

Flight testing of a reconfigurable control system on an unmanned aircraft

David Shore and Marc Bodson

University of Utah

Abstract— A radio-controlled aircraft was built and equipped with air-data and inertial sensors. A radio frequency link was added to transmit data and receive commands from a ground station. Data from several flight tests was used to characterize the dynamic response of the aircraft. Despite the high level of noise associated with the low-cost sensor suite, consistent identification of critical aircraft parameters was obtained. Flight tests were also conducted with actuator failures induced on one elevator, one aileron, and one engine. Recursive parameter identification produced parameters tracking the effects of the failures, such as reduced effectiveness of pitch commands due to a locked elevator, or roll and sideslip due to engine failure. The identified parameters were also used in reconfigurable control experiments, where knowledge of the aircraft parameters was used to compensate for the effect of failures, reducing the pilot’s workload. Overall, the paper demonstrates that recursive identification and reconfigurable control algorithms are implementable in real-time, even in low-cost platforms. They can be designed to effectively compensate for actuator failures and aircraft damage.

NOMENCLATURE

- a_n , body-axis vertical accelerometer signal at the center of gravity
- a_y , body-axis lateral accelerometer signal at the center of gravity
- f_a , input signal that triggers an aileron fault
- f_e , input signal that triggers an elevator fault
- p , roll rate
- p_{com} , roll rate command
- q , pitch rate
- q_{com} , pitch rate command,
- \bar{q} , dynamic pressure
- r , yaw rate
- r_{com} , yaw rate command
- $taildrop$, input signal that triggers the tail-release servo
- thr_{left} and thr_{rt} , left and right throttle commands
- $thr_n = thr/100$, normalized symmetric throttle command
- v , velocity (airspeed)
- w , regressor vector containing measurable signals used for identification
- y , measurable signal used for identification
- α , angle of attack
- β , angle of sideslip
- δ_a , anti-symmetric aileron command
- $\delta_{a,left}$, left aileron command

- $\delta_{a,rt}$, right aileron command
- δ_e , symmetric elevator command
- $\delta_{e,left}$, left elevator command
- $\delta_{e,rt}$, right elevator command
- δ_r , rudder command
- θ^* , vector containing unknown parameters to be identified
- $\theta_{an,\alpha}$, lift coefficient
- $\theta_{an,bias}$, bias term in normal acceleration signal
- $\theta_{ax,qn\alpha^2}$, drag coefficient
- $\theta_{ax,thrn}$, thrust coefficient
- $\theta_{ax,bias}$, bias term in longitudinal acceleration signal
- $\theta_{ay,\beta}$, sideslip coefficient
- $\theta_{ay,bias}$, bias term in lateral acceleration signal
- $\theta_{d,p} = 6.0$, desired value for aileron effectiveness
- $\theta_{d,q} = -2.5$, desired value for elevator effectiveness
- $\theta_{q,el}$, elevator pitch rate effectiveness
- $\theta_{q,bias}$, bias term in pitch rate signal
- $\theta_{p,ail}$, aileron roll rate effectiveness
- $\theta_{p,bias}$, bias term in roll rate signal
- $\theta_{p,el}$, elevator roll rate effectiveness
- $\theta_{p,rud}$, rudder roll rate effectiveness
- $\theta_{\beta,ail}$, aileron sideslip effectiveness

- $\theta_{\beta,bias}$, bias term in sideslip signal
- $\theta_{\beta,r}$, roll rate to sideslip coefficient
- $\theta_{\beta,rud}$, rudder to sideslip coefficient

INTRODUCTION

Recent advances in computing power and electronic hardware have enabled the development of sophisticated programmable flight control systems, and a trend has been set to build such systems into modern aircraft. These aircraft can be designed with a much wider range of functional characteristics than was previously possible, and at the same time with customized control responses giving specific flying qualities.

A subject of recent research has been the design of reconfigurable control systems. Such systems automatically adapt their control strategy in the presence of failures or damages, providing benefits in terms of safety and reliability. Recent articles report data from flight tests of the VISTA F-16¹ and of the X40A², results of the RESTORE program^{3,4}, and simulations of the X-33 hypersonic vehicle⁵. Other studies of a variety of on-line identification and reconfigurable control algorithms can be found, among others, in Refs. 6-13. An interesting overview and comparison of several approaches appeared in Ref. 14.

Most evaluations of reconfigurable control algorithms have been performed in simulations. In the rare cases where flight tests were carried out, safety issues have mandated that failures be simulated, rather than implemented¹⁻³. The first objective of this project was to increase the available experimental data by conducting flight tests on a small radio-controlled aircraft. Such an aircraft offers several benefits to the researcher over a full-size aircraft. The radio link removes the test pilot from the hazardous environment of the experimental aircraft, permitting

tests that would pose unacceptable risks for an on-board pilot. The time-to-build and cost of a small unmanned aircraft are also far below those for manned aircraft, so that a broad range of flight tests can be performed, even with limited resources. In addition, tests can be performed with actual failures, rather than simulated ones.

A second objective of the project was to assess whether reconfigurable control strategies that were evaluated in simulations or in aircraft with high-quality sensors would work with cheap computers and noisy sensors. The angular rate measurements collected in this project were obtained with integrated circuits commonly found in digital cameras and R/C aircraft. The resulting sensor signals exhibited noise and drift significantly greater than those available in the VISTA research aircraft¹. Another source of imperfection was the actuators, which exhibited severe nonlinear dynamics, in particular a large deadzone with hysteresis.

Remote-controlled aircraft have been used for a number of guidance and control research projects. Kaminer and others at the Naval Postgraduate School have developed a “rapid flight test prototyping system”^{15,16}, using a medium-sized R/C aircraft that can be used for a range of guidance, navigation, and control research. Onboard the aircraft is a full sensor suite and a rather sophisticated flight control computer and autopilot capable of autonomous flight.

At Stanford University, the DragonFly UAV project^{17,18} involved the development of a model airplane with sophisticated onboard control capability. The work included the design of an autopilot toward the goal of developing a control system for a coordinated fleet of autonomous UAV’s able to accommodate communication failures while still completing their missions.

At West Virginia University, Napolitano and associates acquired scale models of the F-16, YF-22 and B777 as testbeds for autopilot and fault-tolerant flight control development¹⁹. Their

work has concentrated on detection and accommodation of sensor failures in the absence of sensor redundancy. Further work is expected to focus on detecting and accommodating actuator failures.

A team at the University of Michigan has developed a model aircraft (Solus), capable of intelligent flight control and trajectory planning, with a ground station that does high-level flight management²⁰. Their goal was to investigate the problem of detecting aircraft icing by changes in the dynamic response of the aircraft, and automatically re-planning the flight to accommodate reduced performance and remove the aircraft from the icing conditions. Note that this project and others cited above were concerned with developing sophisticated testbeds for work typically focused on autonomous flight, while the goal here was to demonstrate the applicability of on-line identification and reconfigurable control algorithms, keeping the on-board equipment to a minimum in order to allow risky tests.

The paper is organized as follows. The flight test platform is described in detail in the section titled “Aircraft Construction”. Next, a section titled “Batch Identification” discusses identification results obtained using a least-squares algorithm and data collected in various flights. The objective of this step of the project was to determine whether the signal to noise ratio of the data was sufficient to determine critical aircraft parameters and reliably identify the effect of failures. A positive answer was obtained, but the model had to be adjusted, and a suitable parameterization had to be used. Three actuator failures were simulated in flight tests: a frozen aileron, a frozen elevator, and a failed (idle) engine.

Next, a section titled “Recursive Identification” discusses a recursive implementation of the least-squares algorithm. The objective was to determine whether the critical aircraft parameters could be tracked in real-time. Such an implementation is generally more difficult because the

effective memory of the system must be made shorter than in batch experiments in order to obtain sufficiently fast adaptation in real-time. Fortunately, the results show that the identification algorithm produces estimates of actuator effectiveness that clearly reflect the effect of failures. Finally, the section titled “Reconfigurable Control” shows how the parameters estimated in real-time were combined with a simple control law in order to make the behavior of the aircraft invariant despite slow and abrupt changes in the parameters. The reconfigurable control strategy was found successful at relieving the pilot of the need to compensate for changes in aircraft behavior that resulted from failures. In addition to the previous failures, the system was also tested with a severe failure produced by releasing in flight the whole left side of the horizontal stabilizer/elevator. The paper completes with a “Conclusions” section that reviews the results of the project, and the conclusions that can be drawn from it.

AIRCRAFT CONSTRUCTION

System Architecture

Fig. 1. Radio-controlled model aircraft used for flight tests.

An “Almost-Ready-to-Fly” radio-controlled model airplane kit was selected. The Hobbico Twinstar model, shown in Fig. 1, had a twin-engine airframe made of balsa, with engines on the wings. This arrangement allowed the air-data sensors to be mounted on the nose of the fuselage, in the relatively undisturbed air upstream of the propellers. A twin engine aircraft also enabled flight tests with engine failures. The inertial sensors were mounted inside the fuselage, along with the data handling and telemetry hardware. The flight control system was implemented in a PC at the ground station. A diagram of the data processing system is shown in Fig. 2.

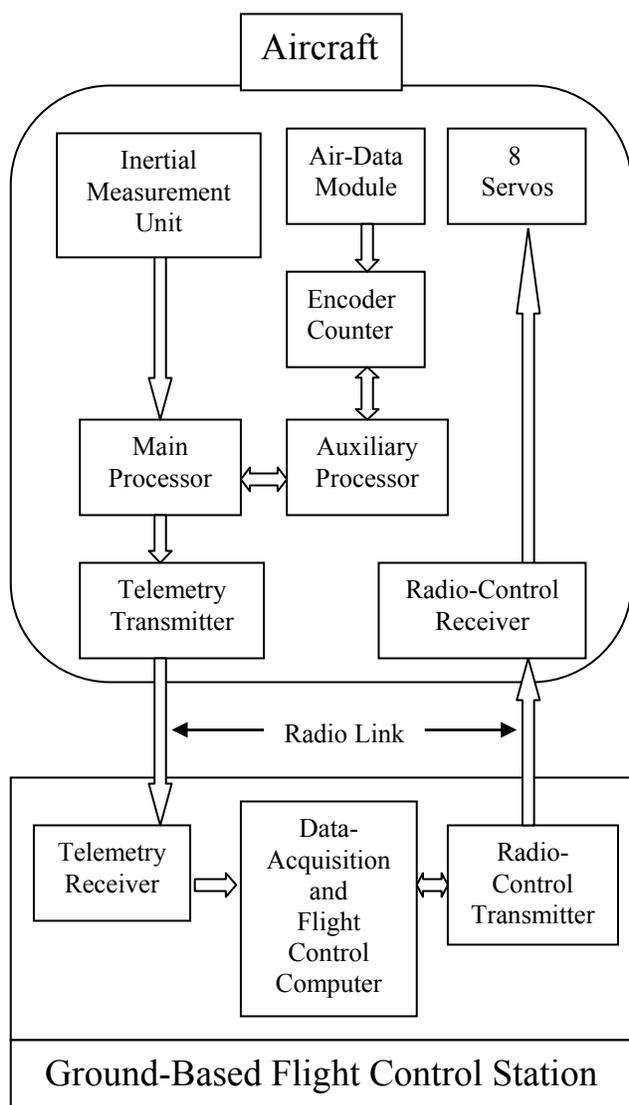


Fig. 2. Diagram of the data processing system.

Air-Data Sensors

Three sensors were assembled into the custom air-data module shown in Fig. 3. For airspeed measurement, a pitot-static tube was connected to a differential-pressure transducer (Omega PX163-2.5BD5V). Speeds up to 114 feet per second could be measured. The wind angles were

measured by optical encoders (US Digital E2-1024) that were assembled with vanes on both ends. The signals were decoded by a 2-axis counter chip (LSICSI LS7266R1).

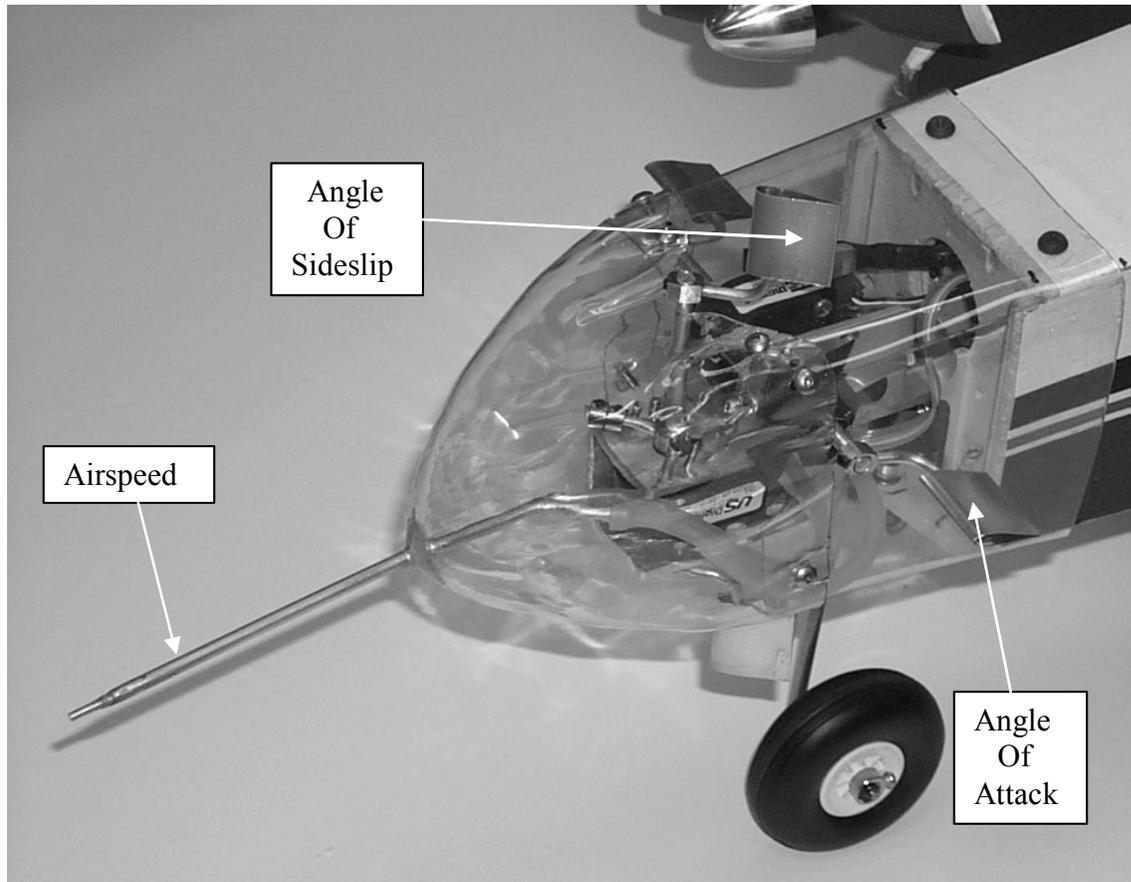


Fig. 3. Air-data module, mounted in the nosecone.

Inertial Measurement Unit

Composing the inertial measurement unit (IMU), gyroscopes (Murata Gyrostar ENC-030JA) and accelerometers (Analog Devices ADXL105) for the three body axes were assembled on a circuit board.

Sensor Calibration

The airspeed and wind-angle sensors were tested in a wind tunnel for dynamic effects and

accuracy. It was discovered that the nosecone caused the direction of the wind vector to change slightly in the vicinity of the vanes. This error was corrected in the decoding algorithm. Dynamic effects were negligible in flight data. There were no significant dynamic effects in the inertial sensors.

Data Acquisition and Telemetry

The main microprocessor board had a NetMedia BX-24 microprocessor, which included on-chip 10-bit analog-to-digital converters (A/D) to read the airspeed sensor and the six IMU sensors. The encoder counter board had a 2-axis counter chip programmed and polled by an auxiliary BX-24 microprocessor. The main microprocessor board encoded the data and sent it as a serial stream to the telemetry transmitter. The objective of encoding the data was to create a protocol of transmission such that transmission errors would only yield short gaps in the data. A thorough discussion of the data handling methods is given in Ref. 21. The telemetry transfer was handled by a MaxStream 100 mW, 900 MHz spread-spectrum digital transceiver pair with built-in data buffering and error checking. The transmission was modulated at 19.2 kbps and this limited the rate of data transmission, and hence the rate of the control system, to 96 Hz.

Actuator Calibration

The pilot's control inputs were measured at the joystick potentiometers, and fed directly into the flight control computer. Actuator positions were not measured, which reduced the on-board hardware and the number of telemetry channels, but introduced errors associated with the actuators' dynamic behavior. To assess the errors between actuator commands and responses, a test jig was made to measure the control surface position with a potentiometer. Tests were run for each control surface in which the position and the corresponding joystick voltage were measured.

Several actuator characteristics were identified in the responses, including a delay, approximately first-order dynamics, a rate limitation, a hysteresis, and some degree of asymmetry. In order to compensate for the first two characteristics, an approximate actuator model was created, with a time delay of 8 steps, and a first-order time constant of 0.05 second. The remaining error constituted unmodeled dynamics, to which the parameter estimation algorithm had to be robust.

Ground Station and Remote Control

The flight test system was designed for minimal on-board data processing, in order to limit aircraft weight and power requirements. On-line parameter identification and control were performed by a PC at the ground station. The PC used a Pentium III 500 MHz processor, and a data acquisition card (Data Technologies DT-2801). This DAQ card included eight A/D converters and two D/A converters. A separate D/A chip (Burr-Brown DAC7625) was addressed by the digital IO on the DAQ card to provide the three additional output channels that were required.

The aircraft's R/C system consisted of a Futaba 8-channel computer radio (R/C transmitter/receiver) controlling eight servos on the aircraft. The servos drove two throttles, two ailerons, two elevators, a rudder, and a mechanism to release the left stabilizer/elevator.

Five of the eight R/C channels (throttle, elevator, aileron, rudder, and fault indicator) were read by A/Ds in the DAQ system of the flight control computer. There, the signals were modified and five servo commands (left elevator, right elevator, left aileron, right aileron, rudder) were sent to the airplane.

Failure Implementation

The following four in-flight failures were planned, and were triggered by a switch on the

pilot's flight control box.

- *Frozen left elevator.* The airframe was modified by splitting the single full-width elevator into two independent half-width elevators, each with its own servo and R/C channel. For an elevator failure, the left elevator was commanded to its neutral position and the right elevator received the usual elevator command.
- *Frozen left aileron.* Each aileron was fitted with a separate servo and R/C channel. For an aileron failure, the left aileron was commanded to its neutral position and the right aileron received the usual aileron command.
- *Left or right engine failure.* The separate throttle servos were given separate R/C channels. For an engine failure, one engine was commanded to a low (idle) throttle setting.
- *Separation of left stabilizer/elevator.* An R/C channel and servo were connected to a mechanism that released the entire left half of the horizontal tail.

BATCH IDENTIFICATION

Linear Parameterization

The simplest and most effective methods of identification of dynamic systems are based on representations of these systems involving linear parameterizations. Specifically, a linear parameterization is defined by a vector product

$$y = \theta^{*T} w \quad (1)$$

For identification, measurements of $y(k)$ and $w(k)$ are taken, which are the values sampled at some period T , and $k = 1, \dots, n$. The problem is to find a parameter estimate $\theta(n)$ satisfying the linear relationship (1) as closely as possible. Because of measurement noise, and because the linear relationship is not an exact representation of the system, an estimate must be found such

that the best fit is obtained, using a sufficiently large number of data points.

For the real-time estimation of aircraft parameters, the following linear parameterization was used in Ref. 13.

$$\begin{aligned}
 y_1 = a_n &= \theta_{11} \bar{q} \alpha + \theta_{12} \bar{q} \\
 y_2 = \dot{q} &= \theta_{21} \bar{q} \alpha + \theta_{22} \frac{\bar{q}}{v} q + \theta_{23} \bar{q} + \theta_{24} \delta_e \\
 y_3 = a_y &= \theta_{31} \bar{q} \beta + \theta_{32} \frac{\bar{q}}{v} p + \theta_{33} \frac{\bar{q}}{v} r + \theta_{34} \bar{q} + \theta_{35} \delta_a + \theta_{36} \delta_r \\
 y_4 = \dot{p} &= \theta_{41} \bar{q} \beta + \theta_{42} \frac{\bar{q}}{v} p + \theta_{43} \frac{\bar{q}}{v} r + \theta_{44} \bar{q} + \theta_{45} \delta_e + \theta_{46} \delta_a + \theta_{47} \delta_r \\
 y_5 = \dot{r} &= \theta_{51} \bar{q} \beta + \theta_{52} \frac{\bar{q}}{v} p + \theta_{53} \frac{\bar{q}}{v} r + \theta_{54} \bar{q} + \theta_{55} \delta_e + \theta_{56} \delta_a + \theta_{57} \delta_r
 \end{aligned} \tag{2}$$

The fundamental dependency of the forces and moments on the dynamic pressure and on the speed of the aircraft, known from aerodynamic theory, was taken into account in the formulation. Note that a similar linear parameterization was successfully used in the identification of the VISTA F-16 aircraft in Ref. 1.

When the linear parameterization (2) was used for the identification of the parameters of the R/C aircraft, significant difficulties were encountered in the estimation of the second, fourth, and fifth equations. Indeed, the rotational accelerations \dot{q} , \dot{p} , and \dot{r} , could not be obtained reliably. In Ref. 1, these accelerations were computed by subtracting the signals from accelerometers placed fore and aft in the aircraft. The rotational rate signals were also sufficiently clean that rotational accelerations could have been obtained by filtered differentiation of the rate signals.

For the R/C aircraft, neither approach was applicable, due to the short length of the fuselage

and to the high level of noise in the accelerometer and gyroscope signals. It also became apparent that the rotational accelerations of the R/C aircraft were small ($\dot{q} \approx 0$, $\dot{p} \approx 0$, and $\dot{r} \approx 0$). This property was attributed to the high level of inherent stability of the R/C aircraft, which had large stabilizing surfaces on the rear of the aircraft, and a wide wing span.

An alternative approach was therefore pursued for the linear parameterization. Specifically, the rotational accelerations were assumed to be zero, and the aircraft model was rearranged so that q , p , and β appeared on the left-hand side of the equations. The resulting equations were obtained

$$a_n = \theta_{a_n, \alpha} (q_n \alpha) + \theta_{a_n, bias} \quad (10)$$

$$q = \theta_{q, el} (v_n \delta_e) + \theta_{q, bias} (10v_n)$$

$$a_y = \theta_{a_y, \beta} (q_n \beta) + \theta_{a_y, bias} \quad (10) \quad (3)$$

$$p = \theta_{p, ail} (v_n \delta_a) + \theta_{p, bias} (10v_n) + \theta_{p, el} (v_n \delta_e) + \theta_{p, rud} (v_n \delta_r)$$

$$\beta = \theta_{\beta, rud} (\delta_r) + \theta_{\beta, ail} (\delta_a) + \theta_{\beta, r} \left(\frac{r}{v_n} \right) + \theta_{\beta, bias} \quad (10)$$

Scaling factors were included in the equations so that the signals would have roughly the same magnitude. In particular, a normalized dynamic pressure was defined as $q_n = \bar{q}/5$, and a normalized velocity as $v_n = v/50$. Note that the bias terms for the accelerometer signals do not include the dynamic pressure as in (3). Slightly better results were obtained with this parameterization, perhaps because of sensor biases.

A parameterization for longitudinal acceleration was also considered, in order to estimate the effects of drag and the response to engine thrust. Specifically, the longitudinal accelerometer signal was fitted to the relationship

$$a_x = \theta_{a_x, \alpha^2} \left(\frac{q_n}{100} \alpha^2 \right) + \theta_{a_x, thr} (thr_n) + \theta_{a_x, bias} (q_n) \quad (4)$$

Each of the equations of (3) and (4) is an independent relation that is a special case of the linear parameterization

$$y = \theta^{*T} w \quad (5)$$

Taking, as an example, the fifth equation of (3), and putting it into the form of (5), the variables are defined as

$$y = \beta \quad \theta^* = \begin{bmatrix} \theta_{\beta, rud} \\ \theta_{\beta, ail} \\ \theta_{\beta, r} \\ \theta_{\beta, bias} \end{bmatrix} \quad w = \begin{bmatrix} \delta_r \\ \delta_a \\ \left(\frac{r}{v_n} \right) \\ 10 \end{bmatrix} \quad (6)$$

Least-Squares Algorithm

The problem of batch identification is to find $\theta(n)$ such that the relation (5) is satisfied. Because the signals y and w contain the contributions of unmodeled dynamics, nonlinearities, and measurement noise, a best fit must be obtained for the measurements $y(k)$ and $w(k)$. For identification of a batch of data, define

$$W = \begin{bmatrix} w(1)^T \\ w(2)^T \\ \vdots \\ w(n)^T \end{bmatrix}, \quad Y = \begin{bmatrix} y(1) \\ y(2) \\ \vdots \\ y(n) \end{bmatrix} \quad (7)$$

W is a matrix of n measurements of the vector w^T and Y is a vector. To obtain the best fit, a least-squares optimization criterion is used

$$J(\theta) = \|Y - W\theta\|^2, \quad (8)$$

whose solution is

$$\theta = [W^T W]^{-1} [W^T Y]. \quad (9)$$

Flight Test Results

Nine representative sections of data, corresponding to roughly two minutes each, were selected for a closer analysis. Different fault conditions and some low power conditions were included, as stipulated in Table I. The flight patterns used in the flight tests were similar for all flights. The pilot used patterns typical of R/C flight, consisting mostly of figure eights with some climbs and descents.

Table I. Conditions of nine sections of flights.

Section:	1	2	3	4	5	6	7	8	9
Normal Flight	X	X	X						
Elevator Fault				X	X				
Aileron Fault						X	X		
Low Power								X	X

Off-line parameter identification using the least-squares algorithm (9) was performed on each section of data, and the results are shown in Fig. 4. Eight significant parameters are shown, each with its own subplot, and the nine sections of flight data are labeled along the x-axes.

Subplot (a) shows the lift coefficient, $\theta_{an,\alpha}$, which is consistently around 0.14 for the first eight sections. The low value at section 9 could be due to decreased lift because of flow separation near a stall condition, since the angle-of-attack was as much as 20.3 degrees in that section.

Subplot (b) shows the elevator pitch effectiveness, $\theta_{q,el}$, whose value is consistently around -

2.6 in the flight sections with no elevator fault (1, 2, 3, 6, 7, and 8). In sections 4 and 5, where only one of the two elevators was actuated, the value is about half the normal value, as expected. Section 9 did not have an elevator fault, but shows reduced effectiveness, possibly because of reduced propeller wash and a high- α flight condition.

Subplot (c) shows the aileron roll effectiveness, $\theta_{p,ail}$, which exhibits properties similar to the elevator pitch effectiveness. Unfailed conditions show a value of 6.6, while sections 6 and 7, where only one of the two ailerons was actuated, show a value reduced in half.

Subplot (d) shows the *elevator* roll-effectiveness, $\theta_{p,el}$. Sections 4 and 5, which have asymmetric elevator (with only the right side working), give a small number with negative sign, as expected. However, the other sections, in which elevator deflections are symmetric, should be closer to zero. It was concluded that the parameter $\theta_{p,el}$ could not be estimated reliably.

Subplot (e) shows the *rudder* roll-effectiveness, $\theta_{p,rud}$. The rudder's force above the center of gravity is such that a positive deflection should impose a positive moment about the longitudinal body-axis. The numbers are mostly negative, though, and are most negative in sections 8 and 9, with high angle-of-attack.

Subplots (f), (g), and (h) are associated with the sideslip (β) which, as expected, is dominated by the rudder. In (f), $\theta_{\beta,rud}$ has a steady value, except in section 9, where low rudder effectiveness may be explained by minimal propeller wash from idling engines. In (g), the aileron-induced sideslip coefficient, $\theta_{\beta,ail}$, is smallest when only one aileron is working (sections 6 and 7). With the aileron failure, the associated drag of aileron deflection is less, and sideslip is accordingly smaller. It is largest with low power/high α (sections 8 and 9). This makes sense in light of the dependency of drag on α^2 . $\theta_{\beta,ail} > 0$ in (g), together with $\theta_{\beta,r} > 0$ in (h), show that a rolling maneuver causes adverse yaw, which must be compensated by rudder to coordinate a turn.

Overall, the results show that the two critical control parameters, namely the pitch effectiveness of the elevators, and the roll effectiveness of the ailerons, could be reliably identified. Their values are affected by failures in a manner consistent with expectations. For some parameters, significant variations are observed that may be explained by the specific flight condition (low speed, high angle-of-attack). Some parameters can only be estimated with marginal or inconsistent results. Fortunately, these parameters represent small effects that do not significantly affect the control performance of the aircraft.

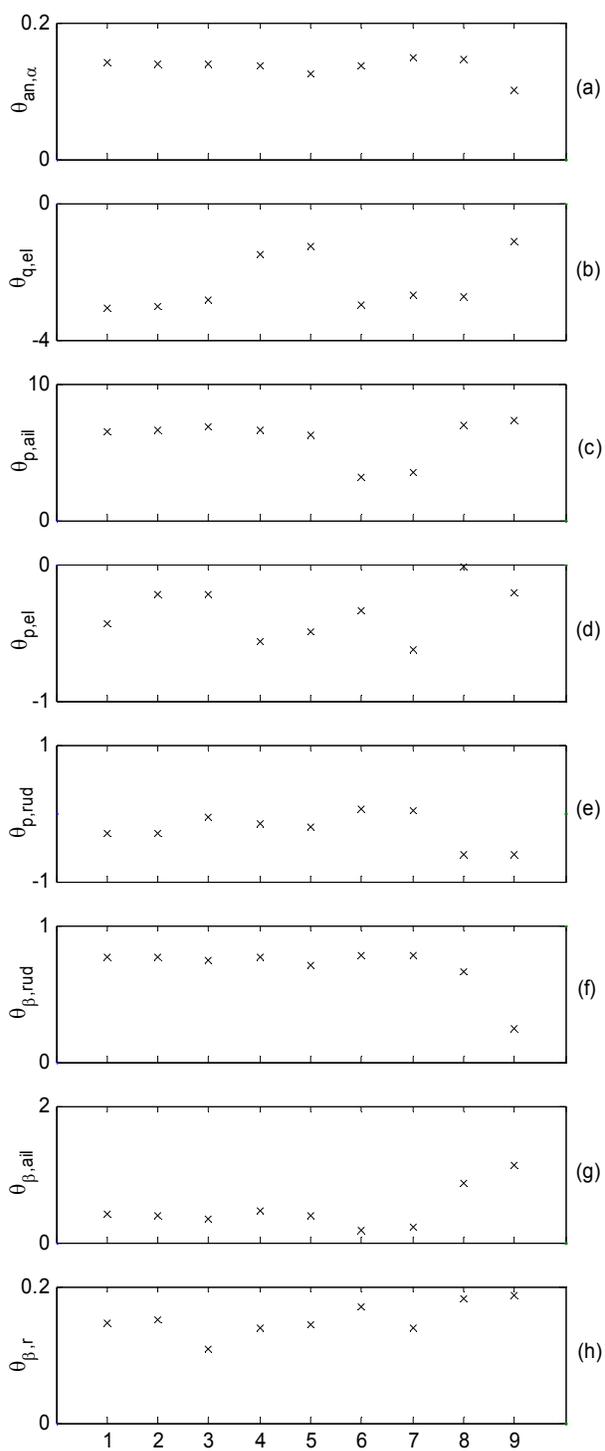


Fig. 4. Parameters from batch identification of nine sections of selected flights. Sections 1-3 have no faults, 4-5 have an elevator fault, 6-7 have an aileron fault, 8-9 have low power.

Data Fit

The quality of the fit of the estimated model to the flight data is shown in this section for a representative data segment. Given a parameter estimate θ , the signal y was compared to the signal

$$y_{fit} = W\theta. \quad (10)$$

A 40-second window of data is displayed in Fig. 5. Flight data is plotted with solid lines, and the fit is shown with dotted lines. Fig. 5 shows that the data fit for a_n , in (a), and for q , in (b), is very good. The angle of attack (alpha), in (c), and the elevator deflection, in (d), are the inputs to the identification algorithm.

The same data window is displayed in Fig. 6, which shows a very good fit of the p channel in (a). The fit of a_y , in (b), is less accurate. This fit depends on two signals, a_y and β , that are both small, so that the signal-to-noise ratio is not good. Fig. 6 also shows a fair fit of the β channel in (c), and the inputs, aileron (d) and rudder (e).

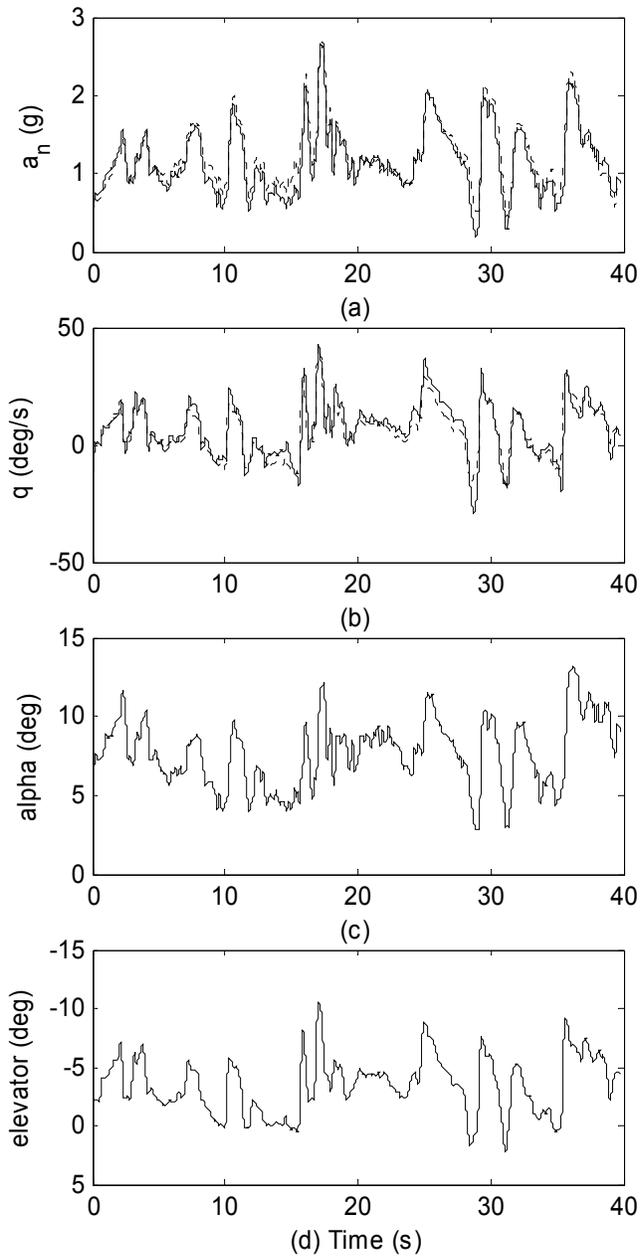


Fig. 5. The fit of the estimated model (dotted lines) to the data (solid lines) for the longitudinal axis.

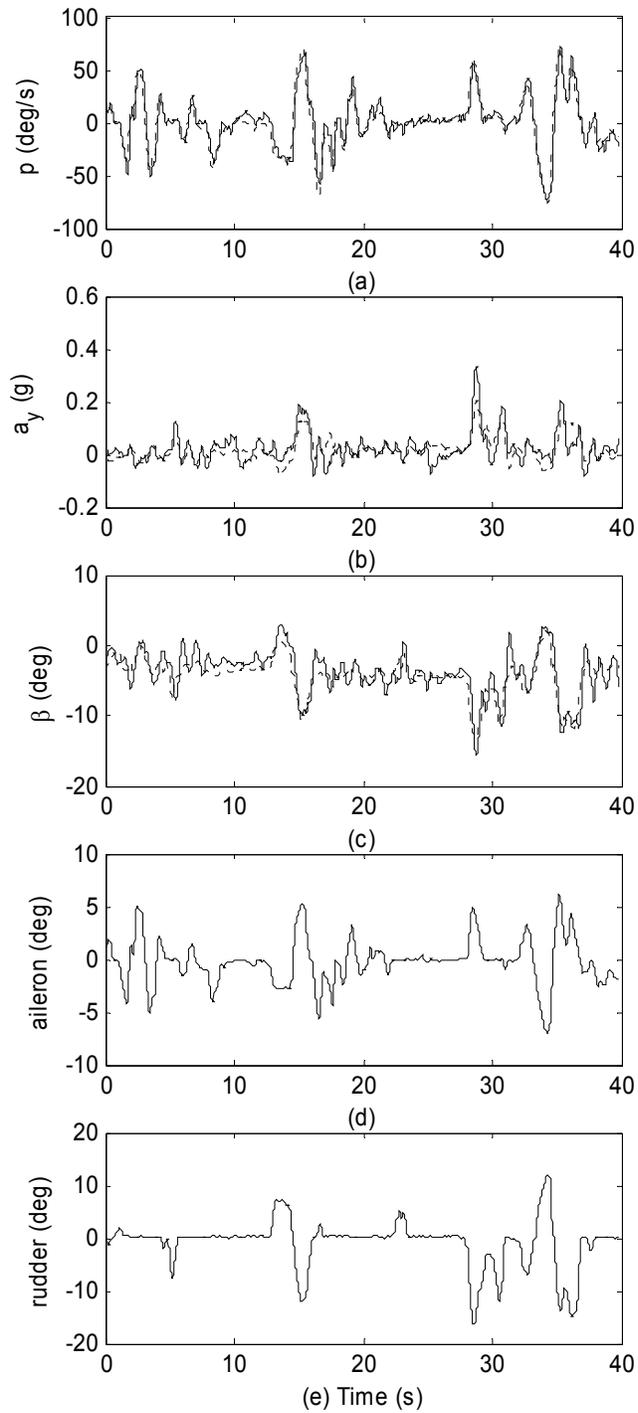


Fig. 6. The fit of the estimated model (dotted lines) to the data (solid lines) for the lateral axis.

RECURSIVE IDENTIFICATION

Recursive Least-Squares Identification

In the previous section, data sets were treated as batches, and parameters were computed that gave the best fit over each set. This approach is useful for model building and for validation, but it assumes that the parameters remain constant over the length of the data set. To track variations associated with systemic changes such as mechanical damage and component failures, the parameters must be continuously updated during regular operation. The desired result can be achieved by using a modified version of the least-squares optimization criterion (8) that is suitable for adaptation. A useful criterion is¹

$$J(\theta(n)) = \sum_{k=1}^n \lambda^{n-k} [y(k) - \theta^T(n)w(k)]^2 + \alpha_w \|\theta(n) - \theta(n-1)\|^2 \quad (11)$$

where λ is known as the “forgetting factor”, and is used to discount old measurements (thereby allowing parameter estimates to change, based on recent data). λ is set to a value $0 < \lambda \leq 1$. Choosing λ close to zero corresponds to the greatest ability to track rapid changes, such as damages and failures, because only the latest data points significantly affect the estimate. Choosing λ close to 1 corresponds to a more slowly adapting algorithm, and to a greater robustness to noise. The second term of (11) is used to ensure the stability of the algorithm, which is needed when the requirement for persistent excitation is not met, such as during cruise conditions. It limits the deviation of the current estimate from the previous estimate, and its influence is adjusted with the factor α_w .

The recursive formulation of (11) is known as the *stabilized recursive least squares algorithm with forgetting factor*¹. An approximate but accurate implementation of the algorithm¹³ includes the update of the so-called covariance matrix

$$\begin{aligned}
P(n) &= \frac{1}{\lambda} P(n-1) \\
&- \frac{1}{\lambda} P(n-1) C(n) (\lambda I + C(n)^T P(n-1) C(n))^{-1} C(n)^T P(n-1)
\end{aligned} \tag{12}$$

where the matrix $C(n)$ is given by

$$C(n) = \left(w(n) \quad \sqrt{n_p \alpha_w (1-\lambda)} \cdot e(n) \right) \tag{13}$$

and $e(n)$ is a sequence of vectors

$$e(1) = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, e(2) = \begin{bmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{bmatrix}, \dots, e(n_p) = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 1 \end{bmatrix}, e(n_p + 1) = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \dots \tag{14}$$

where n_p is the number of parameters in θ , and the number of variables in w . $C(n)$ is a $n_p \times 2$ matrix and $e(n)$ is a $n_p \times 1$ vector. $P(n)$ is initialized to $P(0) = (1/\alpha_w)I$ and used in the parameter update law

$$\begin{aligned}
\theta(n) &= \theta(n-1) + P(n) w(n) (y(n) - w(n)^T \theta(n-1)) \\
&+ \alpha_w \lambda P(n) (\theta(n-1) - \theta(n-2))
\end{aligned} \tag{15}$$

Further discussions of this algorithm and related topics are given in Refs. 1 and 13. The algorithm (12)-(15) was tested on many flights with and without failures, during which the factors $\lambda = 0.998$ and $\alpha_w = 1000$ were chosen. The forgetting factor corresponds to a time constant of 500 samples or 5 seconds.

Fig. 7 shows the results from one flight without failure. Recursive estimates for the primary parameters are plotted in solid lines, the batch estimates of the primary parameters are plotted in dashed lines, and the estimates of the bias parameters are plotted in dotted lines. On subplot (a), the primary parameter is the lift coefficient, $\theta_{an,\alpha}$. The batch estimate of 0.142 agrees well with the recursive estimate. The bias parameter, $\theta_{an,bias}$, is the vertical acceleration (in 0.1g) when the

angle of attack is zero. Subplots (b), (d), and (e) show the parameters associated with the lateral acceleration, the pitch rate, and the roll rate. The recursive estimates are close to the batch estimates for most of the flight. Subplot (c) shows the drag coefficient, $\theta_{ax, qn\alpha^2}$, in solid lines. Here, the recursive value is slow catching up with the batch estimate, but this is not an important parameter for control. Subplot (f) shows rudder effectiveness, $\theta_{\beta, rld}$, with a sudden jump at about $t=30$ seconds. This jump corresponds to the only rudder input during the flight, and shows the need for “sufficient excitation” in parameter identification. In general, it was found that identification of the directional (sideslip) channel was difficult because of the rare use of the rudder in normal R/C flight.

Fig. 8 shows the estimated values of the effectiveness of the aileron and of the elevator during two flights with actuator faults. In each plot, the upper curve is the aileron effectiveness and the lower curve is the elevator effectiveness. In the middle of each plot is the fault indicator line. The fault occurs where the line switches from 0 to 2. In both plots, the parameter associated with the failed actuator shows the expected change to half its no-fault value during the fault, and then a return to its normal value when the fault ends.

A notable feature of the plots is the large fluctuation of the parameters in normal conditions. Some of the variations, particularly the spikes, are thought to be caused by the dead zone mentioned in the actuator calibration section. The variation in the estimates could be partly smoothed out by increasing the forgetting factor λ , or by increasing the stabilization factor α_w , but the cost would be slower adaptation to failures. The choices of λ and α_w used here resulted in convergence of $\theta_{q,el}$ in 6-12 seconds, and convergence of $\theta_{p,ail}$ in 3-9 seconds after a fault. The variations in convergence times are due to differences in excitation of the identified channels.

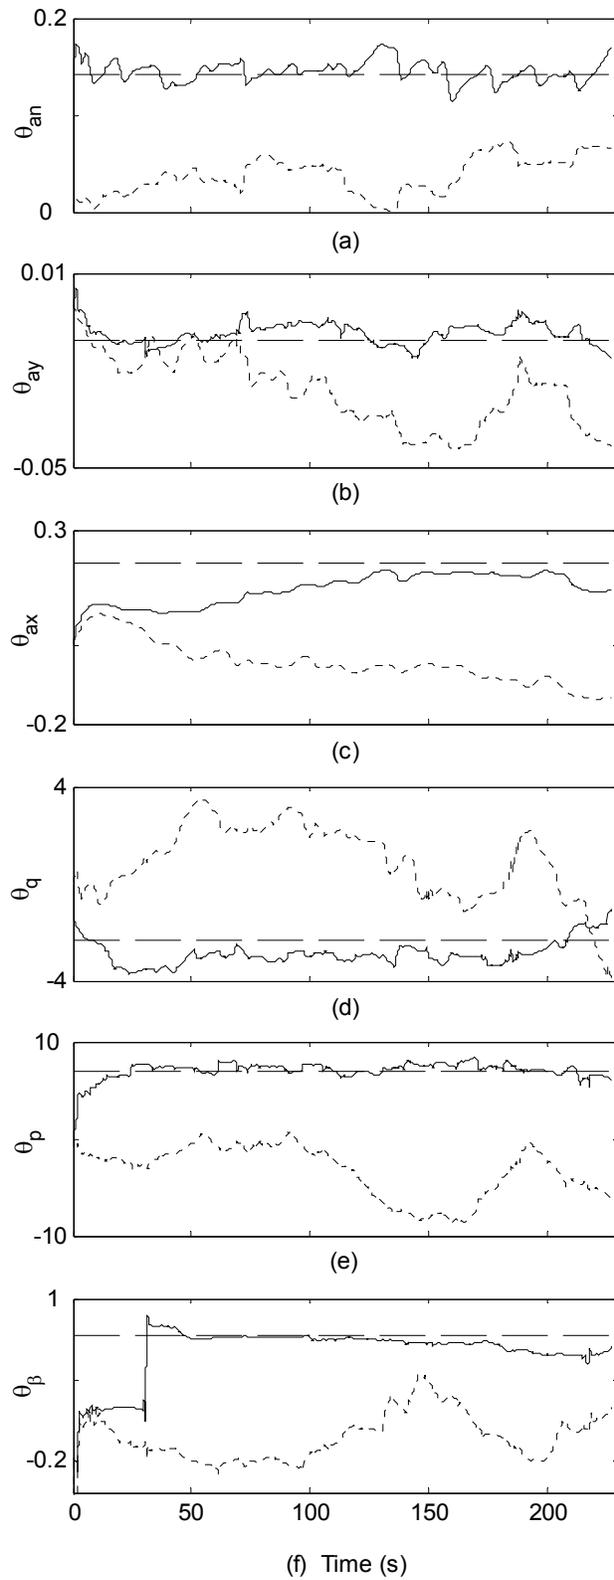


Fig. 7. Recursive identification on a flight with no failures. Primary parameters are plotted in solid lines, with their batch estimates in dashed lines. Bias parameters are shown in dotted lines.

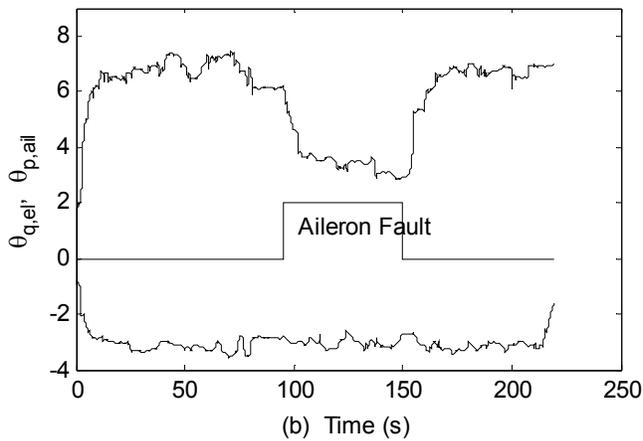
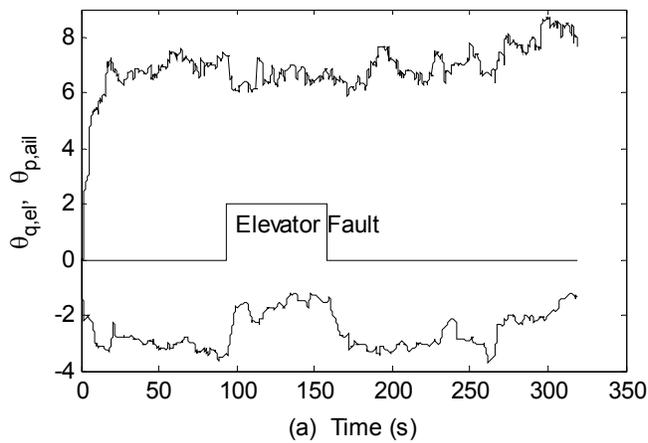


Fig. 8. Results of recursive identification for flights with elevator fault, (a) and aileron fault, (b).

RECONFIGURABLE CONTROL

Control Reconfiguration

Having established that the effect of the failures could be detected in batch and recursive processing of the aircraft data, the final step of the project was to use the estimates in real-time to reduce the impact of failures on the aircraft characteristics seen by the pilot. The reconfigurable control algorithm discussed here can be viewed as a special case of *model reference adaptive control*. In general, model reference adaptive control attempts to modify the closed-loop

dynamics of a system so that its responses track those of a “desirable” system called the reference model²². In off-line identification experiments, it was found that the behavior of the R/C aircraft could be simplified to the point where the dynamics from the control variables (δ_e , δ_a , and δ_r) to the output variables (q , p , and β) reduced to simple gains with biases. In the presence of failures, as well as in some low speed conditions, these gains were found to change significantly. The reconfigurable control law discussed in this section consists in applying gains to the pilot commands and biases so that the overall gains of the system remain constant and equal to some desirable values.

The gain parameter and the bias parameter were estimated each for pitch rate, roll rate, and sideslip. The variables used in the recursive identification algorithm were

$$\begin{aligned}
 y = q, \quad \theta &= \begin{bmatrix} \theta_{q,el} \\ \theta_{q,bias} \end{bmatrix}, \quad w = \begin{bmatrix} v_n \delta_e \\ v_n \end{bmatrix}, \\
 y = p, \quad \theta &= \begin{bmatrix} \theta_{p,ail} \\ \theta_{p,bias} \end{bmatrix}, \quad w = \begin{bmatrix} v_n \delta_a \\ v_n \end{bmatrix}, \\
 y = \beta, \quad \theta &= \begin{bmatrix} \theta_{\beta,rud} \\ \theta_{\beta,bias} \end{bmatrix}, \quad w = \begin{bmatrix} \delta_r \\ 1 \end{bmatrix},
 \end{aligned} \tag{16}$$

With these definitions, the control law consisted in letting

$$\begin{aligned}
 \delta_e &= q_{com} \cdot \frac{\theta_{d,q}}{\theta_{q,el}} - \frac{\theta_{q,bias}}{\theta_{q,el}}, \\
 \delta_a &= p_{com} \cdot \frac{\theta_{d,p}}{\theta_{p,ail}} - \frac{\theta_{p,bias}}{\theta_{p,ail}}, \\
 \delta_r &= r_{com} - \frac{\theta_{\beta,bias}}{\theta_{d,r}}.
 \end{aligned} \tag{17}$$

Since difficulties were encountered in reliably identifying the parameters of the directional channel, adaptive compensation was disconnected for the sideslip channel in the experiments reported in this section.

Some small modifications were added to the control law (17). Since the estimated gains could get very small during ground operations, the bounds $\theta_d/3 \leq \theta_1 \leq 2\theta_d$ were imposed on the estimates before they were used in (17). These bounds were imposed only on the primary parameters (θ_1) and not on the bias parameters (θ_2). Additionally, the commands δ_e , and δ_a were limited to their allowable ranges. More sophisticated command limiting techniques are discussed in Ref. 12.

A block diagram of the reconfigurable control system is shown in Fig. 9. The purpose of the “Filter and Delay” block was to align the control signals in u with the delayed output signals in y_{pid} . The control signals propagate through the PWM encoding/decoding and the mechanical response of the actuators, resulting in delay. The signal misalignment is also aggravated because the output signals were delayed during on-board data acquisition, transmission, and decoding in the flight control program. Overall, the Filter and Delay block applied a filter to the control variables that matched the total delay.

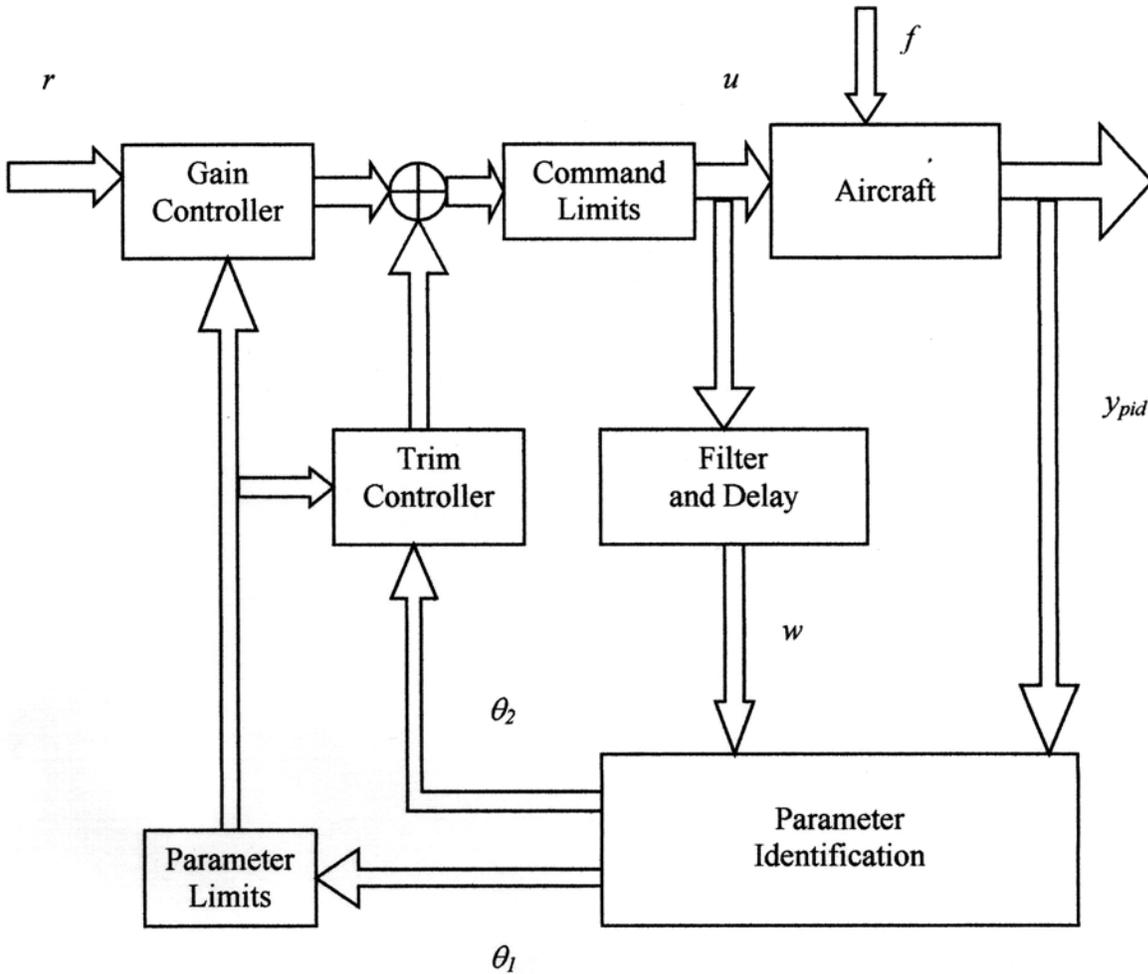


Fig. 9. Block diagram of the reconfigurable control system.

For discussion of the results with the reconfigurable control system, the following terminology is used:

- the open-loop (OL) parameters are the gains from the actuator commands (u) to the output variables (y_{pid}), as determined by the identification algorithm. They are the gains that the pilot would feel if there were no reconfigurable control algorithm.
- the closed-loop (CL) parameters are the gains from the pilot's commands (r) to the output variables (y_{pid}), as determined by a similar recursive algorithm, but after the flight. These are the gains that the pilot perceives, and are the result of the actuator effectiveness multiplied by

the control law gains.

The definitions for the variables in Fig. 9 are

$$\begin{aligned}
 r &= \begin{bmatrix} q_{com} \\ f_e \\ p_{com} \\ f_a \\ r_{com} \end{bmatrix} & u &= \begin{bmatrix} \delta_{e,left} \\ \delta_{e,rt} \\ \delta_{a,left} \\ \delta_{a,rt} \\ \delta_r \end{bmatrix} & f &= \begin{bmatrix} thr_{left} \\ thr_{rt} \\ taildrop \end{bmatrix} & y_{pid} &= \begin{bmatrix} vt \\ q \\ p \\ \beta \end{bmatrix} \\
 w &= \begin{bmatrix} \delta_{e,rt} \\ \delta_{a,rt} \\ \delta_r \end{bmatrix} & \theta_1 &= \begin{bmatrix} \theta_{q,el} \\ \theta_{p,ail} \\ \theta_{\beta,rud} \end{bmatrix} & \theta_2 &= \begin{bmatrix} \theta_{q,bias} \\ \theta_{p,bias} \\ \theta_{\beta,bias} \end{bmatrix} & & (18)
 \end{aligned}$$

Three variables are grouped into a separate input vector f because they do not interact with the control system, but are sent from the pilot directly to the airplane. Normally, the elevator command in (17) was $\delta_e = \delta_{e,rt} = \delta_{e,left}$. During an elevator fault, however, $\delta_e = \delta_{e,rt}$ and $\delta_{e,left} = 0$. Aileron faults were treated similarly. For engine faults, either thr_{left} or thr_{rt} was set to 25%, and the remaining engine received the thr command from the joystick.

Compensation of Actuator Effectiveness

The first set of experiments aimed at compensating variations of actuator effectiveness due to failures. The estimated biases were not used ($\theta_2 = 0$). In Fig. 10, the OL parameters (solid lines) and the CL parameters (dotted lines) are shown for two flights. In the first flight, (a), an elevator failure occurred. In the second flight, (b), an aileron failure occurred. The results show a marked improvement of the behavior of the reconfigurable system over the uncompensated system. In general, reconfiguration brought the effective gains close to the desired values (dashed lines), despite significant variations in the effectiveness of the actuators. At the beginning and ending of both faults, the CL parameters show transient excursions, but the variations did not reach those

of the OL values.

The pilot's evaluation of the reconfigurable system indicated that the airplane exhibited more consistent control responses, particularly after failures. Responses were also found to be more uniform during approaches to landing, an indication that the algorithm also compensated for changes associated with low speeds.

The stabilizer/elevator release fault was implemented during two flights. Fig. 11 shows pictures of the in-flight tail release. A plot of the elevator effectiveness during a tail release is shown in Fig. 12. The OL parameter is shown with a solid line and the CL parameter is shown with a dotted line. Surprisingly, the parameters did not exhibit any visible reduction in the elevator effectiveness after the release, although such changes were clearly visible in the previous, locked-elevator experiments. Reduction was also observed in previous flight tests where identification was performed using the pitch acceleration signal¹. A possible explanation for this unexpected result is that the failure produced a reduction of the control authority of the pitch command *together with* a comparable reduction of the stabilizing effect of the horizontal tail. Therefore, while the dynamics of the system were altered by the failure, the steady-state gain was not. The pilot indeed reported that the airplane was less stable in pitch after the tail release, but was *not* less responsive to elevator commands.

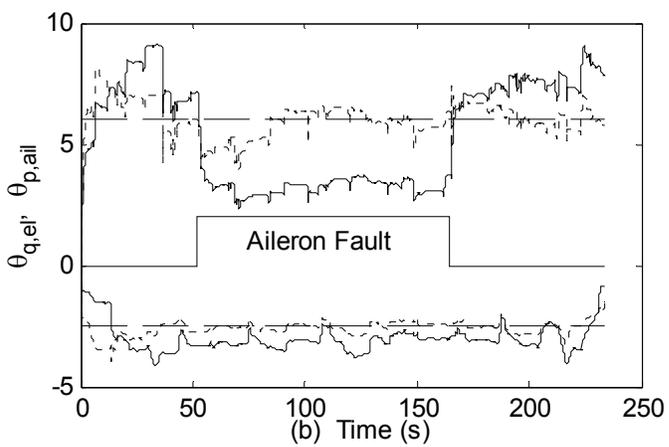
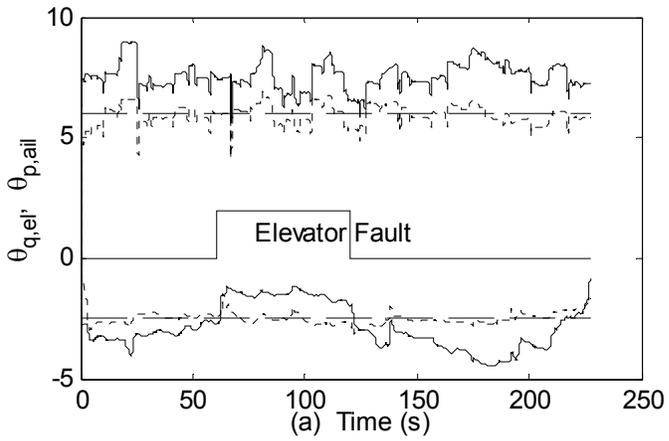
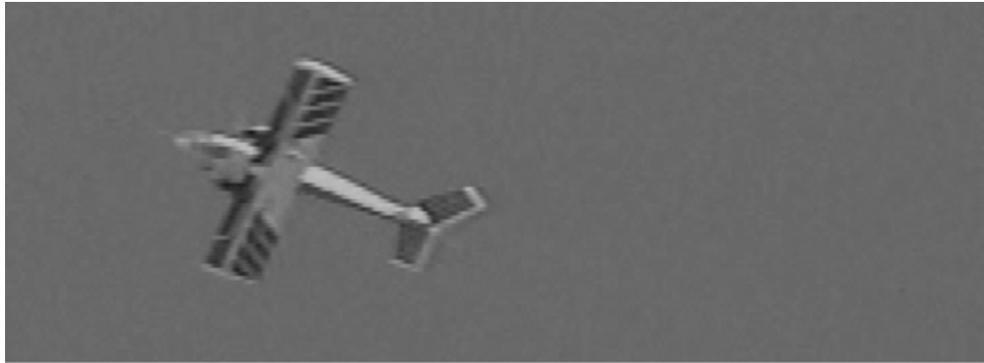


Fig. 10. Results of control reconfiguration during two flights. The algorithm attempted to stabilize the open-loop parameters (solid lines) at the desired values of -2.5 in (a), and 6.0 in (b) (dashed lines). The result was the closed-loop parameters (dotted lines).



(a)



(b)

Fig. 11. The left stabilizer/elevator is beginning to separate from the tail in (a), and is well behind the airplane in (b).

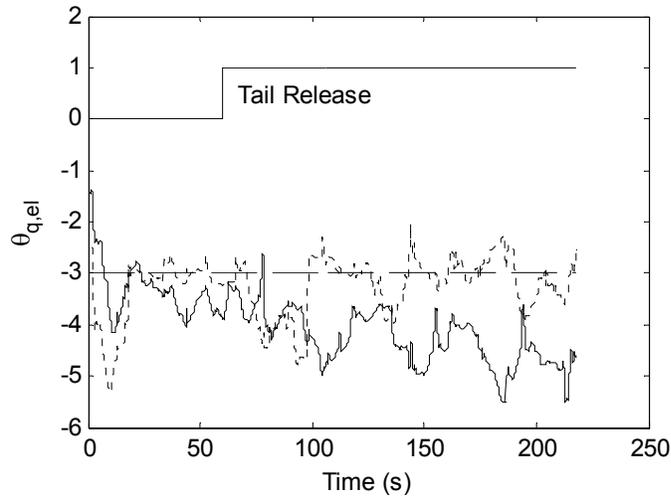


Fig. 12. Elevator effectiveness during a tail release fault. The desired value of $\theta_{q,el}$ for this flight was -3.0 (dashed line), the OL parameter is in solid lines, and the CL parameter is in dotted lines.

Automatic Trim Compensation

The control law was also implemented using the estimates of the bias parameters for automatic trim compensation. The roll bias term, $\theta_{p,bias}$, is a measure of the aircraft's roll rate when $\delta_a = 0$. The pitch bias term, $\theta_{q,bias}$, gives the pitch rate when $\delta_e = 0$, and the sideslip bias term, $\theta_{\beta,bias}$, is the sideslip with $\delta_r = 0$. These biases can be eliminated by the pilot, using manual trim controls, but this task can distract the pilot from other duties, especially after a mechanical failure. In particular, after an engine failure, a twin engine aircraft can experience large yawing and rolling moments due to asymmetric thrust.

Fig. 13 shows the results of a flight test where trim compensation was found useful in handling periods of time where the right engine was brought to idle, similar to an engine failure. As before, the reader should interpret "OL" as what the pilot would have perceived if there had been no compensation, and "CL" as what the pilot perceived due to the actions of the reconfigurable control algorithm.

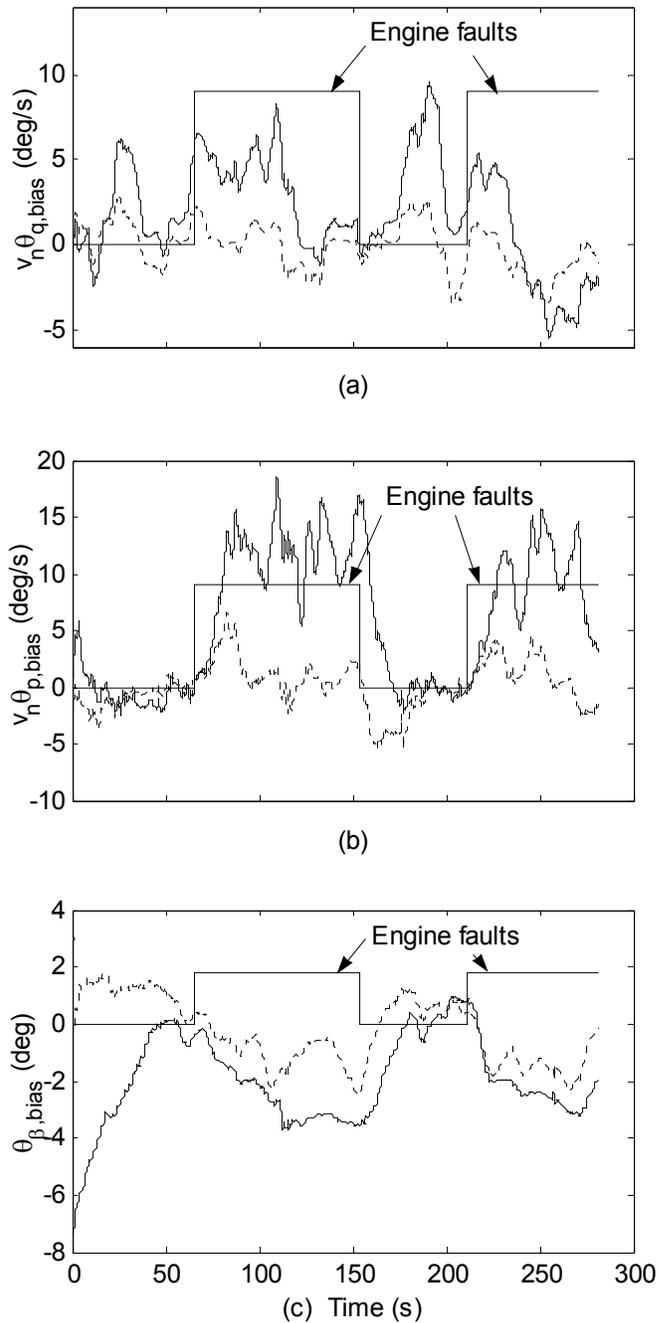


Fig. 13. Results of automatic trim compensation during two right engine faults.

The OL bias is shown with solid lines, the CL bias with dotted lines. The pitch bias in (a) is unaffected by the engine failure, but the automatic trim system removes much of the nose-up pitch bias present for most of the flight. The roll bias in (b) is close to zero except during engine

failures, where it increases to roughly 18 deg/sec. This is a roll to the right, which is reduced significantly by the trim system. A roll to the right is expected after a right engine failure because the asymmetric thrust causes the aircraft to yaw to the right. Both the yaw angle β and the lack of propeller wash from the right engine, which decreases the lift of the right wing, cause the aircraft to roll to the right. The OL sideslip bias in (c) starts at a large negative value. This is because the rotating propeller wash around the nose, combined with low airspeed, caused the sideslip sensor to indicate -30 deg during ground operations. However, by the time takeoff speed was reached, the effect of the propeller wash was overcome by the airspeed, and the sideslip was close to zero. With the correct sideslip reading, the bias term in (c) slowly converges to zero after takeoff. It becomes negative (nose right ~ 4 deg), as expected, during both of the (right) engine failures. Much of the sideslip was removed by the trim system during all phases of flight. The aircraft handled normally for the entire flight, while the remaining engine continued operating at full power.

The pilot's evaluation of this experiment indicated that the trim compensation system was effective in relieving the pilot of the need to manually trim the aircraft. This was beneficial during all phases of flight, and was crucial in compensating for the rolling effects of a failed engine (the R/C aircraft was found to be very difficult to control with a single engine, and an early flight test with engine failure, but no reconfiguration, resulted in a crash). The yawing effects did not cause difficulty of control, and the automatic sideslip trim was used for some flights but not for all because of the difficulty of identifying the sideslip channel.

CONCLUSIONS

The paper reported the results of parameter identification and reconfigurable control experiments performed on a radio-controlled aircraft. An ambitious goal was to develop the test platform in such a manner that risky tests could be performed, including some that would not be acceptable in a piloted aircraft. Failure modes were implemented that resulted in a frozen elevator, a frozen aileron, an engine failure, and the separation of an entire tail surface.

Overall, the data showed that reliable identification of critical aircraft parameters could be obtained off-line and in real-time. Remarkably, this result was obtained despite the low cost of the test platform, which resulted in high sensor noise, biases, and strong actuator nonlinearities. The effect of actuator failures could be observed on the estimated parameters, which tracked variations of effectiveness and bias. Compared to experiments performed earlier with piloted aircraft, it was found useful to reduce the model of the radio-controlled aircraft to a set of gains from the control surface deflections to the pitch rate, roll rate, and sideslip variables. Finer identification of the aircraft dynamics may not be feasible with low-cost sensors, and is possibly not useful for an aircraft with a high degree of internal stability.

Control reconfiguration experiments were also performed, and showed that a continuously adaptive control system was successful in compensating for parameter variations due to failures and changing flight conditions. Specifically, the algorithm computed the commands required to accommodate changes in the gains and biases of the system, reducing the workload of the pilot. The code for the adaptive algorithm, including the ground telemetry operations, was implemented on a single 500 MHz personal computer at an update rate of 96 Hz. Therefore, one could easily imagine such a system implemented in piloted commercial aircraft and unmanned air vehicles designed with fault tolerance in mind.

For future research, it would be interesting to explore how the reconfigurable control law can be incorporated into a nonlinear autopilot and applied to a small UAV such as the one used in this project. A more sophisticated sensor suite would be needed to obtain absolute position measurements. Nevertheless, the simplified reconfigurable control law discussed in this paper could be useful as an inner loop, instead of a more conventional system based on rotational accelerations, if the aircraft had a high degree of inherent stability.

REFERENCES

¹Ward, D., Monaco, J., Bodson, M., “Development and Flight Testing of a Parameter Identification Algorithm for Reconfigurable Control”, *Journal of Guidance, Control, and Dynamics*, vol. 21, no. 6, 1998, pp. 948-956.

²Oppenheimer, M.W., Doman, D.B., “Reconfigurable Control Design for the X-40A With In-Flight Simulation Results,” *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Providence, RI, 2004, AIAA-2004-5017.

³Brinker, J., Wise, K., “Flight Testing of Reconfigurable Control Law on the X-36 Tailless Aircraft,” *Journal of Guidance, Control, and Dynamics*, vol. 24, no. 5, 2001, pp. 903-909.

⁴Eberhardt, R., Ward, D., “Indirect Adaptive Flight Control of a Tailless Fighter Aircraft,” *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Portland, OR, 1999, pp. 466-476.

⁵Doman, D., Ngo, A., Leggett, D., Saliers, M., Pachter, M., “Development of a Hybrid Direct-Indirect Adaptive Control System for the X-33,” *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Denver, CO, 2000.

⁶Seanor, B., Song, Y., Napolitano, M., Campa, G., “Comparison of On-line and Off-line Parameter Estimation Techniques using the NASA F/A-18 HARV Flight Data”, *Proceedings of the AIAA Atmospheric Flight Mechanics Conference*, 2001, AIAA-2001-4261.

⁷Campa, G., Napolitano, M., Seanor, B., Fravolini, M., Song, Y., “Application of an Improved LWR Method to Real-Time Aircraft Parameter Identification Problems”, *Proceedings of the American Control Conference*, Anchorage, AK, 2002, pp. 197-202.

⁸Calise, A., Lee, S., Sharma, M., “Development of a Reconfigurable Flight Control Law for Tailless Aircraft,” *Journal of Guidance, Control, and Dynamics*, vol. 24, no. 5, 2001, pp. 896-902.

⁹Elgersma, M., Enns, D., Shald, S., Voulgaris, P., “Parameter Identification for Systems with Redundant Actuators,” *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Boston, MA, 1998, pp. 109-117.

¹⁰Siwakosit, W., Hess, R., “Multi-Input/Multi-Output Reconfigurable Flight Control Design,” *Journal of Guidance, Control, and Dynamics*, vol. 24, no. 6, 2001, pp. 1079-1088.

¹¹Shtessel, Y., Buffington, J., Banda, S., “Tailless aircraft flight control using multiple time scale reconfigurable sliding modes,” *IEEE Transactions on Control Systems Technology*, vol. 10, no. 2, 2002, pp. 288 – 296.

¹²Bodson, M., Pohlchuck, W., “Command Limiting in Reconfigurable Flight Control”, *Journal of Guidance, Control, and Dynamics*, vol. 21, no. 4, 1998, pp. 639-646.

¹³Bodson, M., “A Reconfigurable Nonlinear Autopilot”, *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 5, 2003, pp. 719-727.

¹⁴Steinberg, M., “Comparison of Intelligent, Adaptive, and Nonlinear Flight Control Laws,” *Journal of Guidance, Control, and Dynamics*, vol. 24, no. 4, 2001, pp. 693-699.

¹⁵Hallberg, E., Kaminer, I., Pascoal, A., “Development of the Rapid Flight Test Prototyping System for Unmanned Air Vehicles”, *Proceedings of the American Control Conference*, Philadelphia, PA, 1998, pp. 699-703.

¹⁶Kaminer, I., Yakimenko, O., Dobrokhodov, V., Lim, B., “Development and Flight Testing of GNC Algorithms using a Rapid Flight Test Prototyping System”, *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, 2002, AIAA-2002-4653.

¹⁷Jang, J., Tomlin, C., “Autopilot Design for the Stanford DragonFly UAV: Validation through Hardware-in-the-Loop Simulation”, *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Montreal, Canada, 2001, AIAA-2001-4179.

¹⁸Jang, J., Tomlin, C., “Design and Implementation of a Low Cost, Hierarchical and Modular Avionics Architecture for the DragonFly UAVs”, *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, 2002, AIAA-2002-4465.

¹⁹Napolitano, M., “F-16 and YF-22 Scale Models with On-board On-line Learning Microprocessors-based Neural Algorithms for Autopilot and Fault-Tolerant Flight Control Systems”, Department of Mechanical and Aerospace Engineering, West Virginia University, 1999, pp. 198-203.

²⁰Atkins, E., Miller, R., Pelt, T., Shaw, K., Ribbens, W., Washabaugh, P., Bernstein, D., “Solus: An Autonomous Aircraft for Flight Control and Trajectory Planning Research”, *Proceedings of the American Control Conference*, Philadelphia, PA, 1998, pp.689-693.

²¹Shore, D., “Accommodation of Mechanical Faults with Reconfigurable Flight Control Using a Model Aircraft”, Master of Science Thesis, Department of Mechanical Engineering, University of Utah, Salt Lake City, Utah, Dec. 2003.

²²Sastry, S., Bodson, M., *Adaptive Control: Stability, Convergence, and Robustness*, Prentice Hall, Englewood Cliffs, NJ, 1989. The book may be viewed as a PDF file through the web page: <http://www.ece.utah.edu/~bodson/acscr>.