$ polynomial matrix of degree $n - n_{\mathrm{U}} - d$, and each Smith zero of $\beta(\mathbf{z})$ counting multiplicity that lies on or outside the unit circle is a Smith zero of $\beta_{\mathrm{U}}(\mathbf{z})$. More precisely, if $\lambda \in \mathbb{C}$, $|\lambda| \geq 1$, and rank $\beta(\lambda) <$ normal rank $\beta(\mathbf{z})$, then rank $\beta_{\mathrm{U}}(\lambda) <$ normal rank $\beta_{\mathrm{U}}(\mathbf{z})$ and rank $\beta_{\mathrm{S}}(\lambda) =$ normal rank $\beta_{\mathrm{S}}(\mathbf{z})$. Furthermore, we can write

$$\beta_{\mathrm{U}}(\mathbf{z}) = \beta_{\mathrm{U},0}\mathbf{z}^{n_{\mathrm{U}}} + \beta_{\mathrm{U},1}z^{n_{\mathrm{U}}-1} + \cdots + \beta_{\mathrm{U},n_{\mathrm{U}}-1}z + \beta_{\mathrm{U},n_{\mathrm{U}}},$$

where $\beta_{\mathrm{U},0} \stackrel{\triangle}{=} \beta_d$.

Now, we present a variation of RCAC that uses the coefficients of $\beta_{\mathrm{U}}(\mathbf{z})$ instead of the coefficients of $\beta(\mathbf{z})$. In this case, we let $\nu = n_{\mathrm{U}} + d$ and for $i = d,\ldots,n_{\mathrm{U}}+d$, $\bar{\beta}_i = \beta_{\mathrm{U},i-d}$.

If the transfer function from u to z is minimum phase (i.e., the invariant zeros of (A,B,E_1,E_2) are contained inside of the unit circle), then $\beta_{\mathrm{U}}(\mathbf{z}) = H_d$. In this case, RCAC requires only a single Markov parameter, namely, H_d. More specifically, we let $\nu = d$ and $\bar{\beta}_d = H_d$.

If the transfer function from u to z is nonminimum phase, then it can be difficult in practice to identify the coefficients of $\beta_{\mathrm{U}}(\mathbf{z})$. This motivates a third construction of the cumulative retrospective cost adaptive controller, which uses a finite number of Markov parameters from u to z.

IV.C. Controller Construction Using Markov Parameters

Consider the μ-MARKOV model of (4) obtained from μ successive back-substitutions of (4) into itself, and given by

$$
\begin{aligned}
z(k) = & -\sum_{i=1}^{n} \alpha_{\mu,i} z(k-\mu-i) + \sum_{i=d}^{\mu} H_{zu,i} u(k-i) + \sum_{i=1}^{n} \beta_{\mu,i} u(k-\mu-i) \\
& + \sum_{i=0}^{\mu} H_{zw,i} w(k-i) + \sum_{i=1}^{n} \gamma_{\mu,i} w(k-\mu-i),
\end{aligned} \tag{14}
$$

where $\alpha_{\mu,i} \in \mathbb{R}$, $\beta_{\mu,i} \in \mathbb{R}^{l_z \times l_u}$, $\gamma_{\mu,i} \in \mathbb{R}^{l_z \times l_w}$, $H_{zu,i} \in \mathbb{R}^{l_z \times l_u}$, $H_{zw,i} \in \mathbb{R}^{l_z \times l_w}$, and $\mu \geq d$. Thus, the μ-MARKOV transfer function from u to z is given by

$$
G_{zu,\mu}(\mathbf{z}) = \frac{1}{p_\mu(\mathbf{z})} \left(H_{zu,d} \mathbf{z}^{\mu+n-d} + \cdots + H_{zu,\mu} \mathbf{z}^{n} \right) + \frac{1}{p_\mu(\mathbf{z})} \left(\beta_{\mu,1} \mathbf{z}^{n-1} + \cdots + \beta_{\mu,n} \right), \tag{15}
$$

where $p_\mu(\mathbf{z}) \stackrel{\triangle}{=} \mathbf{z}^{\mu+n} + \alpha_{\mu,1}\mathbf{z}^{n-1} + \cdots + \alpha_{\mu,n}$. This system representation is nonminimal, overparameterized, and has order $n+\mu$. Note that the coefficients of the terms $\mathbf{z}^{n+\mu-1}$ through $\mathbf{z}^{n}$ in the denominator are zero.

The Laurent series expansion of $G_{zu}(\mathbf{z})$ about $\mathbf{z} = \infty$ is

$$
G_{zu}(\mathbf{z}) = \sum_{i=d}^{\infty} \mathbf{z}^{-i} H_{zu,i}.
$$

Truncating the numerator and denominator of (15) is equivalent to the truncated Laurent series expansion of $G_{zu}(\mathbf{z})$ about $\mathbf{z} = \infty$. Thus, the truncated Laurent series expansion of $G_{zu}(\mathbf{z})$ is

$$
\bar{G}_{zu,\mu}(\mathbf{z}) \stackrel{\triangle}{=} \sum_{i=d}^{\mu} \mathbf{z}^{-i} H_{zu,i}.
$$

Note that, for a single-input, single-output system, a subset of the roots of the polynomial

$$
H(\mathbf{z}) \stackrel{\triangle}{=} \mathbf{z}^{\mu-d} H_{zu,d} + \mathbf{z}^{\mu-d-1} H_{zu,d+1} + \cdots + \mathbf{z} H_{zu,\mu-1} + H_{zu,\mu} \tag{16}
$$

can be shown to approximate the nonminimum-phase zeros from u to z that lie outside of a circle in the complex plane centered at the origin with radius equal to the spectral radius of A. Thus, knowledge of $H_{zu,d}, \ldots, H_{zu,\mu}$ encompasses knowledge of the nonminimum-phase zeros from u to z that lie outside of the spectral radius of A.

Therefore, we present a variation of RCAC that uses only the Markov parameters $H_{zu,d}, \ldots, H_{zu,\mu}$. In this case, we let $\nu = \mu$ and for $i = d, \ldots, \mu$, $\bar{\beta}_i = H_{zu,i}$. This choice of $\bar{\beta}_d, \ldots, \bar{\beta}_\nu$ works well provided that $\mu \geq d$ is chosen large enough so that roots of (16) approximate the nonminimum-phase zeros from u to z.

V. Cumulative Retrospective Cost Adaptive Control With Amplitude and Rate Saturation

In this section, we adjust the retrospective cost adaptive controller (11)-(13) to account for amplitude and rate saturation. More specifically, we assume that the adaptive algorithm has access to a measurement of the saturated control signal, $u_{\rm s}(k) = \mathcal{N}(u(k))$, where $\mathcal{N}(\cdot)$ is the input nonlinearity. In this case, we adjust the retrospective cost adaptive controller (11)-(13) to account for the amplitude-or-rate-saturation nonlinearity. Specifically, the amplitude-or-rate-saturated control signal $u_{\rm s}(k)$ is used to construct ϕ in place of the unsaturated control $u(k)$, that is, we replace ϕ given by (7) with

$$
\phi(k) \stackrel{\triangle}{=} \left[\begin{array}{cccccc} y^{\rm T}(k-1) & \cdots & y^{\rm T}(k-n_{\rm c}) & u_{\rm s}^{\rm T}(k-1) & \cdots & u_{\rm s}^{\rm T}(k-n_{\rm c}) \end{array} \right]^{\rm T}.
$$

No additional alterations to the controller construction are required. Note that it follows from (6) that the resulting control law includes amplitude-or-rate-saturating nonlinearities and is thus nonlinear even in the case where the adaptive gains $\theta(k)$ have converged and are constant.

VI. Retrospective Cost Adaptive Control of the NASA GTM using Markov Parameter Modeling Information

In this section, we present numerical examples to demonstrate adaptive control of GTM using the cumulative retrospective cost adaptive controller (11)-(13). All examples considered in this section are based on the fully nonlinear GTM. Furthermore, the examples considered in this section are based on a nominal flight condition with the following parameters:

1. Flight path angle and angle of attack of 0 and 3 degrees, respectively.
2. Body x-axis, y-axis, and z-axis velocities of 161.66, 0, and 7.12 feet/sec, respectively.
3. Angular velocities in roll, pitch, and yaw of 0, 0, and 0 degrees/sec, respectively.
4. Latitude, longitude, and altitude of 0 degrees, 0 degrees, and 800 feet, respectively.
5. Roll, pitch, and yaw angles of 0.07, 3, and 90 degrees, respectively.
6. Elevator, aileron, and rudder angles of 2.7, 0, and 0 degrees, respectively.

We assume that the performance equals the output measurement, that is, $z = y$. Furthermore, z consists of the deviations in the roll angle, yaw angle, and altitude from the commanded values. The control input u consists of the commanded deviations in elevator angle, left aileron angle, right aileron angle, and rudder angle from the nominal flight condition.

In this section, we control GTM using RCAC, where the parameters $\bar{\beta}_d, \ldots, \bar{\beta}_\nu$ are chosen to be a finite number of Markov parameters as discussed in Section IV.C. More specifically, we let $\nu = \mu = 10$ and let $\bar{\beta}_2, \ldots, \bar{\beta}_{10}$ be the first 9 nonzero Markov parameters of GTM linearized about the nominal trim condition. GTM is sampled at a frequency of 4 Hz, and the adaptive controller (11)-(13) is implemented in feedback with $n_c = 6$, $R = I_3$, and $P(0) = 0.001 I_{126}$. In all examples we initialize the adaptive controller to be zero, that is, $\theta(0) = 0$, and we do not use a forgetting factor in the adaptive controller, that is, $\lambda = 1$. Finally, we stress that RCAC is the only controller implemented in the feedback loop, that is, no baseline or nominal controller is used. The sampling frequency of 4 Hz is above the Nyquist frequency associated with all modes and thus is sufficiently fast to meet the performance objectives. In particular, since RCAC is a discrete-time adaptive-control methodology, there is no need to implement fast sampling to emulate a continuous-time adaptive-control law.

VI.A. Crosswind Disturbance Rejection

Consider the disturbance rejection problem where the control objective is to maintain the altitude, yaw, and roll at their trim conditions in the presence of a cross wind. In particular, after 200 seconds, a 40 knot crosswind impacts the aircraft. This crosswind remains at 40 knots for the duration of the simulation. Figure 1 shows the time history of the closed-loop performance and control. After a transient due to the crosswind disturbance, the altitude, roll angle, and yaw angle return to the trim conditions. However, the aircraft drifts sideways since the controller is not provided with the measurements needed to reject the crosswind disturbance from the latitude or latitudinal-velocity states.

VI.B. Altitude Command Following

Consider the altitude command following problem where the control objective is to have the altitude follow a sequence of 10-foot altitude drops while maintaining the trim roll angle and trim yaw angle. Figure 2 shows the time history of the closed-loop performance and control. The altitude follows the commanded trajectory with limited overshoot in the step response. Furthermore, the roll and yaw angles remain close to their trim conditions.

VI.C. Altitude Command Following when Elevator Effectiveness is Reduced to 60%

Consider the control objective of following a sequence of altitude doublets, while maintaining the trim roll angle and trim yaw angle. Furthermore, assume that after 550 seconds, the elevator experiences damage so that the elevator effectiveness is reduced to 60%. To allow the adaptive controller to adjust quickly to

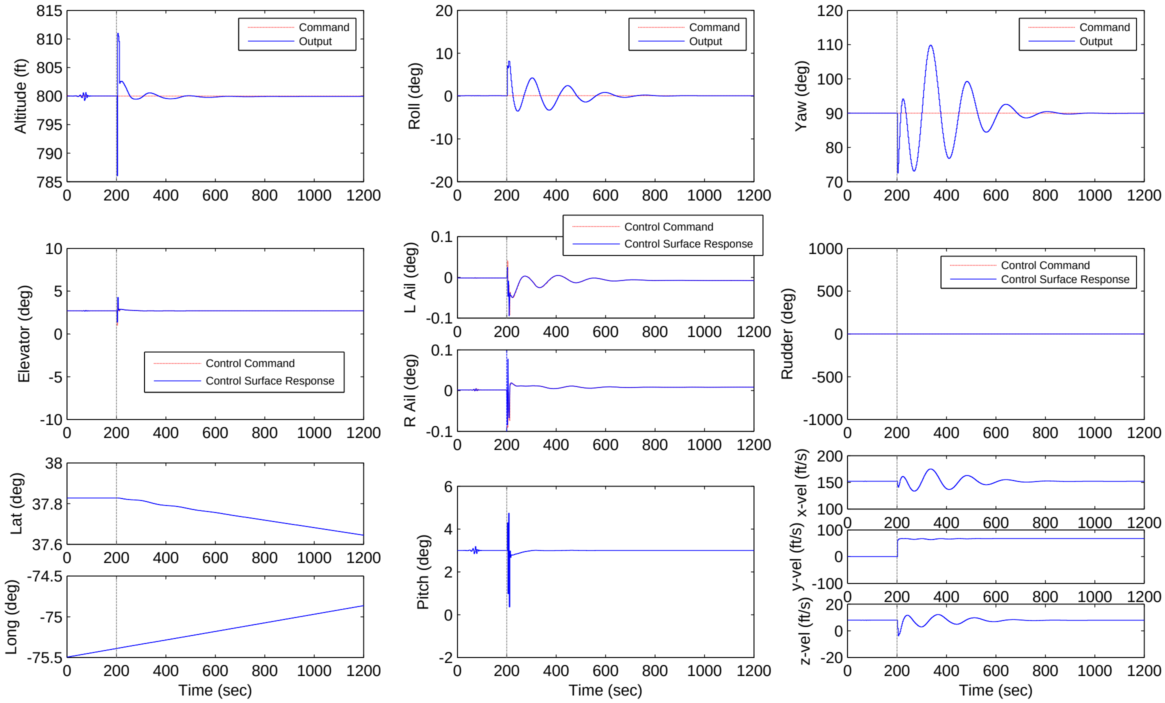

Figure 1. Crosswind Disturbance Rejection: The adaptive control (11)-(13) is implemented in feedback with GTM where $n_c = 6$, $R = I_3$, $P(0) = 0.1I_{126}$, and $\bar{\beta}_2, \ldots, \bar{\beta}_{10}$ chosen to be the first 9 nonzero Markov parameters of the linearized GTM. A 40 know crosswind disturbance is introduced at 200 seconds, and the controller adapts to reject the disturbance. Specifically, the altitude, roll angle, and yaw angle return to the trim conditions.

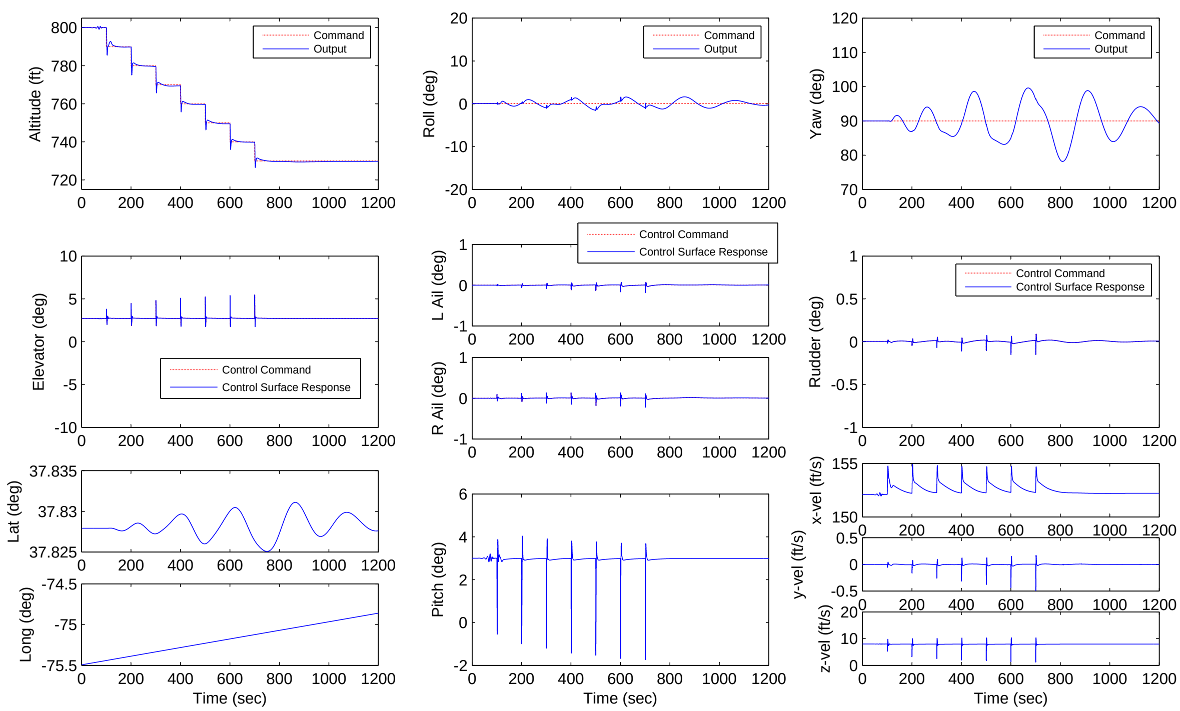

Figure 2. Altitude Command Following: The adaptive control (11)-(13) is implemented in feedback with GTM where $n_c = 6$, $R = I_3$, $P(0) = 0.1I_{126}$, and $\bar{\beta}_2, \ldots, \bar{\beta}_{10}$ chosen to be the first 9 nonzero Markov parameters of the linearized GTM. The altitude follows the commanded trajectory, while the roll and yaw angles remain close to the trim conditions.

the elevator damage, the recursive-least-squares covariance matrix $P(k)$ is reset to $P(0)$ after the elevator damage occurs. Figure 3 shows the time history of the closed-loop performance and control. Prior to the elevator damage, the altitude follows the commanded trajectory with limited overshoot in the step response, while the roll and yaw angles remain close to their trim conditions. After the elevator damage occurs at 550 seconds, GTM experiences transient flight behavior, after which the altitude follows the commanded trajectory. Furthermore, after the damage occurs, the adaptive controller correctly scales the elevator control signal by about 167% to account for the 60% drop in elevator effectiveness.

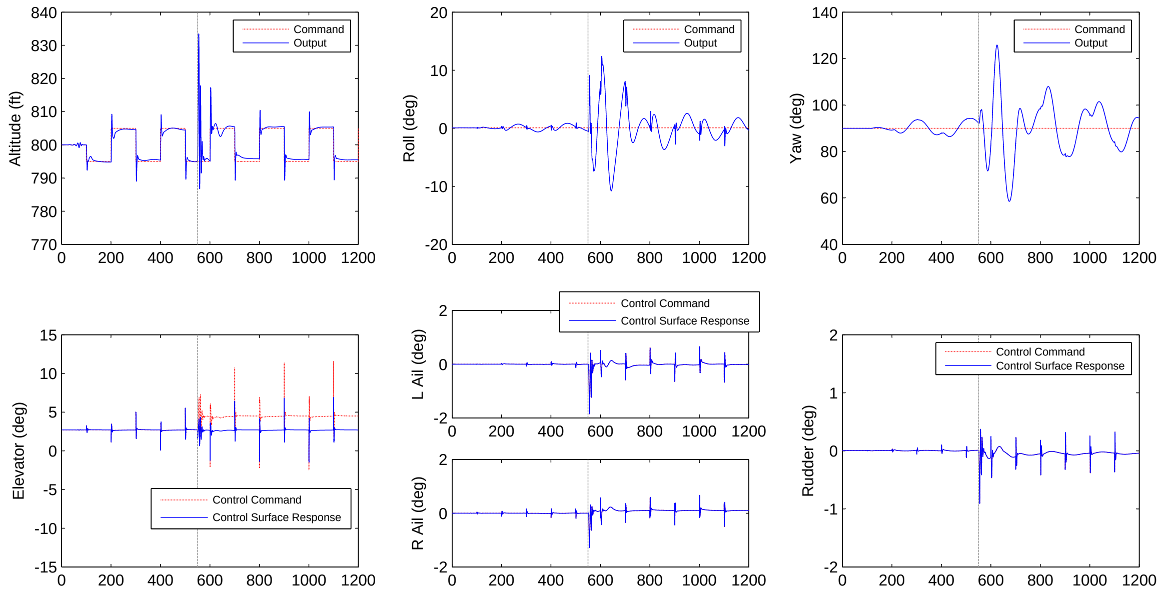

Figure 3. Altitude Command Following with Reduced Elevator Effectiveness: The adaptive control (11)-(13) is implemented in feedback with GTM where $n_c = 6$, $R = I_3$, $P(0) = 0.1I_{126}$, and $\bar{\beta}_2, \ldots, \bar{\beta}_{10}$ are chosen to be the first 9 nonzero Markov parameters of the linearized GTM. Prior to the elevator damage, the altitude follows the commanded trajectory. At 550 seconds, the elevator damage occurs and $P(k)$ is reset to $P(0)$. After the damage occurs, GTM experiences transient flight behavior, after which the altitude follows the commanded trajectory.

VI.D. Yaw Angle Command Following when the Rudder is Locked

Consider the control objective of following a constant yaw angle, while maintaining the trim altitude and roll angle. Furthermore, assume that after 600 seconds, the rudder experiences damage and is locked at an off-nominal angle. Specifically, we consider the cases where the rudder is locked at -0.1, -0.3, and -0.5 degrees. To allow the adaptive controller to adjust quickly to the rudder damage, the recursive-least-squares covariance matrix $P(k)$ is reset to $P(0)$ after the rudder damage occurs. Figure 4 shows the time history of the closed-loop performance and control for the cases where the rudder is locked at -0.1, -0.3, and -0.5 degrees. Prior to the rudder damage, GTM follows the constant commanded yaw angle, while the altitude and roll angle remain close to their trim conditions. After the rudder damage occurs at 600 seconds, GTM experiences transient flight behavior, after which the yaw angle converges to the commanded value. Note that the left and right aileron control angles converge to values that are not equal in magnitude. The difference in the left and right aileron control angles allows the closed-loop system to compensate for the rudder damage.

VII. Retrospective Cost Adaptive Control of the NASA GTM using a Model Reference Architecture

In this section, we present numerical examples to demonstrate control of GTM using RCAC with a model reference architecture. All examples considered in this section are based on the nominal flight conditions described in Section VI. Furthermore, the examples considered in this section are based on the linearized GTM about the nominal flight conditions. Future work will extend the results of this section to the nonlinear GTM. Finally, for all examples in this section, GTM is sampled at a frequency of 50 Hz.

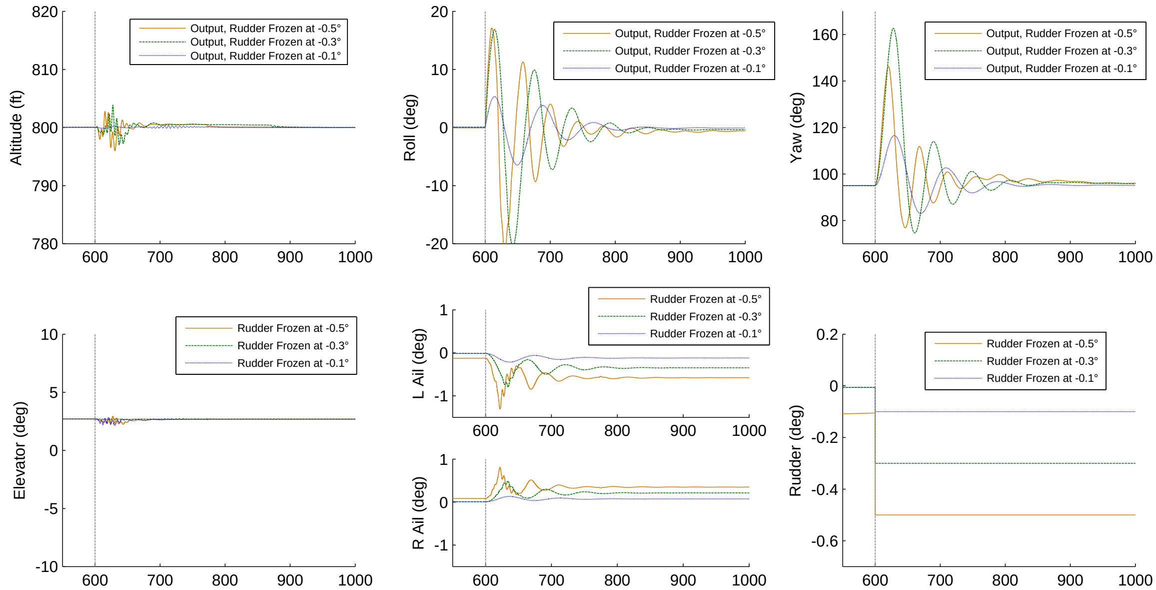

Figure 4. Yaw Angle Command Following when the Rudder is Locked: The adaptive control (11)-(13) is implemented in feedback with GTM, where $n_c = 6$, $R = I_3$, $P(0) = 0.1I_{126}$, and $\bar{\beta}_2, \ldots, \bar{\beta}_{10}$ are chosen to be the first 9 nonzero Markov parameters of the linearized GTM. Prior to the rudder damage, the yaw angle follows the constant command trajectory. At 600 seconds, the rudder damage occurs and $P(k)$ is reset to $P(0)$. After the damage occurs, GTM experiences transient flight behavior, after which the yaw angle converges to the commanded value. Note that the difference in the magnitude of the left and right aileron control angles allows the closed-loop system to compensate for the rudder damage.

For the results in this section, we implement two separate RCAC loops. The first loop is a single-input, single-output loop that controls GTM's behavior from elevator to altitude. The second loop is a two-input, two-output loop that controls GTM's behavior from rudder and ailerons to roll and yaw. Both adaptive control loops are run simultaneously, allowing for combined altitude-roll-yaw maneuvers. Note that we ignore (from a control design perspective only) the cross-coupling between certain inputs and outputs by implementing separate adaptive controllers for the elevator-to-altitude channel and the rudder-and-ailerons-to-roll-and-yaw channels. More specifically, the two-RCAC-loop architecture ignores the rudder-and-aileron-to-altitude channels as well as the elevator-to-roll-and-yaw channels, which numerical results suggest are weakly coupled.

Next, we briefly describe the model reference architecture adopted for RCAC in this section. Let $y_{\rm alt}(k)$, $y_{\rm roll}(k)$, and $y_{\rm yaw}(k)$ denote the altitude, roll, and yaw perturbations about the nominal flight conditions, respectively.

First, we describe model reference RCAC for the elevator-to-altitude control loop. Consider the discrete-time reference model given by

$$y_{\rm alt,ref} = G_{\rm alt,ref}(\mathbf{z})r_{\rm alt},$$

where $r_{\rm alt}(k)$ is the input to the reference model, $y_{\rm alt,ref}(k)$ is the output from the reference model, and the reference model is given by

$$G_{\rm alt,ref}(\mathbf{z}) \stackrel{\triangle}{=} \frac{\xi_1(\mathbf{z}-\zeta)}{(\mathbf{z}-0.7)^{10}},$$

where $\zeta \stackrel{\triangle}{=} 1.7062$ is an estimate of the nonminimum-phase zero from elevator to altitude (which cannot be moved using feedback) and $\xi_1 \stackrel{\triangle}{=} 8.361 \times 10^{-6}$ is chosen so that the reference model has unity gain at dc (i.e., $|G_{\rm alt,ref}(1)| = 1$). The objective of this control problem is to use a variation of RCAC to force $y_{\rm alt}$ to follow $y_{\rm alt,ref}$. More specifically, we use the adaptive controller (11)-(13), where the performance $z(k) = y_{\rm alt}(k) - y_{\rm alt,ref}(k)$, the feedback measurement $y(k) = \begin{bmatrix} y_{\rm alt}(k) & r_{\rm alt}(k) \end{bmatrix}^{\rm T}$, and additional modifications are made based on the model reference architecture. Note that the inclusion of $r_{\rm alt}(k)$ in $y(k)$ results in feedforward control as is standard in model reference architecture. For all the examples considered in this section, the model reference RCAC from elevator to altitude is implemented with $n_c = 9$, $R = 1$, and $P(0) = 10^{15} I_{27}$. Furthermore, RCAC uses only two parameters $\bar{\beta}_d \triangleq \beta_d$ and $\bar{\beta}_{d+1} \triangleq -\beta_d \zeta$. Thus, we require knowledge of only the first nonzero Markov parameter and an estimate of the nonminimum-phase zero from elevator to altitude. In all examples, $\theta(0) = 0$ and $\lambda = 1$.

Next, we describe model reference RCAC for the rudder-and-ailerons-to-roll-and-yaw control loop. Consider the discrete-time reference model given by

$$\begin{bmatrix} y_{\rm roll,ref} \\ y_{\rm yaw,ref} \end{bmatrix} = G_{\rm r-y,ref}(\mathbf{z}) \begin{bmatrix} r_{\rm roll} \\ r_{\rm yaw} \end{bmatrix},$$

where $r_{\rm roll}(k)$ and $r_{\rm yaw}(k)$ are the inputs to the reference model, $y_{\rm roll,ref}(k)$ and $y_{\rm yaw,ref}(k)$ are the outputs from the reference model, and the reference model is given by

$$G_{\rm r-y,ref}(\mathbf{z}) \triangleq \xi_2 \frac{1}{(\mathbf{z} - 0.95)^2} I_2,$$

where $\xi_2 \triangleq 2.5 \times 10^{-3}$ is chosen so that the reference model has unity gain at dc (i.e., $\|G_{\rm r-y,ref}(1)\| = 1$). The objective of this control problem is to use a variation of RCAC to force $y_{\rm roll}$ to follow $y_{\rm roll,ref}$ and $y_{\rm yaw}$ to follow $y_{\rm yaw,ref}$. We use the adaptive controller (11)-(13), where

$$z(k) = \begin{bmatrix} y_{\rm roll}(k) - y_{\rm roll,ref}(k) \\ y_{\rm yaw}(k) - y_{\rm yaw,ref}(k) \end{bmatrix} \qquad y(k) = \begin{bmatrix} y_{\rm roll}(k) \\ y_{\rm yaw}(k) \\ r_{\rm roll}(k) \\ r_{\rm yaw}(k) \end{bmatrix}, \tag{17}$$

and several other technical changes are implemented for the model reference architecture. For all the examples considered in this section, the model reference RCAC from rudder and ailerons to roll and yaw is implemented with $n_c = 9$, $R = 1$, and $P(0) = 10^6 I_{108}$. Furthermore, we use only one parameter, namely, the first nonzero Markov parameter, because the linearized transfer function from rudder and ailerons to roll and yaw is minimum phase. In all examples, $\theta(0) = 0$ and $\lambda = 1$. Finally, we stress that the two-loop model reference RCAC is the only controller implemented, that is, no baseline or nominal controller is used.

VII.A. Altitude Doublet and Sinusoid Command Following

Consider the altitude command following problem, where the control objective is to have GTM's altitude follow a sequence of 20-second doublets overlaid with a 1-Hz sinusoid, while maintaining the roll and yaw near trim. More specifically, we let $r_{\rm alt}$ be a sequence of 20-second doublets with amplitude of 5 feet plus a 1-Hz sinusoid with amplitude of 1 foot, and $r_{\rm roll} = r_{\rm yaw} = 0$. Figure 5 shows the time history of the commands and the closed-loop altitude, roll, and yaw. The roll and yaw responses stay close to the trim position, while the altitude follows the reference model with negligible transient. Note that the altitude response and reference model are indistinguishable in Figure 5.

VII.B. Altitude, Roll, and Yaw Sinusoidal Command Following

Consider the simultaneous altitude, roll, and yaw command following problem, where the control objective is to have GTM's altitude, roll, and yaw follow sinusoid commands. We let $r_{\rm alt}$ be a 1-Hz sinusoid with amplitude of 5 feet, $r_{\rm roll}$ be a 0.5-Hz sinusoid with amplitude of 0.3 radians, and $r_{\rm yaw}$ be a 0.5-Hz sinusoid with amplitude of 0.3 radians. Figure 6 shows the time history of the commands (i.e., the outputs of the reference models) and the closed-loop altitude, roll, and yaw.

VII.C. Altitude, Roll, and Yaw Doublet Command Following

Consider the simultaneous altitude, roll, and yaw command following problem, where the control objective is to have GTM dive 10 feet while rolling to 0.2 radians and yawing to 0.2 radians for 10 seconds, then have GTM climb 10 feet while rolling and yawing back to trim for 10 seconds, and finally repeat this maneuver.

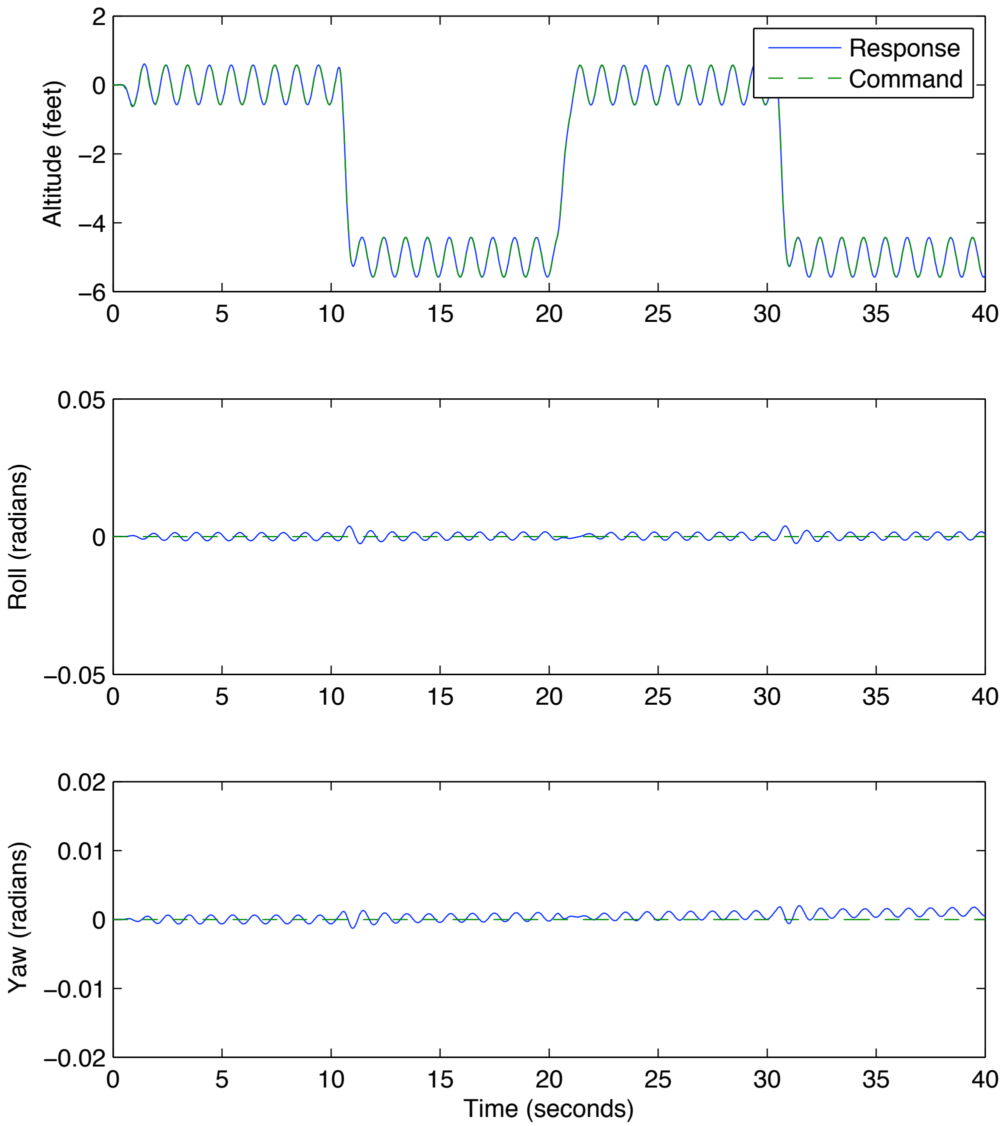

Figure 5. Altitude Doublet and Sinusoid Command Following: The model reference RCAC is used to follow a sequence of doublets and a 1-Hz sinusoid in altitude, while maintaining the trim roll and yaw positions. The adaptive controllers use only the first nonzero Markov parameter and the nonminimum-phase from the elevator-to-altitude channel, and the first nonzero Markov parameter from the rudder-and-ailerons-to-roll-and-yaw channels.

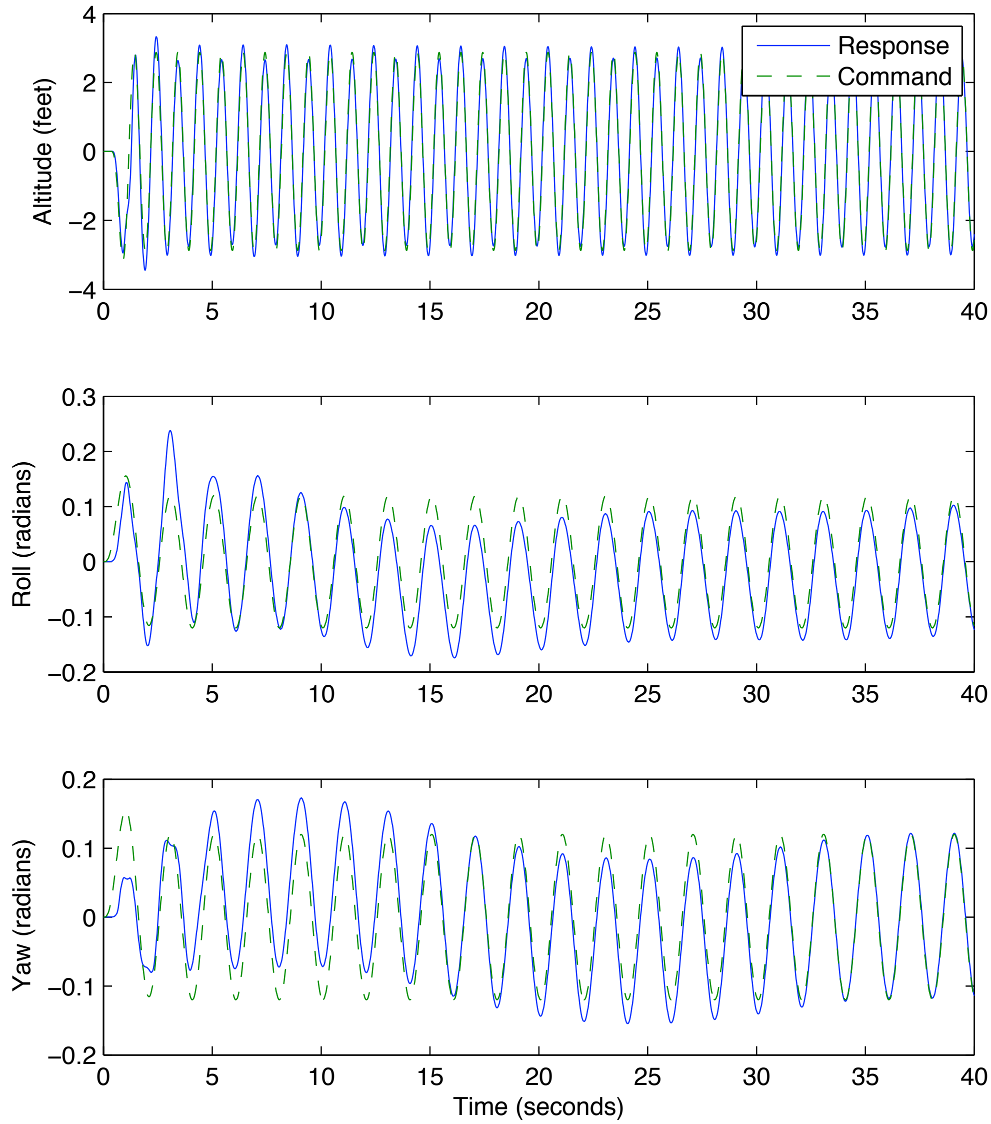

Figure 6. Altitude, Roll, and Yaw Sinusoidal Command Following: The model reference RCAC is used to follow sinusoids in altitude, roll, and yaw. The adaptive controllers use only the first nonzero Markov parameter and the nonminimum-phase from the elevator-to-altitude channel, and the first nonzero Markov parameter from the rudder-and-ailerons-to-roll-and-yaw channels.

More specifically, we let r_{alt} be a sequence of 20-second doublets with amplitude of 10 feet, r_{roll} be a sequence of 20-second doublets with amplitude of 0.2 radians, and r_{yaw} be a sequence of 20-second doublets with amplitude of 0.2 radians. Figure 7 shows the time history of the commands and the closed-loop altitude, roll, and yaw. Note that the altitude follows the reference model well from the start of the first doublet, whereas the roll and yaw require a couple of cycles to converge to the commanded signal.

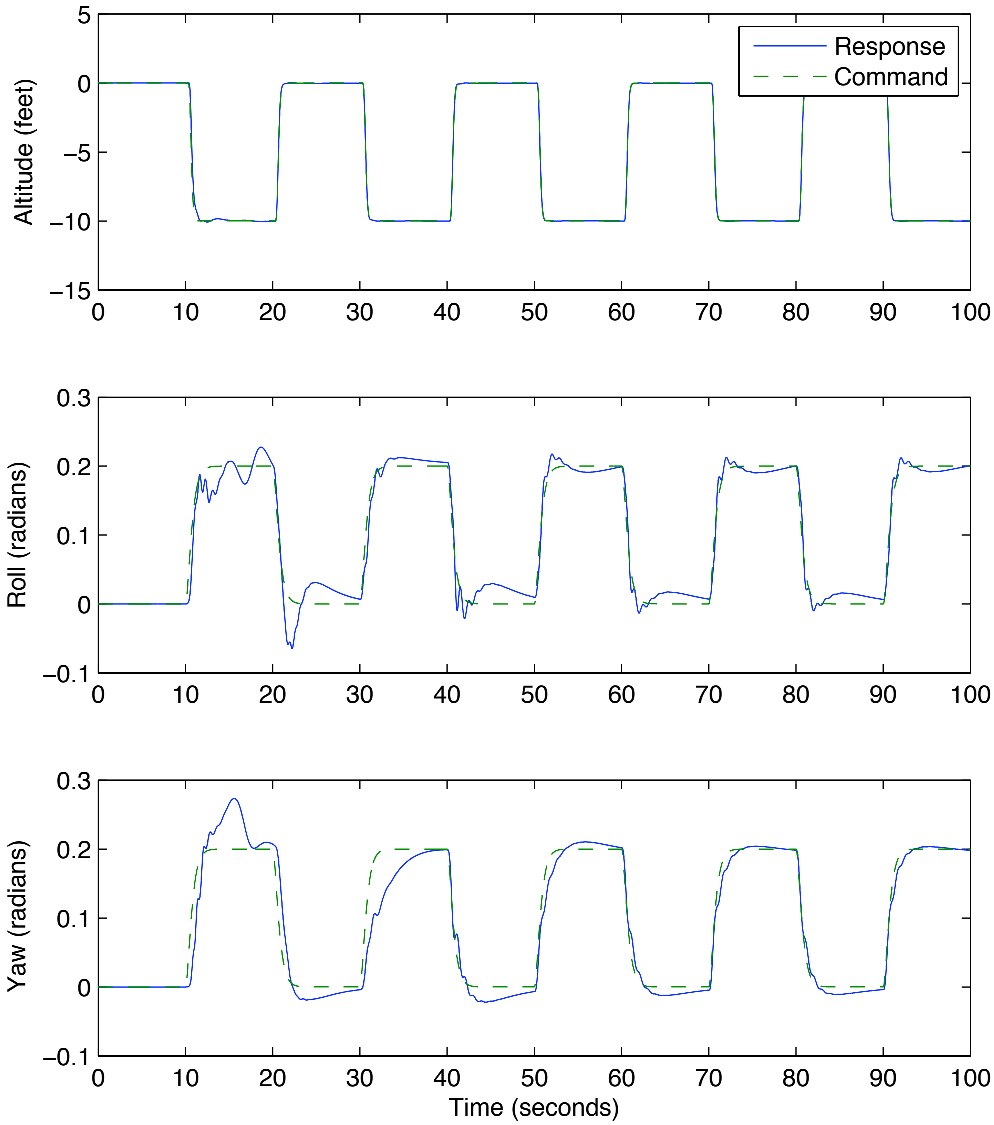

Figure 7. Altitude, Roll, and Yaw Doublet Command Following: The model reference RCAC is used to follow a sequence of doublets simultaneously in altitude, roll, and yaw. The adaptive controllers use only the first nonzero Markov parameter and the nonminimum-phase from the elevator-to-altitude channel, and the first nonzero Markov parameter from the rudder-and-ailerons-to-roll-and-yaw channels.

VII.D. Altitude Doublet and Sinusoid Command Following when Elevator Effectiveness is Reduced to 60%

Reconsider the altitude command following problem described in Section VII.A, where the control objective is to have GTM's altitude follow a sequence of 20-second doublets overlaid with a 1-Hz sinusoid, while maintaining the roll and yaw near trim. Furthermore, assume that the elevator experiences unknown (i.e., the recursive-least-squares covariance matrix is not reset) damage at 20 seconds such that the elevator effectiveness is reduced to 60% of its nominal effectiveness. Figure 8 shows the time history of the commands (i.e., the outputs of the reference models) and the closed-loop altitude, roll, and yaw. The roll and yaw responses stay close to the trim position through the entire simulation. Prior to the elevator damage, the altitude follows the reference model with negligible transient behavior. After the elevator damage, the altitude experiences a small transient and converges to the commanded signal. The associated elevator, aileron, and rudder control signals are shown in Figure 9. Note that the elevator control surface never exceeds ±20 degrees.

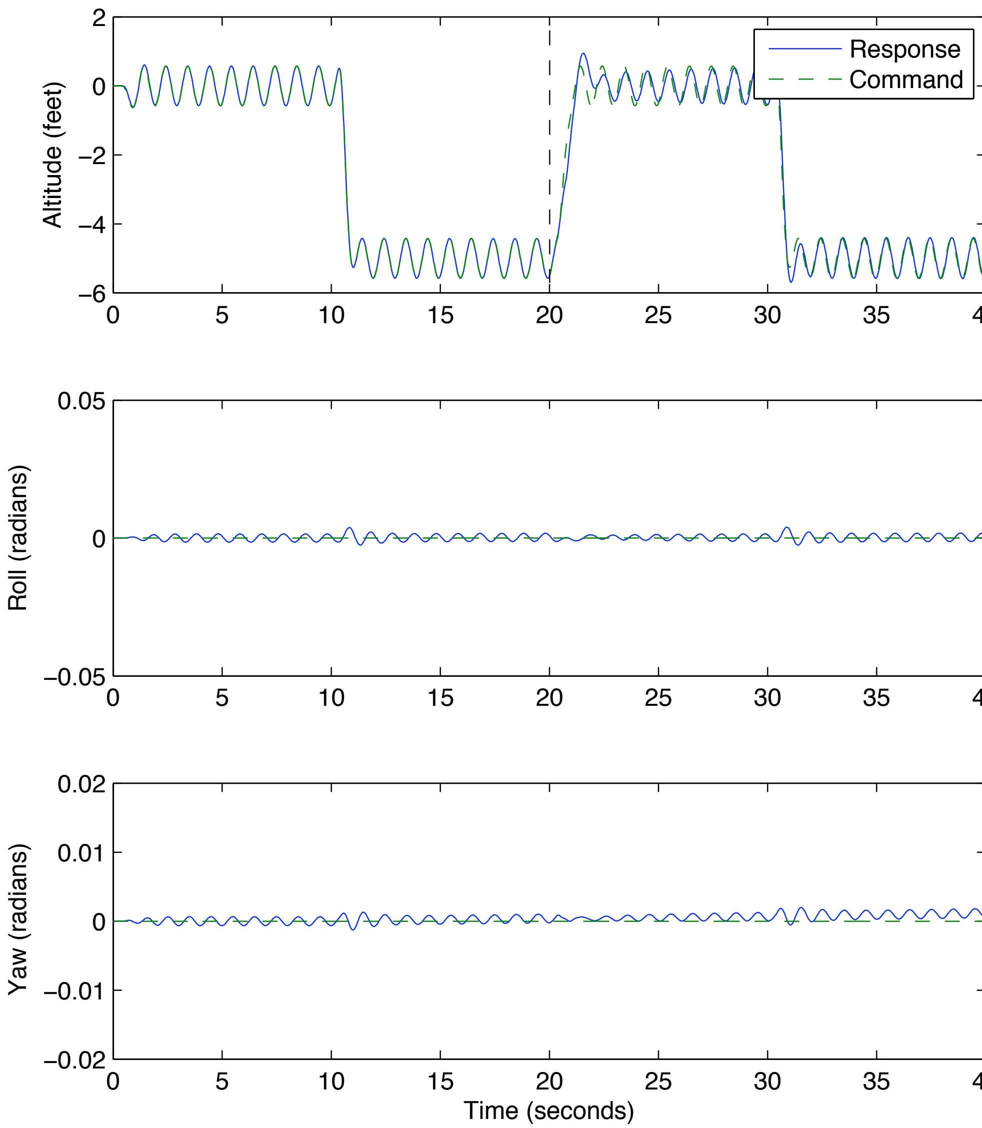

Figure 8. Altitude Doublet and Sinusoid Command Following when Elevator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow a sequence of doublets and a 1-Hz sinusoid in altitude, while maintaining the trim roll and yaw positions. The elevator experiences unknown damage at 20 seconds such that the elevator effectiveness is reduced to 60% of its nominal effectiveness. The plots show that RCAC corrects for the damage.

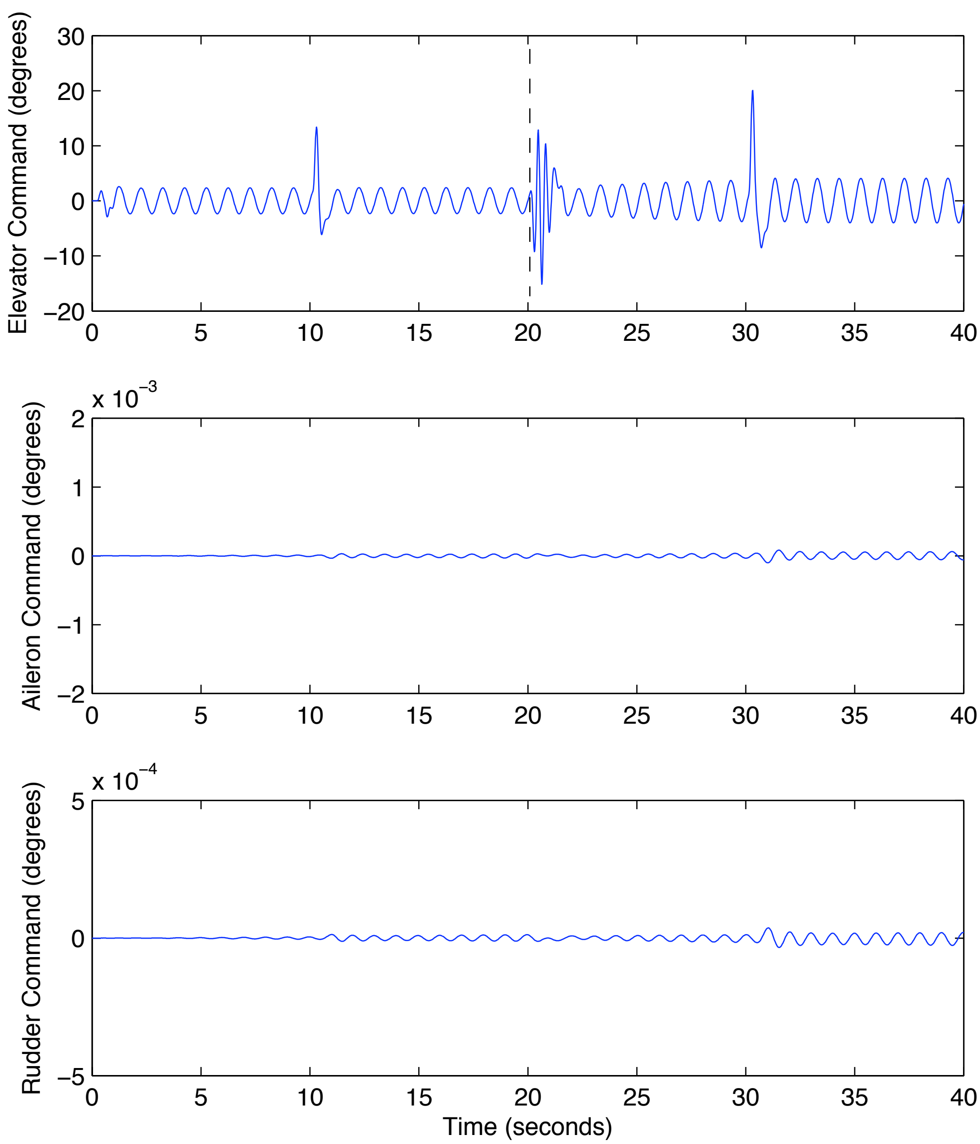

Figure 9. Altitude Doublet and Sinusoid Command Following Control Signals when Elevator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow a sequence of doublets and a 1-Hz sinusoid in altitude, while maintaining the trim roll and yaw positions. The elevator experiences unknown damage at 20 seconds such that the elevator effectiveness is reduced to 60% of its nominal effectiveness. The elevator control surface never exceeds ± 20 degrees, and the aileron and rudder control surfaces never exceed ± 0.001 degrees. Note that the magnitude to the elevator control signal increases after 20 seconds to compensate for the elevator damage.

VII.E. Altitude, Roll, and Yaw Sinusoidal Command Following when Actuator Effectiveness is Reduced to 60%

Reconsider the simultaneous altitude, roll, and yaw command following problem described in Section VII.B, where the control objective is to have GTM's altitude, roll, and yaw follow sinusoid commands. Furthermore, assume that the elevator, ailerons, and rudder all experience unknown damage at 40 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. Figure 10 shows the time history of the commands and the closed-loop altitude, roll, and yaw. Prior to the elevator damage, RCAC drives the altitude, roll, and yaw toward the commands. After the damage, the altitude, roll, and yaw exhibit small transient responses before RCAC drives the altitude, roll, and yaw toward the commands. The associated elevator, aileron, and rudder control signals are shown in Figure 11. Note that the magnitude of each control signal increases after 40 seconds to compensate for the damage.

VII.F. Altitude, Roll, and Yaw Doublet Command Following when Actuator Effectiveness is Reduced to 60%

Reconsider the simultaneous altitude, roll, and yaw command following problem described in Section VII.C, where the control objective is to have GTM's altitude, roll, and yaw follow a sequence of doublet commands. Furthermore, assume that the elevator, ailerons, and rudder all experience unknown damage at 100 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. Figure 12 shows the time history of the commands and the closed-loop altitude, roll, and yaw. Prior to the elevator damage, RCAC drives the altitude, roll, and yaw toward the commands. After the damage, the altitude, roll, and yaw exhibit transient responses before RCAC drives the altitude, roll, and yaw toward the commands. The associated elevator, aileron, and rudder control signals are shown in Figure 13. Note that the magnitude of each control signal increases after 100 seconds to compensate for the damage.

VIII. Conclusion

In this paper, we reviewed the cumulative retrospective cost adaptive controller, which was presented in Ref. 15, for stabilization, command following, and disturbance rejection. The algorithm requires limited information about the open-loop system. Specifically, it requires knowledge of the first nonzero Markov parameter and the nonminimum-phase zeros from the control to the performance measurement. Furthermore, we presented an extension to RCAC in Section V that helps to manage the effect of input nonlinearities (e.g., rate and amplitude saturation). We presented results on the use of RCAC to control GTM under various operational environments and flight scenarios, including nominal flight scenarios (i.e., undamaged flight) and damaged flight scenarios (i.e., reduced rudder, aileron, and elevator effectiveness as well as rudder lock). In addition, we demonstrated that RCAC (without the use of any baseline control scheme) can be effective on GTM with limited model information, specifically, a finite number of Markov parameters. Finally, we presented results on the use of RCAC within a model reference architecture to control the linearized GTM using only knowledge of the first nonzero Markov parameter and the nonminimum-phase zero in the transfer function from elevator to altitude. Future work will include extending the model reference RCAC results to the nonlinear GTM.

References

[1] Narendra, K. S. and Annaswamy, A. M., *Stable Adaptive Systems*, Prentice Hall, 1989.

[2] Åström, K. J. and Wittenmark, B., *Adaptive Control*, Addison-Wesley, 2nd ed., 1995.

[3] Ioannou, P. and Sun, J., *Robust Adaptive Control*, Prentice Hall, 1996.

[4] Tao, G., *Adaptive Control Design and Analysis*, Wiley, 2003.

[5] Goodwin, G. C. and Sin, K. S., *Adaptive Filtering, Prediction, and Control*, Prentice Hall, 1984.

[6] Åström, K. J., Hagander, P., and Sternby, J., "Zeros of Sampled Systems," *Automatica*, Vol. 20, 1984, pp. 31–38.

[7] Goodwin, G. C., Ramadge, P. J., and Caines, P. E., "Discrete-time multivariable adaptive control," *IEEE Trans. Autom. Contr.*, Vol. 25, 1980, pp. 449–456.

[8] Hoagg, J. B., Santillo, M. A., and Bernstein, D. S., "Discrete-Time Adaptive Command Following and Disturbance Rejection with Unknown Exogenous Dynamics," *IEEE Trans. Autom. Contr.*, Vol. 53, 2008, pp. 912–928.

[9] Akhtar, S. and Bernstein, D. S., "Lyapunov-stable discrete-time model reference adaptive control," *Int. J. Adaptive Contr. Signal Proc.*, Vol. 19, 2005, pp. 745–767.

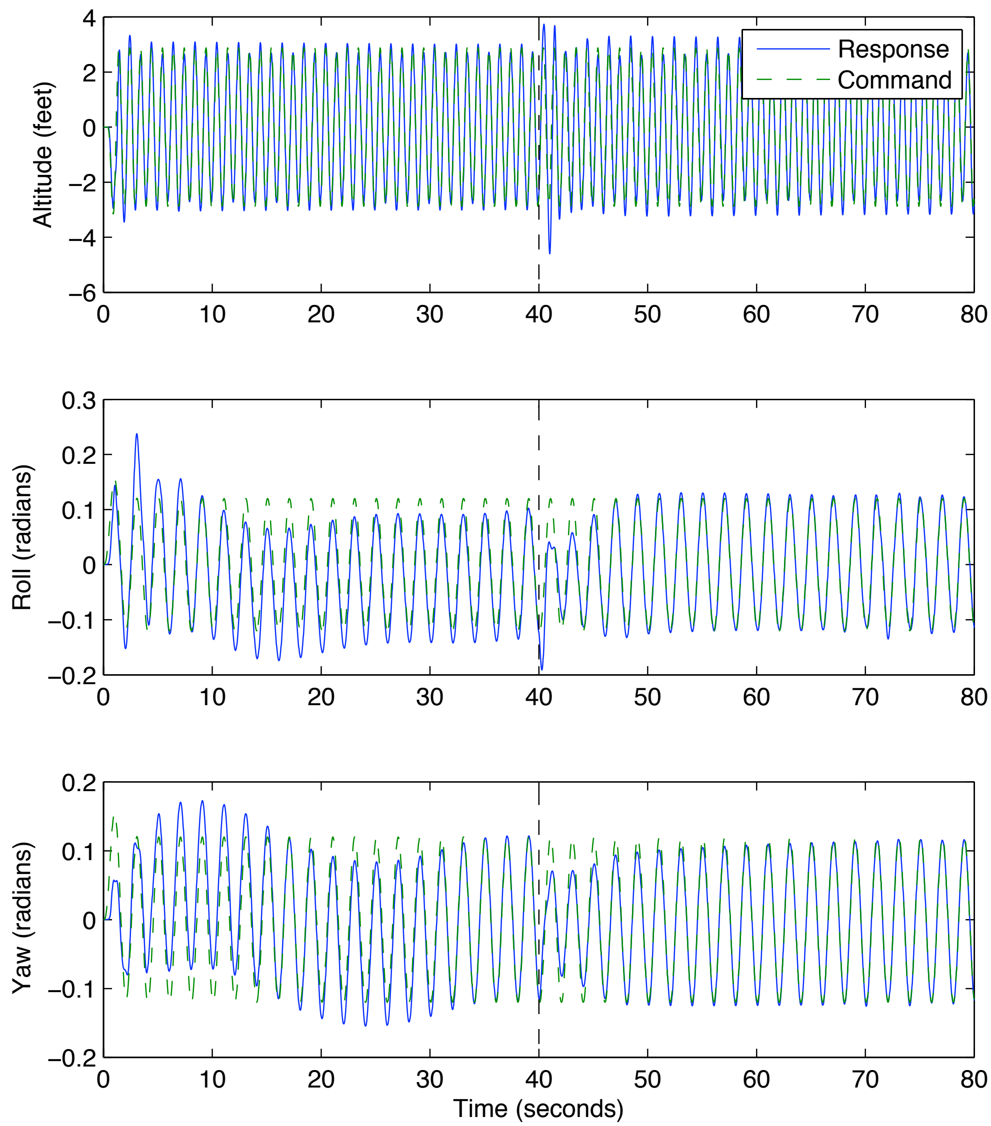

Figure 10. Altitude, Roll, and Yaw Sinusoidal Command Following when Actuator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow sinusoids in altitude, roll, and yaw. The elevator, ailerons, and rudder experience unknown damage at 40 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. The plots show that RCAC corrects for the damage.

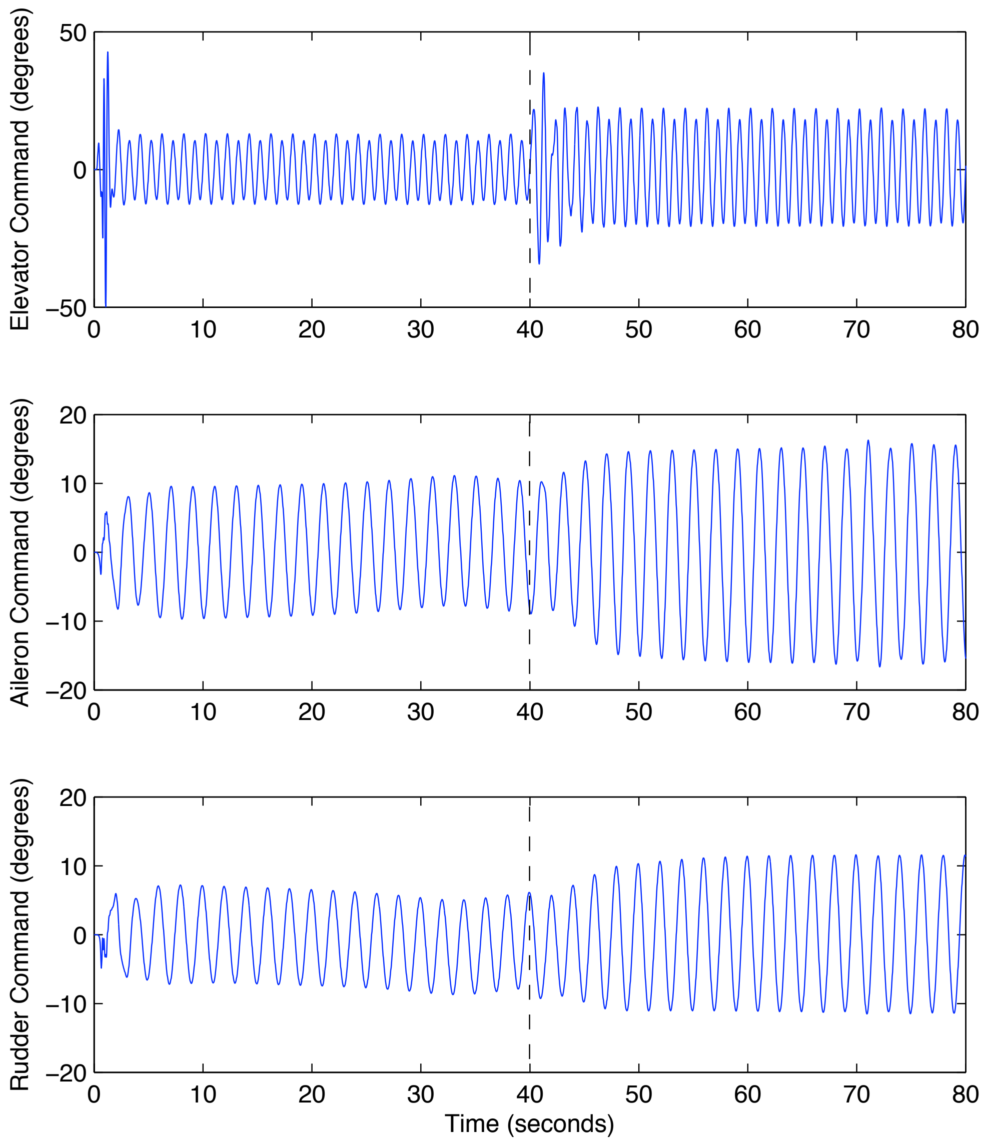

Figure 11. Altitude, Roll, and Yaw Sinusoidal Command Following Control Signals when Actuator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow sinusoids in altitude, roll, and yaw. The elevator, ailerons, and rudder experience unknown damage at 40 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. The magnitude of each control signal increases after 40 seconds to compensate for the damage.

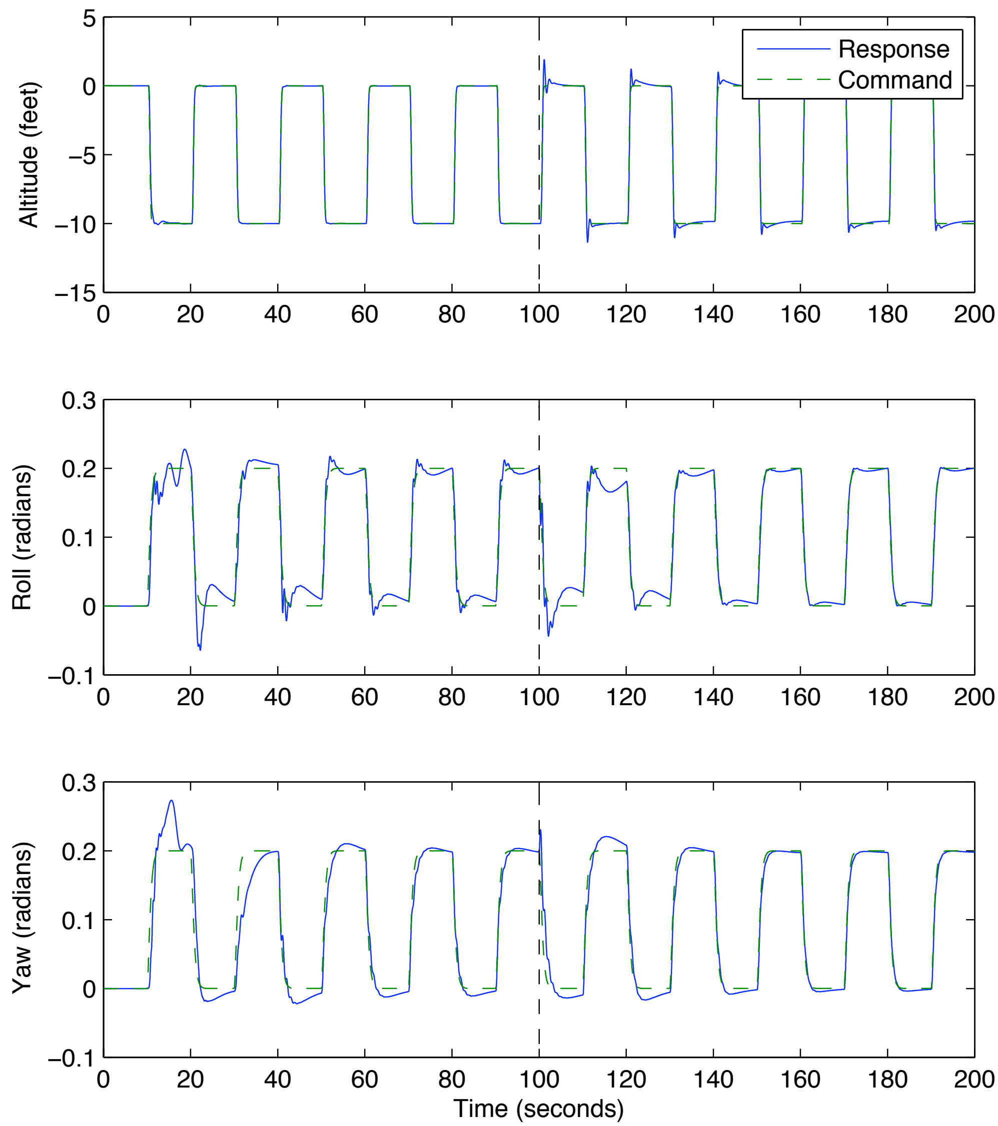

Figure 12. Altitude, Roll, and Yaw Doublet Command Following when Actuator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow sinusoids in altitude, roll, and yaw. The elevator, ailerons, and rudder experience unknown damage at 100 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. The plots show that RCAC corrects for the damage.

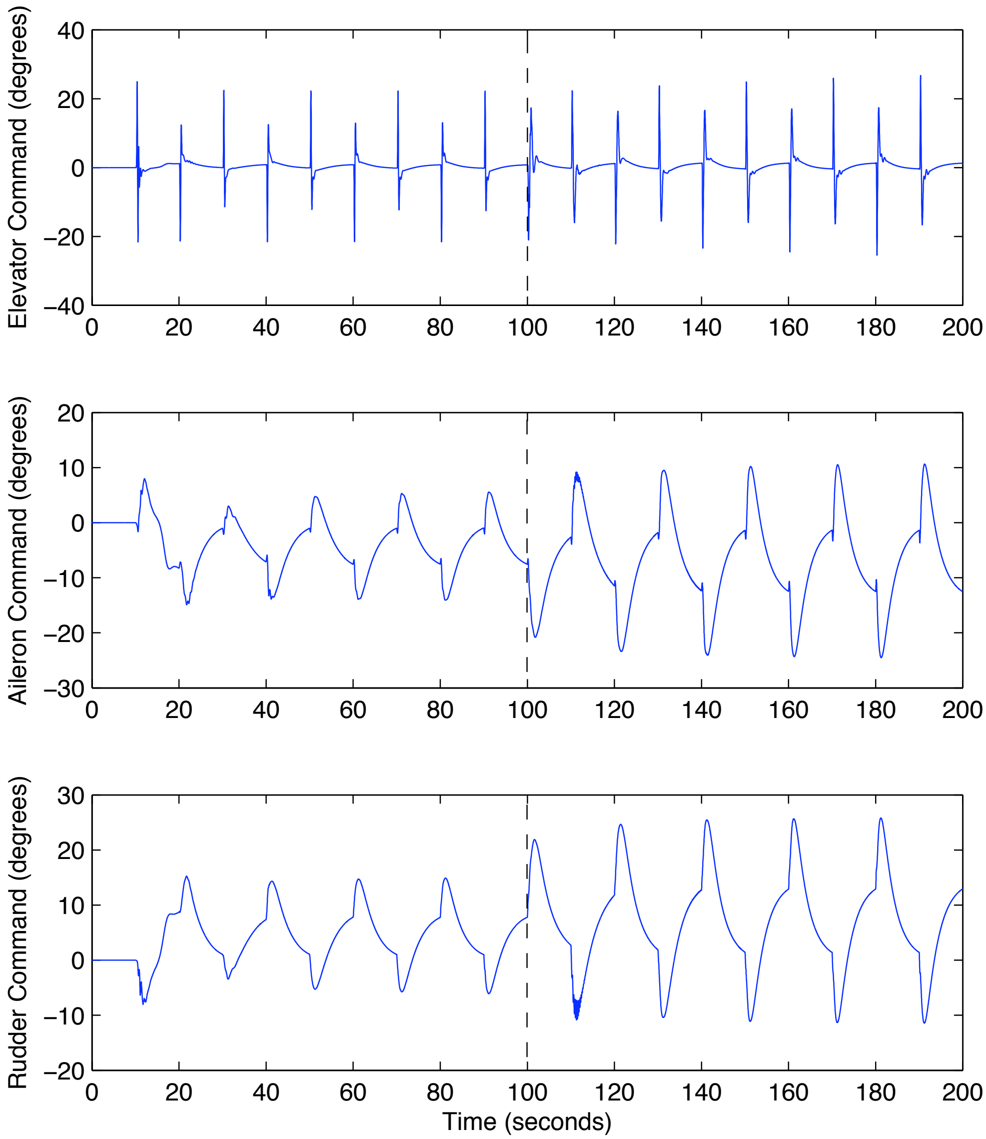

Figure 13. Altitude, Roll, and Yaw Doublet Command Following Control Signals when Actuator Effectiveness is Reduced to 60%: Model reference RCAC is used to follow sinusoids in altitude, roll, and yaw. The elevator, ailerons, and rudder experience unknown damage at 100 seconds such that the effectiveness of each control surface is reduced to 60% of its nominal effectiveness. The magnitude of each control signal increases after 100 seconds to compensate for the damage.

[10] Johansson, R., "Global Lyapunov stability and exponential convergence of direct adaptive control," *Int. J. Contr.*, Vol. 50, 1989, pp. 859–869.

[11] Akhtar, S. and Bernstein, D. S., "Logarithmic Lyapunov functions for direct adaptive stabilization with normalized adaptive laws," *Int. J. Contr.*, Vol. 77, 2004, pp. 630–638.

[12] Bayard, D. S., "Extended horizon liftings for stable inversion of nonminimum-phase systems," *IEEE Trans. Autom. Contr.*, Vol. 39, 1994, pp. 1333–1338.

[13] Venugopal, R. and Bernstein, D. S., "Adaptive disturbance rejection using ARMARKOV/Toeplitz Models," *IEEE Trans. Contr. Sys. Tech.*, Vol. 8, 2000, pp. 257–269.

[14] Santillo, M. A. and Bernstein, D. S., "Adaptive Control Based on Retrospective Cost Optimization," *AIAA J. Guid. Contr. Dyn.*, Vol. 33, 2010, pp. 289–304.

[15] Hoagg, J. B. and Bernstein, D. S., "Cumulative retrospective cost adaptive control with RLS-based optimization," *Proc. Amer. Contr. Conf.*, June 2010.

[16] Hoagg, J. B., Lacy, S. L., and Bernstein, D. S., "Broadband adaptive disturbance rejection for a deployable optical telescope testbed," *Proc. Amer. Contr. Conf.*, Portland, OR, 2005, pp. 4953–4958.

[17] Holzel, M. S., Santillo, M. A., Hoagg, J. B., and Bernstein, D. S., "Adaptive Control of the NASA Generic Transport Model Using Retrospective Cost Optimization," *Proc. AIAA Guid. Nav. Contr. Conf.*, August 2009, AIAA-2009-5616.

[18] Fledderjohn, M. S., Cho, Y.-C., Hoagg, J. B., Santillo, M. A., Shyy, W., and Bernstein, D. S., "Retrospective Cost Adaptive Flow Control Using a Dielectric Barrier Discharge Actuator," *Proc. AIAA Guid. Nav. Contr. Conf.*, August 2009, AIAA-2009-5857.

[19] Bailey, R. and Jordan, T., "NASA Langleys AirSTAR Testbed: A Subscale Flight Test Capability for Flight Dynamics and Control System Experiments," *roc. Guid. Nav. Contr. Conf.*, Honolulu, HI, August 2008, AIAA-2008-6660.

[20] Murch, A., "A Flight Control System Architecture for the NASA AirSTAR Flight Test Infrastructure," *roc. Guid. Nav. Contr. Conf.*, Honolulu, HI, August 2008, AIAA-2008-6990.