

# IMPLEMENTATION AND FLIGHT TEST ASSESSMENT OF AN ADAPTIVE, RECONFIGURABLE FLIGHT CONTROL SYSTEM

J. Monaco, D. Ward, R. Barron, and R. Bird

Barron Associates, Inc.  
The Jordan Building, 1160 Pepsi Place, Suite 300  
Charlottesville, Virginia 22901

## ABSTRACT

During the summer of 1996 a series of flight tests demonstrated a new indirect-adaptive approach to reconfigurable flight control known as the self-designing controller (SDC). The SDC achieves improved, appropriately decoupled responses during arbitrary effector or airframe impairment scenarios, and successful SDC flight tests culminated with smooth landing of the VISTA/F-16 in cross-wind conditions with a (simulated) missing primary control surface (left horizontal tail). The SDC couples model-following receding-horizon optimal control with an on-line parameter identification (ID) algorithm designed to provide smooth, accurate estimates of possibly time-varying system parameters, even under conditions of low excitation. The adaptive model-following approach is designed to reduce control law development costs and improve system performance in the presence of gradual or abrupt changes, including unforeseen events. This paper provides (1) a brief summary of the SDC algorithms, (2) a discussion of SDC implementation on the VISTA/F-16 flight control hardware, (3) a summary of flight test results, and (4) suggestions for further research in reconfigurable/adaptive controls.

## INTRODUCTION

It has long been a goal of flight-control research to achieve first-rate flying qualities across constantly expanding operational envelopes and for novel stores configurations. Simultaneously, there has been interest in reducing the time and cost for developing new flight control systems. In recent years, attention has also focused on flight-control system robustness, i.e., the ability to operate well under off-nominal or unexpected conditions. The emphasis on robustness has led, logically, to studies of reconfigurable control, which is intended to adapt quickly to control effector and airframe damage as well as less traumatic events such as release of stores or gradual component aging.

*This work has been supported by the Directorate of Mathematical and Computer Sciences, Air Force Office of Scientific Research, Bolling AFB, DC, under SBIR Phase II Contract F49620-94-C-0087 with Barron Associates, Inc. Technical supervision was supplied by the USAF Flight Dynamics Directorate, Wright Laboratory, Air Force Materiel Command (ASC), Wright-Patterson AFB, OH. Lockheed Martin Tactical Aircraft Systems (Ft. Worth, TX) and Calspan SRL Corporation (Buffalo, NY) participated as subcontractors and performed batch simulation evaluations and flight test support, respectively.*

Copyright ©1997 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

Traditional control design methods frequently require specialized knowledge, substantial off-line analysis, and extensive in-flight tuning and validation (often accompanied by numerous iterations of large and complex pre-specified gain schedules). Design and development expenses mushroom when one requires *a priori* control modes designed to handle the large universe of possible anomalies that may be experienced during the flight-vehicle service life. Rather than rely on numerous control system implementations based on pre-hypothesized impairment or airframe damage scenarios, the SDC uses an indirect-adaptive control architecture as shown in the Fig. 1 below.

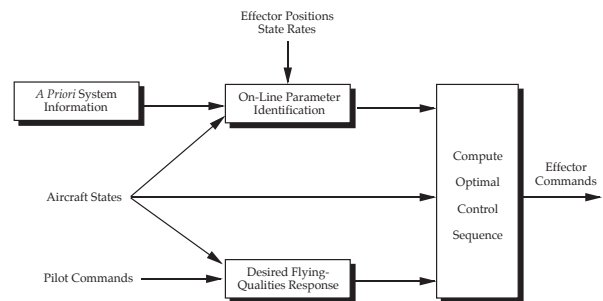


Figure 1: SDC Architecture

The SDC computes a time-varying model of the current aircraft dynamics using a novel on-line system identification technique that can rapidly track time-varying parameters and is robust to adverse conditions such as low excitation or correlated inputs. The identified model of the current aircraft dynamics is then communicated to an optimal control module that computes the effector commands to give the desired aircraft responses as specified by a set of flying qualities models. The ability of the SDC to reconfigure rapidly after single or multiple impairments, making effective use of residual effector authority, greatly enhances aircraft survivability. Additionally, because the parameter ID and on-line control gain computations are independent processes, these modules can be extracted from the SDC and incorporated into other control or diagnostic systems to provide additional capabilities. Finally, the model-following design approach has the potential to reduce significantly the time and expense of developing flight control systems for new aircraft, operational conditions, stores configurations, flying-qualities requirements, and failure modes. This is

borne out, in part, by the fact that a full-envelope reconfigurable F-16 primary flight controller was developed and flight tested within the time and financial constraints of a Phase II SBIR contract.

## ALGORITHM SUMMARY

### MSLS Parameter Identification

The aircraft parameter identification problem is formulated to obtain composite (lumped) dimensionalized stability and control derivatives required for the computations of controller gains. Toward this end, the equations of motion are defined as:

$$\dot{x}_{lon} - NL_{lon}(\dots) = \theta_{lon}^T \phi_{lon} \quad (1)$$

$$\dot{x}_{latdir} - NL_{latdir}(\dots) = \theta_{latdir}^T \phi_{latdir} \quad (2)$$

where  $x_{lon}^T = [q, \alpha]$  and  $x_{latdir}^T = [p, r, \beta]$  represent the aircraft states;  $NL(\dots)$  represents the nonlinear terms due to inertial cross coupling and orientation-dependent gravitational effects;  $\phi_{latdir}^T = [p, \alpha, \delta_{ts}, \delta_{ta}, 1]$  and  $\phi_{lon}^T = [p, r, \beta, \delta_{ts}, \delta_{ta}, \delta_{fa}, \delta_{rud}, 1]$  are vectors of observations, consisting of aircraft states, effector positions, and bias terms; and  $\theta_{lon}^T = [\theta_q, \theta_\alpha]$  and  $\theta_{latdir}^T = [\theta_p, \theta_r, \theta_\beta]$  are matrices containing the lumped parameters to be identified. By analytically linearizing the nonlinear terms in (1), collecting terms, and simplifying, one can express the fundamental nonlinear equations of motion in linear time-invariant (LTI) form suitable for use by a number of control design techniques. For detailed derivations of these equations, consult [1].

There are two significant difficulties with on-line estimation of aircraft stability and control derivatives: (1) data collinearities and (2) time-varying parameters. Data collinearities occur when any of the input variables to a system being identified can be represented as linear combinations of other input variables. Such conditions can occur in cruise flight or result from linear state feedback and/or the use of “ganged” effectors (e.g., combining both asymmetric flap and asymmetric tail to generate rolling moment). These data collinearities are a result of *insufficient excitation* and lead to singularities in the regression equations unless one ceases adaptation or extends the data-collection window far enough into the past to reduce data correlations.

Time-varying parameters can arise either during changes in flight condition (slow variations) or during impairment, stores release, or flight in a nonlinear regime (fast variations). To track time-varying parameters, one can incorporate a forgetting factor in Recursive Least Squares (RLS) estimation or increase the size of the parameter variance model in Kalman or Extended Kalman Filter (EKF) techniques. However, these approaches lead to two significant difficulties. As the data window gets smaller, the likelihood of data collinearities within the data window increases significantly, leading to the problems discussed above. Additionally, for a given level of excitation, there is a certain rate of parameter variation above which it is

impossible to track accurately the time-varying parameters. This is because the data window cannot be made small enough, and the RLS and EKF algorithms break down if the system being identified is not stationary within the data window.

To overcome the problem of identifying time-varying parameters in a system which, frequently, is insufficiently excited, Barron Associates, Inc. (BAI) developed a modified sequential least squares (MSLS) algorithm. This is an iterative batch parameter identification algorithm (hence the use of the word *sequential* as opposed to *recursive*) that incorporates additional constraints to prevent covariance matrix singularities that may occur when one requires relatively small data windows to track parameters that vary at moderate or higher rates. The process to prevent these singularities is known as *regularization* [2] and is the principle behind a number of successful batch and nonlinear regression techniques such as ridge regression [3] and the Levenberg-Marquardt algorithm [4, 5].

MSLS augments the squared-error cost-function prescription with an additional penalty or penalties consisting of linear equality constraints on the parameters. These constraints are of the form:

$$K(t) = M\theta(t) \quad (3)$$

where  $\theta \in \mathbb{R}^{P \times o}$  is a matrix of estimated parameters,  $K(t) \in \mathbb{R}^{\ell \times o}$  and  $M \in \mathbb{R}^{\ell \times P}$  define the linear constraint relationship,  $\ell$  is the number of constraints,  $P$  is the number of parameters to be estimated, and  $o$  is the number of outputs (or equations) that share a common input vector,  $\phi$ .

The above constraint is added to the standard squared-error cost function and results in a total empirical loss function of the form:

$$J(\theta) = \frac{1}{2} \sum_{\nu=t_0}^t \lambda^{t-\nu} (y(\nu) - \theta^T \phi(\nu))^T (y(\nu) - \theta^T \phi(\nu)) + \frac{\nu}{2} \|W_0^{1/2} (K(t) - M\theta)\|_F^2 \quad (4)$$

where  $y \in \mathbb{R}^{o \times 1}$  is the vector of system outputs;  $\phi \in \mathbb{R}^{P \times 1}$  is the vector of system inputs;  $W_0 \in \mathbb{R}^{\ell \times \ell}$  is the relative penalty associated with each constraint;  $\|\cdot\|_F$  is the  $F$  (or Frobenius) norm; and  $0.0 < \lambda \leq 1.0$  is the “forgetting factor” used to discount prior observations. Here  $\nu$  is the area of the window over which the cost function is computed. This area is included in the penalty term so that the relative influence of a fixed weight remains essentially the same regardless of forgetting factor used.


In MSLS, both *temporal* constraints and *spatial* constraints. *Temporal* constraints penalize parameter estimates that deviate from their previous estimate. This results in a smoothing over time, but, as is shown below, does not hinder the ability to track rapidly varying parameters. In [6] it is shown that *temporal* constraints have the effect of reducing the data window size during periods of high excitation, resulting in rapid parameter tracking. *Spatial* constraints

penalize parameter estimates that diverge from *a priori* estimates of their true values. These estimates can be constant, or they can be computed by on-board nonlinear models.

To show how *temporal* constraints can improve dramatically upon RLS in ill-conditioned problems, a simple linear system was created of the form:

$$y(k) = [\theta_1(k), \theta_2(k)] \begin{bmatrix} \phi_1(k) \\ \phi_2(k) \end{bmatrix} + \nu(k) \quad (5)$$

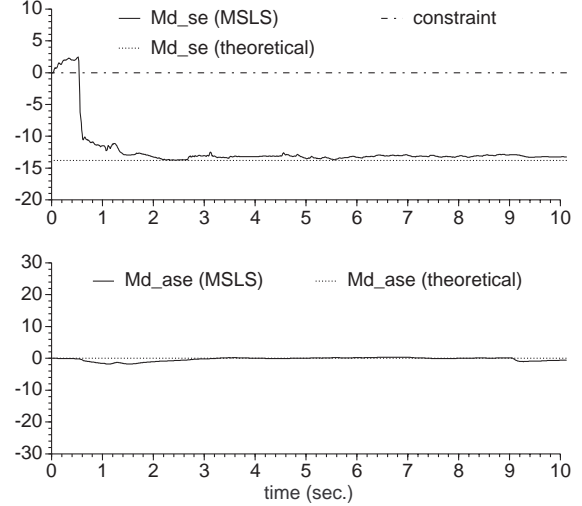
where  $\nu(k)$  is zero-mean Gaussian white noise with variance  $\sigma^2 = 0.1$ . To create a poorly conditioned observation matrix, the measurements  $\phi_1$  and  $\phi_2$ , which would correspond to state measurements and/or surface deflections in a flight-control context, were driven with *highly-correlated* Gaussian signals. Fig. 2 shows a comparison of MSLS and RLS estimates of the system parameters. The MSLS algorithm faithfully tracks both slowly and rapidly changing parameters. Note that near-instantaneous tracking of the abrupt change in  $\theta_2$  is not compromised by the use of temporal constraints. The RLS results were the best that the authors were able to achieve by adjusting the forgetting factor.



**Figure 2: MSLS vs. RLS Parameter Estimation**

Evaluating MSLS on flight data yields similar results. Fig. 3 shows MSLS parameter identification results on a pitch-axis maneuver. In this data the asymmetric elevon surface,  $\delta_{ta}$  is highly correlated with the bias term in the regressor,  $\phi$ . Despite the absence of appreciable asymmetric elevon activity, the MSLS estimate is stable and well-contained, because the use of *temporal* and *spatial* constraints serves to regularize  $\frac{\partial \dot{Q}}{\partial \delta_{ta}}$  during periods of low excitation. Typically one is not concerned with identifying the contribution of  $\delta_{ta}$  to the pitching moment; however, in the case of

failures, these “off-diagonal” terms become extremely important. When RLS was evaluated on the same data used for Fig. 3 the  $\frac{\partial \dot{Q}}{\partial \delta_{ts}}$  estimate was only slightly degraded; however, the  $\frac{\partial \dot{Q}}{\partial \delta_{ta}}$  estimate was markedly degraded to the above-mentioned collinearities; the RLS estimate of this parameter wandered between +500 and -500.



**Figure 3: MSLS Estimates of  $\frac{\partial \dot{Q}}{\partial \delta_{ts}}$  and  $\frac{\partial \dot{Q}}{\partial \delta_{ta}}$**

One might be tempted to argue that the *spatial* constraints simply provide the algorithm with the correct answer. To show that this is not the case, an identical *spatial* constraint was also applied to the symmetric tail contribution to pitching moment. However, because there was sufficient symmetric elevon activity, MSLS ignored the spatial constraint and tracked the true parameter faithfully. The ability to ignore spatial constraints during periods of sufficient excitation is an important characteristic of the algorithm. Additional parameter identification results on flight data are presented later.

### RHO Control Law

The SDC uses a receding-horizon optimal (RHO) control law architecture in concert with a regularized parameter ID algorithm for estimation of aircraft stability and control derivatives. Receding-horizon control was originally developed in the process controls industry under a variety of names, the most common of which is generalized predictive control (GPC) [7, 8]. There are a number of variations of GPC methodologies, including discrete and continuous time versions, tracking formulations, and adaptive algorithms that combine the GPC algorithm with an on-line system identification technique. However, all of the variants work in very much the same way; a finite-time optimal control solution is computed that minimizes the predicted error between predicted plant dynamics and the desired plant response. This minimization results in a sequence of optimal control commands, and the first command, correspond-

ing to the current time, is applied to the system. At the next control update, rather than applying the second command in the open-loop optimal command sequence, the finite horizon optimization is completely recalculated using up-to-date estimates of the plant dynamics, desired control, and system state. In this way, the open-loop finite-horizon optimal control problem becomes closed-loop, and the optimization horizon is said to *recede* because the controller never applies the commands corresponding to the end of the horizon.

A receding-horizon controller shares a number of advantages with LQ control techniques, especially stability and robustness [9]. However, unlike infinite-time LQ control, a receding-horizon controller *anticipate* desired response and optimizes tracking over a short horizon. This quality adds desirable phase lead to the controller and makes it attractive for multi-input multi-output (MIMO) control problems where one is interested in achieving a specific *transient* response as well as a specific *steady-state* response [10]. In fact, one paper recently argued that, for problems that are not inherently linear-time-invariant (LTI), receding-horizon control is “the only viable controller synthesis method” [11].

The RHO control law makes use of an augmented system of equations consisting of the aircraft dynamics and integrator states. It is also possible to incorporate effector dynamics and flying-qualities models into the augmented system; however, this was not done for the flight test program.

### Desired Flying-Qualities Responses

Modern aircraft typically respond as high-order systems because of their fundamental aero-inertial properties, as well as behaviors owing to control law formulation, actuator and sensor dynamics, and digital hardware components. Often, these high-order responses are not adequately represented by traditional low-order models of aircraft characteristics. However, low-order equivalent system (LOES) parameters, e.g., damping ratio and undamped natural frequency, do provide valuable information for modeling desired aircraft handling qualities. For pre-defined levels of handling-qualities performance, MIL-STD-1797A guidelines specify acceptable ranges of the LOES parameters [12].

The RHO control law utilizes models of the desired aircraft transient responses so that it can *predict and anticipate* these behaviors and generate commands to track the models optimally. MIL-STD-1797A sets forth desired responses in the form of low-order equivalent system linear differential equations.

The specific command gradients and flying-qualities models used in the SDC evolved from prescriptions provided by Lockheed Martin Tactical Aircraft Systems (LMTAS) under a Neural Network Flight Control System (NNFCS) contract [13]. The NNFCS work emphasized high-agility maneuvering at large, post-stall angles of attack; however, the models were designed to achieve Level I handling-qualities at all flight conditions and angles of at-

tack. For a more complete discussion of these models and the associated low-order equivalent system (LOES) parameters, see [13, 1].

## SIMULATION-BASED EVALUATION

### Time-Domain Simulation Analysis

The SDC was evaluated off-line at BAI using the LMTAS nonlinear VISTA/F-16 six-degree-of-freedom Aircraft Trim, Linearization, and Simulation (*ATLAS*) program. This simulation includes full aerodynamic data, fourth-order actuator dynamics, a propulsion model, and a model of the software that controls and monitors the research flight control law. To validate successive versions of the SDC software, several thousand simulation runs were conducted. For new versions of the software, the simulated closed-loop responses were evaluated for exhaustive combinations of varied mass and inertial properties, stores configurations, initial trim flight conditions, single and multi-axes maneuvers, levels of turbulence, types and magnitudes of effector impairments, and amounts of pure time delay inserted in the actuator paths.

Using these data, SDC performance was assessed for all simulation experiments. Here, the criteria for adequate performance were defined to be: (1) flying qualities comparable to those of the Block 40 primary flight control law during normal, unimpaired flight and subject to various levels of simulated turbulence; (2) flying qualities superior to the Block 40 control law under failure conditions - among other traits, this required that aircraft motions remain uncoupled after impairments; (3) close tracking of LOES models for all axes; (4) minimal occurrences of effector position and/or rate saturation; (5) absence of unwarranted effector oscillations; and (6) robustness to pure time delays on the order of 75 to 100 *ms* in the actuator paths.

Because the F-16 airframe is statically unstable about its pitch axis, additional attention was given to the longitudinal flying qualities following an impairment. Although the Block 40 leading-edge flap schedules stabilize the pitch-axis response of the unimpaired aircraft, the ability of these effectors to maintain longitudinally-stable flight subsequent to a control surface failure is not known. Therefore, simulation experiments were conducted to assess the ability of the SDC to reconfigure and control the unstable aircraft, i.e., with the leading-edge flaps disabled. Simulations revealed that SDC control of the statically-unstable airframe was qualitatively similar to that when the leading-edge flap schedules were used.

Figs. 4 and 5 illustrate the closed-loop response at 260 *KCAS* and 15*kft*. In this maneuver a 100 percent left horizontal tail impairment is introduced 11 *sec*. into the maneuver. For all of the SDC work, a missing actuator was simulated by commanding the surface to the negative of the local angle of attack (which is different than the aircraft angle of attack). Unlike a locked surface, this type of failure

changes the effective  $M_\alpha$  and destabilizes the aircraft.

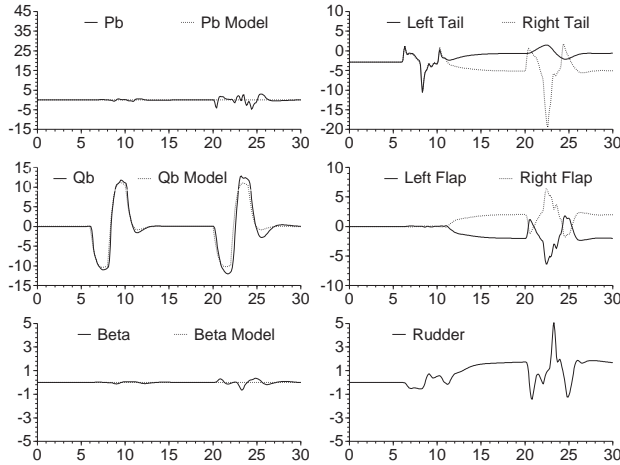


Figure 4: SDC Responses with Failure

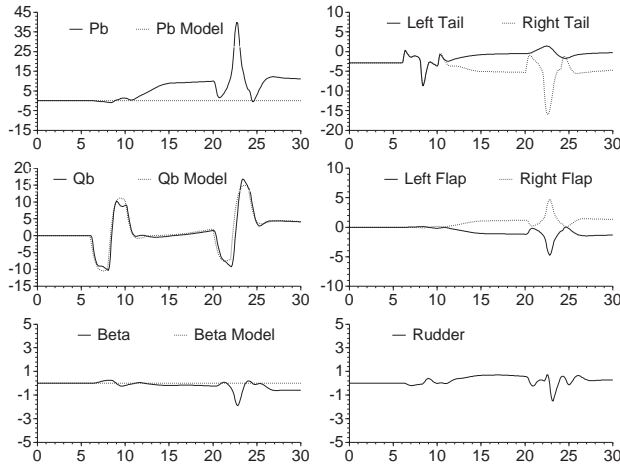


Figure 5: Block 40 Responses with Failure

As seen in Fig. 4, the SDC exhibits faithful model following prior to the introduction of the impairment at 11 sec. During the insertion of the failure, the SDC begins to compensate for the roll generated by the asymmetric tail deflection by commanding an opposing asymmetric flaperon deflection. After approximately 1 - 2 sec., any small pitch-rate, roll-rate, and sideslip transients have been removed by the reconfigured SDC. A second, moderately aggressive pitch doublet sequence ( $\Delta N_z \doteq 3.5 g$ ), of equal magnitude and duration to the first, is executed at 20 sec. Even with only one effective tail surface, the SDC is still able to track the flying-qualities models quite well. Lateral and directional motions are reasonable. Body-axes roll-rate and sideslip excursions are kept to  $\pm 4 \text{ deg./sec.}$  and  $\pm 0.5 \text{ deg.}$ , respectively, during the most aggressive maneuvering. Straight, wings-level flight is preserved, with the bank angle remaining essentially zero for the duration of the simulation.

As seen in Fig. 5, the Block 40 controller also exhibited good pitch response after the impairment had occurred; however, the lateral and directional responses were no longer decoupled from the longitudinal motion. After the failure, a steady-state body-axes roll rate, on the order of  $10 \text{ deg./sec.}$  resulted. An unabated, nonzero sideslip was also generated. During the second pitch-doublet sequence, the *uncommanded* lateral motions were approximately 5 times greater than the corresponding SDC excursions.

### Piloted Simulation Analysis

It would have been desirable to compare further the flying qualities of the SDC controller vis-a-vis the robust, but non-reconfigurable, baseline F-16 controller in actual flight tests. However, for safety reasons, it was not possible to fly the Block 40 system with a simulated impairment. Therefore, LMTAS conducted real-time pilot-in-the-loop simulation studies to compare the two control laws. As was the case in the batch simulation, all pilots reported that the SDC axes decoupling during impairments was significantly improved over the baseline Block 40 controller. With the non-reconfigurable Block 40 controller, the pilots consistently reported difficulties manually counteracting the adverse axes coupling during impairments. The following are representative pilot comments:

- *Pilot A*: "... [the baseline controller has] noticeable pitch-roll interaction."
- *Pilot B*: "... [the baseline controller has] more significant roll coupling than with [the] SDC ... [during pitch captures, the] least little input can drive a low frequency PIO due to roll coupling"
- *Pilot C*: "... [the baseline controller has] a lot bigger tendency to roll in pitch than with the SDC ... There is a lot of roll in the pitch up during pitch attitude captures. It's really bad, but not uncontrollable."

Thus, it is seen that responses following a simulated impairment are significantly different for the two control laws (SDC and Block 40). Both piloted and batch simulation results revealed that the most appreciable differences between the two control systems occurred for horizontal-tail and multiple-effector impairments. Trailing-edge-flap failures resulted in less-pronounced, albeit significant differences. In the Block 40 system, differences between unimpaired and impaired responses typically observed during a flap failure were: (1) severely limited and asymmetric roll responses and (2) increased lateral/directional coupling. When compared to Block 40 control of the impaired aircraft, the SDC represented a marked improvement; the Block 40 controller (which was not developed for reconfigurable control) exhibited coupling of longitudinal and lateral/directional motions not produced by the SDC. While the pilots noted that some of the steady-state couplings

could be “trimmed-out,” there were still significant cross-couplings during maneuvers that made the Block 40 system more difficult to control and resulted in post-failure Level III pilot ratings.

Table 1 summarizes pilot ratings for the SDC for a variety of maneuvers and failure conditions. While these ratings are based on a limited number of pilots, test runs, and maneuvers, they generally coincide with the ratings observed in the flight tests. The pilots noticed areas in which the SDC could be improved, and, as a result, BAI made some adjustments of the SDC models prior to the flight tests.

**Table 1: SDC Cooper-Harper Ratings**

Maneuver	No Fail	Missing LHT	Missing RHF
Bank-to-Bank Rolls	3	4	4
Loaded Rolls	3.5	3	3
Pitch Attitude Captures	4	3	4


**Simulation-Based Stability Analysis**

Because the SDC is adaptive, traditional linear gain and phase-margin analyses cannot be conducted. Therefore, the stability margins were determined by LMTAS using 6-DOF nonlinear simulations in which an interface system, with the ability to multiply or delay control surface signals, was inserted ahead of the VISTA/F-16 aircraft model. For each of the control surface loops, the gains and delays in this subsystem were adjusted while applying small doublet inputs (pitch stick, roll stick, and rudder pedal, respectively) that perturb the aircraft dynamics. Through trial and error, the stability margins were then estimated by determining the gains and delays that caused the pitch-rate, body-axes roll-rate, and sideslip responses to become neutrally stable.

For high-performance military aircraft, desired gain and phase margins are 6 dB and 45 deg., respectively. However, the production F-16, with Block 40 control laws, does operate with phase margins as low as 30 deg. at several points in the flight envelope. SDC stability margins for each of nine mass properties/flight condition configurations were computed. The SDC gain margins were proper in all axes. Longitudinal phase margins were comparable to those of the Block 40 controller, ranging from approximately 30 deg. to in excess of 50 deg. Phase margins for the lateral axis surpass flight control design specifications at all test points. In the directional axis, phase margins were proper for all cases except one in which the sideslip phase margin was below design specifications. For the flight tests, the tracking and integral-error penalties were ultimately reduced significantly. Thus, it is reasonable to expect that the phase margins, especially in the sideslip axis, were substantially better in the flight-tested configuration than the configuration used for the stability analysis. Time and funds did not permit a follow-up analysis to confirm this assertion.

**VISTA/F-16 IMPLEMENTATION**

The VISTA/F-16 Aircraft is sketched in Fig. 6. The VISTA contains a Variable Stability System (VSS) comprised of three Rolm Hawk/32 computers which typically serve to compute effector commands that enable the VISTA responses to imitate the behavior of some other aircraft. Each of these Hawk/32 computers has (approximately) the throughput of an Intel 80286 processor. In the SDC, the VSS software was replaced with the adaptive SDC software.



**Figure 6: VISTA F-16 Aircraft**

The VISTA digital flight control computer (DFLCC) has additional logic that continuously monitors the VSS-computed commands and reverts control to the primary F-16 Block 40 control laws if research control-law commands are deemed unsafe. This logic is known as the Vehicle Integrity Monitor (VIM) and was retained for the SDC. The VIM is responsible for ensuring that the VSS system does not violate any pre-specified flight envelope or structural safety limits. These limits include constraints on normal load factor, lateral accelerations, angle of attack, angle of sideslip, altitude, dynamic pressure, roll rate, pilot commands, effector positions, and effector rates. Of particular interest were structural limits that prevent any actuator commands from putting twist on the fuselage when the velocity exceeds 275 knots. However, a typical reconfiguration scenario is one in which, during a tail failure, the SDC uses asymmetric flaperon to counteract rolling moments generated by asymmetries in the tail deflections. Therefore, all flight tests were conducted below 275 knots to avoid this constraint. It is interesting to note that the VIM introduces an additional transport delay in the control loop; however, the final SDC design proved robust to this unanticipated delay, and overall performance was only slightly degraded.

The SDC software was segmented so that its computations could be shared by the three Hawk computers. Hawk 1 implemented the RHO control law and failure simulation. Hawk 2 resources were used for parameter identification, and Hawk 3 performed integration of the Riccati equations.\* Each Hawk flight control computer accommodated foreground and background processes. A foreground

\*Note that approximately two-thirds of the Hawk computer resources are used for sensor and data management and are, thus, unavailable for the SDC.

process was initiated and terminated once each 40Hz iteration. Estimation of the aircraft stability and control derivatives and calculations of the Riccati  $K$  matrices and  $k$  vectors were done as background tasks and allowed to continue from one iteration to the next. Foreground tasks supplied the most recent estimates of the stability and control derivatives and the current Riccati gains to the optimal control calculations. The RHO control law simply used these parameter sets until new estimates became available. With a 40 Hz system rate, analysis of flight data revealed that the background tasks ran at approximately 13 Hz, roughly twice as slow as the 25 Hz rates that had been anticipated from timing benchmarks conducted at BAI and Calspan. However, the reconfigurable controller continued to perform well at these slower rates.

## VISTA/F-16 FLIGHT TESTS

### Introduction

The self-designing controller underwent flight testing during the period May 14, 1996, to July 9, 1996. Five test flights were performed, each lasting approximately one hour. The goals of the flight tests were to:

- demonstrate improved handling qualities (HQ) *with* simulated impairments,
- achieve Level-I and high Level-II HQ *without* impairments,
- land *without* impairments, and
- land *with* impairments.

Due to the limited amount of flight time, the SDC tests did not employ a *build up* approach whereby various components of the software and algorithm would have been tested sequentially. Rather, the test plan for the first flight called for the entire algorithm and associated software to be engaged; then, depending on the results, experiments of increasing or decreasing difficulty would be tried.

Prior to each flight, the SDC software was tested extensively in the VISTA/F-16 ground simulator. This simulation uses the aircraft as a test-bed, with all flight control and Hawk software functioning as it would in the air. Effector positions are processed through aircraft equations of motions, and computed state information is sent to the sensors as appropriate analog voltages. Using the ground simulation, pilots and engineers were able to evaluate VSS performance from the cockpit and record and analyze real-time data. The ground simulation provided essential pre-flight evaluation and validation of SDC software. Configuration changes could be rapidly assessed, allowing more productive use of the limited flight time.

The ground simulation was also used to test disengage transients that would result any time aircraft control reverted to the baseline (Block 40) F-16 flight control system. Intentional SDC disengagements from a variety of flight

conditions and orientations, with and without impairments, were used to ensure that disengage transients (e.g., linear and angular acceleration excursions) were acceptable and that rapid recovery from all scenarios was achievable. The disengage transients received substantial attention from the SDC design team and pilots because emphasis in the flight tests was placed on failed performance during powered approaches and landing tasks. During all ground simulation check-outs, at-altitude evaluations, and low-approach validations, the transients resulting from intentional and unintentional disengages were observed to be smooth and well within the acceptable limits.

### Flight 1

The initial SDC flight took place on May 14, 1996. The flight crew comprised John Ball (Safety Pilot/Calspan) and Jeff Peer (Evaluation Pilot/Calspan). Total flight time was 1.4 hours from take-off to landing; turbulence was minimal to nonexistent for the duration of the flight. The goals for this flight were to (1) engage the SDC, (2) collect flying-qualities and parameter identification data for nominal (unimpaired) flight, (3) assess handling characteristics and obtain parameter ID data during impaired flight, and (4) evaluate SDC responses for various configuration and parameter changes to facilitate algorithm refinement in subsequent flight software.

Data were collected at two distinct test points 8.5 *kft.*, 11 *deg.* AOA and 15 *kft.*, 260 *KCAS*, with a majority of the flight devoted to the latter condition. The SDC was able to engage the first time it was requested to do so and remained engaged for up to ten minutes at a time. While this was an important milestone, flying qualities were generally unacceptable and inconsistent. Pilots noted flight periods characterized by pitch “spikes” and several occasions of roll and/or yaw oscillations. Some oscillations had the potential to produce divergent response, and the pilots executed manual disengage sequences more than once. Post-flight data analysis revealed an inter-Hawk data communication problem that caused improper parameter estimates to be passed to the Riccati gain computation. Additionally, several sensed parameters being supplied to SDC routines were in error. Of specific importance were values used in the MSLS spatial constraint routines. These spatial constraints are simple functions of aircraft states and observables, e.g., angle of attack, angle of sideslip, aircraft weight, and dynamic pressure. Data from Flight 1 revealed that the remaining fuel quantity, which was being used to estimate the total aircraft weight, was improperly supplied; the magnitude of the signal was not self-consistent and remained constant for the duration of the flight. These problems were fixed prior to the second flight.

Analysis of Flight 1 also resulted in the following modifications: (1) the pitch-rate flying-qualities model was augmented to include (optionally) a roll-angle correction term ( $g \cos \Phi \cos \Theta$ ), designed to provide a normal-acceleration feel similar to the Block 40  $N_z$  command system; (2) the

roll-rate flying qualities model was also modified to incorporate (optionally) dihedral effects (roll due to sideslip commands); (3) optional yaw damping logic was added to “stiffen” the SDC directional response based on pilot comments that the Flight 1 SDC was “fishy” in yaw; (4) some MSLS and RHO parameters were adjusted to account for slower update rates and additional time delays.

### Flight 2

The second SDC flight occurred on May 30, 1996. During ground simulation and checkout prior to Flight 2, BAI also uncovered a VISTA hardware implementation problem that was responsible for errors in the yaw-channel parameter estimates. Ground simulation data revealed that the  $\dot{r}$  signal computed by the VISTA flight control computers was inconsistent with numerical differentiation of the  $r$  signal. The VISTA/F-16 estimates of vehicle angular accelerations ( $\dot{p}$ ,  $\dot{q}$ ,  $\dot{r}$ ) were obtained using differences between *two* appropriate linear-accelerometer measurements. These approximations yield reasonable results for the  $\dot{p}$  and  $\dot{q}$  computations, but because of the collocation properties of the VISTA accelerometer suites, a weighted difference of *three* linear accelerometer-measurements gives a much more accurate estimate of  $\dot{r}$ . Making this change, yaw channel derivatives and responses were markedly improved.

The crew for Flight 2 were Jeff Peer (Safety Pilot/Calspan) and Joseph Sweeney (Evaluation Pilot/LMTAS Chief Test Pilot). Jeff Peer indicated that the second flight was “orders of magnitude better than the previous flight.” In Flight 2, the pilots were able to evaluate SDC performance with a missing (simulated) flap, elevon, and combination elevon/partial rudder. In all cases, the SDC performance resulted in decoupled responses that were comparable to the unimpaired performance. During pitch maneuvering with a missing left horizontal tail, the aircraft response was so like the unfailed response that there was some question as to whether the failure simulation logic, which had not been tested in Flight 1, was working properly. And, the pilots had to visually verify the lack of motion on the left tail surface during moderate to aggressive maneuvering.

While the impaired handling qualities (HQ) were comparable to the unimpaired HQ, thus achieving one of the major project goals, there were a number of areas in which the basic SDC performance could be improved. Specifically, while up-and-away unimpaired pitch response was rated Level I, there was some slight pitch-axis drift during powered approach. The cause of this drift was identified as slow convergence of integral Riccati gains. The SDC was modified, therefore, to initialize the Riccati gains at *a priori* expected values and adapt from there as opposed to initializing all gains at values prescribed by the RHO transversality conditions. This resolved the pitch drift problem. Additionally, the pilots noted that the dihedral model felt artificial due to the lack of lag between sideslip and roll responses. BAI incorporated a lag in this model so that the

aircraft did not roll immediately when the pilots kicked rudder pedal. Finally, BAI adjusted the cost function penalties to provide tighter tracking in the power-approach (PA) configuration.

### Flight 3

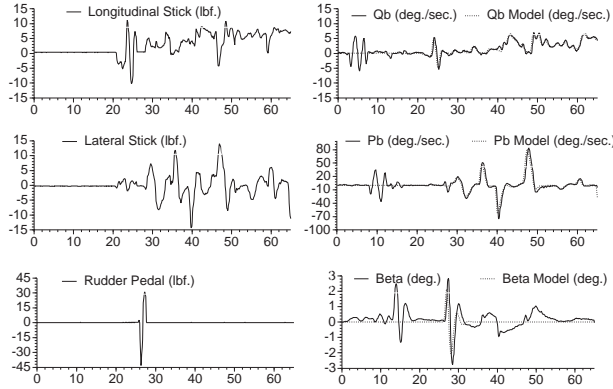
The third SDC flight occurred on June 9, 1996. The flight crew were John Ball (Safety Pilot/Calspan) and Jeff Peer (Evaluation Pilot/Calspan). Overall, the flight was deemed “a significant improvement over Flight 2” and the pilots were able to perform a number of up-and-away HQ maneuvers, including bank-to-bank rolls, pitch captures, and loaded rolls. Several single- and multiple-impairment experiments were conducted with favorable comments, and, as with Flight 2, the impaired HQ were comparable to the unimpaired case. Even for a multiple surface impairment (0 percent effective LHT and simultaneous 50 percent effective rudder) there was only slight coupling of longitudinal and lateral/directional axes. In all cases, the SDC received Level I and high Level II Cooper-Harper ratings. In one instance, loaded rolls, the impaired HQ were actually better than the unimpaired case. This is because the baseline SDC roll model was considered slightly too aggressive and the impairment reduced the roll sensitivity. During Flight 3, data were gathered on the F-16 baseline roll response during powered approach, and BAI fitted a new roll model designed to generate F-16-like behavior.

### Flight 4

The fourth SDC flight occurred on June 14, 1996, five days after the previous flight, crewed by Jeff Peer (Safety Pilot/Calspan) and Maj. Kevin Christensen (Evaluation Pilot/USAF). Though some familiarization and informal handling-qualities up-and-away maneuvers were conducted, most of this flight was devoted to powered-approach assessment and the landing tasks. Several simulated landings with aggressive offset corrections, were conducted at altitude. Intentional disengages were executed to evaluate the SDC transient response at this flight condition. Impairment responses were also evaluated. Four low approaches with the SDC were made without effector impairments. The system performed well during these approaches, and the transient responses to intentional disengages were also well within acceptable limits. Pilots considered the SDC performance acceptable, and a landing was performed, possibly the first time such has been accomplished with “full-up” parameter identification used in conjunction with a continuously adaptive flight control law.

By the fourth flight, most of the HQ models and SDC tracking penalties had been adjusted to pilot likings (although there was some difference of opinion between pilots). Fig. 7 depicts the SDC responses to general pilot inputs for the unimpaired aircraft at 268 *KCAS*; 14,560 *ft*. The SDC engagement occurred at 19.8 sec. Due to the nature of the data recording routines, channels output


by the SDC were only active while the VSS was engaged; therefore, flying-qualities models and pilot-input data prior to 19.8 sec. are to be ignored.



**Figure 7: SDC Unimpaired Response**

In Fig. 7 the first column shows the stick and rudder pedal inputs, and the corresponding pitch-rate, roll-rate, and sideslip responses are shown with their respective models in the second column. Note that the SDC provided faithful tracking of all flying-qualities models. Roll-rate model following was particularly good.

Fig. 8 is a continuation of the previous record (without a VSS disengage). After approximately 4 seconds of evaluation, a 0 percent effective LHT impairment (100interval of several seconds, so as not to trip the VIM). As before, the first column shows pilot inputs, and the second column gives the corresponding responses of the aircraft and the associated flying-qualities models.



**Figure 8: SDC Responses with Failure**

Observe that in Fig. 8 there were no significant transients upon activation of the impairment. The SDC was able to trim out the failure and achieve a high level of axes decoupling almost instantly. However, as noted by the test pilot, it took approximately 10 sec. for the SDC to remove completely the slight pitch/roll couplings encountered during

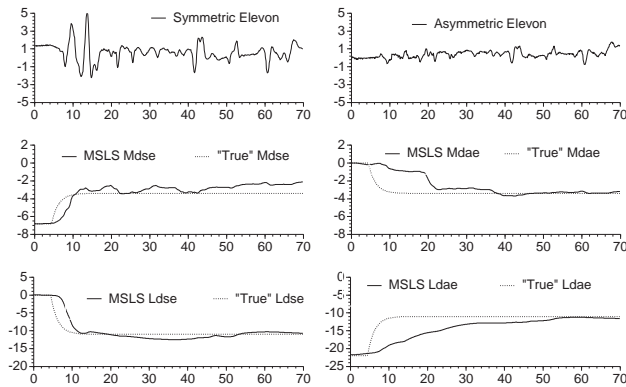
maneuvering. The SDC was stable in the presence of this impairment and was able to exhibit flying qualities comparable to those demonstrated prior to the impairment.

Pilot inputs and model-following characteristics are very similar for Figs. 7 and 8. However, at the conclusion of Flight 4, Record 4, there was a divergent nose-down transient exhibited by the SDC which led to a manual disengage of the VSS by the Safety Pilot. This was noted in Maj. Christensen’s flight report and was of particular concern to the SDC team. Post-flight analysis at BAI revealed that these pitch excursions resulted from parameter estimation error when the contribution of the symmetric elevon to pitch acceleration  $\frac{\partial \dot{Q}}{\partial \delta_{se}}$  went to zero. Inspection and analysis of these data revealed that the problem was manifested only during impaired flight following very specific multi-axes maneuvers. The erroneous pitch stability and control derivatives were readily explained when the inertial cross-coupling compensation routines were evaluated. These routines subtract known effects and nonlinear inertial terms from the equations of motion prior to identification and then add these terms back subsequent to the identification. Most of these inertial-cross-coupling (ICC) terms are proportional to ratios and combinations of flight-vehicle products and moments of inertia. Though the products and moments of inertia vary significantly with fuel quantity, the fundamental lumped inertial parameters were found to change only slightly during a flight; thus, appropriate constant values were chosen. During multi-axes maneuvers, when appreciable pitch, roll, and yaw rates existed simultaneously, the ICC terms dominated, and discrepancies between the true and assumed lumped inertial parameters translated into subtracting too much (or too little, depending on the maneuver) from the equations of motion prior to identification. The discrepancies then had to be compensated for by the MSLS algorithm, which modified its estimates to absorb these errors. (The pitch-rate equation, which was most susceptible to problems associated with the constant lumped inertial parameter assumptions, encountered problems when relatively large roll and yaw rates existed simultaneously.)

To remedy this, two modifications of the parameter identification routines were implemented. Prior to Flight 5, adjustments of forgetting factors and the spatial and temporal weighting matrices were made. The new settings were obtained automatically (off-line) using a guided-random search algorithm. Optimal MSLS settings (those which minimized the weighted-sum-squared error between key estimates and true parameters) were obtained using Flight 4, Record 4 data and evaluated on several other flight records. Flight 4, Record 4 was chosen because there was a LHT failure, and, thus, the affected parameters were known to vary significantly. Post-processed parameter identification results, using the optimal MSLS settings, are shown in Fig. 9.

The results shown in Fig. 9 have been obtained using MSLS at the same update rates realized in flight.

Note that the parameter estimates are reasonable predictors of the assumed parameters, and, more importantly, any tendency of  $\frac{\partial \hat{Q}}{\partial \delta_{se}}$  to approach zero is mitigated. Pitch and roll symmetric-elevon stability derivatives adapted rapidly to the failures, and even though asymmetric-tail motion was generally small during this record, the corresponding asymmetric-elevon stability derivatives were stable. Furthermore, despite the relatively brief periods of asymmetric-tail activity, these derivatives were able to converge to proper values. With these MSLS settings, similar improvements were observed on many other flight records at both test points.



**Figure 9: SDC Stability Derivatives Estimates**


A second modification to the identification software was an adjustment in the MSLS “clamping” routines that limit parameter estimates to predefined minimum and maximum values. These extrema were chosen to encompass all values that the “true” parameters can possibly take on at any combination of loading, mass property configurations, flight condition, and potential failure scenarios. Previously, the upper extremum for the contribution of symmetric elevon to pitch acceleration was set at 0.0. The only circumstance where this parameter has a zero contribution to pitch acceleration is when both left and right elevons are simultaneously 0 percent effective. Because such a failure scenario does not allow reconfiguration on the F-16 and was not in the scope of these flight tests, the upper limit for this parameter was modified and set at  $-1.00$ ; other parameter limits were modified similarly. Nevertheless, with the new parameter identification settings, the parameters rarely encountered these limits on flight data.

### Flight 5

Due to mechanical problems with the VISTA aircraft, the fifth SDC flight did not occur until July 9, 1996. For the final flight, the crew were John Ball (Safety Pilot/Calspan) and Jeff Peer (Evaluation Pilot/Calspan). Weather conditions were overcast with isolated pockets of light turbulence, confined mainly near scattered cloud formations. There was an appreciable (15 *kts.*) crosswind at Niagara International Airport, where the landing with 0 percent effective

left horizontal-tail surface was to be performed; however, this was just within the safety limits agreed upon in the flight test plan. Landing build-down tasks with simulated failures were performed at altitude and demonstrated impressive results. After numerous multi-axes maneuvers with impairments, it was determined that all pitch-divergence tendencies had been eliminated. Three low approaches with a 0 percent effective LHT were flown; offset and correction maneuvers were evaluated, and the disengage transient responses were assessed. These transients were well within the comfort level of both pilots. A final approach was made, which culminated in *landing the VISTA with a 0 percent effective LHT under SDC control with full parameter identification*. To the author’s knowledge, this is the first time a damaged aircraft has been landed under reconfigurable control.

Fig. 10 shows the pitch-axis maneuvering and elevon commands during landing with (simulated) 100 percent missing left elevon. Time histories begin 35 sec. into the flight record with SDC engagement and failure activation and terminate at touchdown approximately 150 sec. into the record. Throughout the approach to final touchdown, the SDC demonstrated faithful tracking of the body-axes pitch-rate model, in spite of a (simulated) missing half-tail surface. Notice that the left elevon is floating at the local angle of attack. The pitching moments for control were being generated by the healthy right-half-tail surface. With the ineffective control surface, a straight-in landing was executed with minimal lateral stick and pedal inputs.



**Figure 10: SDC Landing with Missing Left Elevon**

The SDC appropriately removed couplings between the longitudinal, lateral, and directional axes resulting from the simulated impairment. Careful inspection of Flight 5, Record 12 data reveals that body-axes roll-rate and sideslip excursions were well below  $\pm 1.5 \text{ deg./sec.}$  and  $\pm 1.0 \text{ deg.}$ , respectively. Despite the effects of a simulated failure being compounded by a significant crosswind (15 *kts.*) and gusty conditions during the landing, flying qualities were satisfactory for all axes, and, as noted by the flight

crew, SDC performance was consistent and predictable for the entire flight.

## CONCLUSIONS

All of the technical objectives were achieved with the experimental SDC. Successful flight tests of the SDC culminated in a smooth landing of the VISTA/F-16 in cross-wind conditions with a (simulated) missing primary control surface (left horizontal tail). Up-and-away flight tests demonstrated excellent reconfiguration capabilities for numerous impairments, including a (simulated) completely missing flaperon, a completely missing half-tail surface, a partially missing rudder, and simultaneously missing half-tail and (partial) rudder. In both the up-and-away flight and the landings, the pilots noticed little difference between the unimpaired and impaired flying qualities, and adverse axial cross-couplings were almost completely eliminated.

The SDC computes a time-varying model of the current aircraft dynamics using a novel on-line system identification technique that can rapidly track time-varying parameters and is robust to adverse conditions such as low excitation or correlated inputs. This identified model of the current aircraft dynamics is then communicated to an optimal control module that computes the effector commands which will give the desired aircraft response as specified by a set of desired flying-qualities models. The combination of reconfigurability and model-following capabilities in the SDC can result in improved flight safety with, potentially, less development time. The SDC approach offers particular advantages over reconfigurable flight control methods that rely on *a priori* design of both detection filters and control modes tailored for specific failure scenarios.

### Recommendations for Future Research

While BAI believes that the SDC represents an important milestone toward the automated design of adaptive and reconfigurable controllers, further research is needed to take full advantage of the potential of these algorithms to provide a useful and cost-effective control design methodology for future aircraft. Specifically, in the area of on-line parameter identification, large numbers of effectors may give rise to data collinearities that will be significantly worse than anything encountered in the F-16 SDC development effort. Research into automated “ganging” of effectors, possible use of active noise injection, and further parameter identification algorithm regularization will be required. Effector nonlinearities and interactions may require changes in the way one constructs the model of the plant to be identified. On-line learning or adjustment of *a priori* aircraft models could be incorporated to “fine-tune” the spatial constraints used for parameter identification. Currently, the MSLS algorithm will slowly drift toward the spatial constraints, if they are specified and if there is not sufficient excitation to override them. Finally, automatic adjustment of parameter identification settings would help in the implementation of the ID algorithm. BAI took a step in this direction

prior to the final flight, where a random optimization technique, similar to simulated annealing or evolutionary programming, was used to find optimal forgetting factors, temporal constraint weights, and spatial constraint weights using data from the first four flights. While MSLS is very robust to the specific values chosen, the SDC can benefit from increased automaticity in this area.

Regarding the optimal control law, additional design constraints can be included in computation of the optimal control signals. Possible constraints include minimization of radar cross section, maneuver-load alleviation, or drag reduction. Inclusion of such constraints will require modifications to the RHO control law derivation, and their impact on the stability, robustness, and flying qualities of the control law will need to be investigated. Automated command limiting is also possible as part of the control law optimization. Research into using inner-loop constrained optimization to limit pilot commands automatically is currently in progress [14]; further work is needed to examine implementation and flying-qualities issues associated with these approaches. For instance, while inner-loop stability may be improved, will the overall system be subject to the same types of pilot-induced oscillation (PIO) tendencies that are presently associated with actuator limits? Control allocation and control-axes prioritization issues have not been fully addressed in the SDC; these will become increasingly important as novel future aircraft are designed that are unstable in several axes and utilize effectors capable of generating moments about multiple control axes simultaneously. Automatic adjustment of LQ tracking and integral penalty terms would significantly improve the “self-designing” aspect of the controller. Currently, there are no clear guidelines for setting these parameter, and for establishing the horizon length. The SDC did, however, incorporate an effective heuristic method for adjusting actuator penalties on-line, and it may be possible to extend the method to cover all of the RHO design parameters. Finally, flying-qualities specifications that are designed specifically to be utilized by model-following controllers would be a significant help to the overall controller design procedure. While the SDC demonstrated that adaptive model-following control was possible, the models that were being followed did not always result in Level I flying qualities, and, as summarized above, many of the adjustments to the algorithm between flight tests consisted of modifications to the models in response to pilot comments. Nevertheless, the model-following approach provided a natural and intuitive way to *tune* the controller in response to pilot comments.

Finally, during the course of the SDC development, a number of practical steps were taken to prevent undesirable behavior on the part of the parameters and controller gains. The efficacy of these methods was demonstrated in the successful flight tests, which increased the confidence in this type of adaptive control law. Future efforts will benefit from further research into both theoretical and practical means

of ensuring the stability, robustness, and overall safety of adaptive controllers of this nature.

### Summary

The SDC is the fruit of decades of research by BAI and many others into parameter identification, optimal control, reconfigurable control, adaptive control, and nonlinear modeling. The techniques described in this report are suitable for control of many high-performance systems that utilize complex effector suites, such as high-agility aircraft, flexible aircraft with significant airframe-effector coupling, experimental aircraft with novel effectors, robotic systems, and chemical processes. The adaptive model-following approach can reduce control law development costs and improve system performance in the presence of gradual or abrupt changes, including unforeseen events. BAI believes that the successful SDC design and flight test results reported herein represent an important demonstration leading to a new generation of adaptive model-following controllers.

### REFERENCES

- [1] D. Ward, J. Monaco, R. Barron, R. Bird, J. Virnig, and T. Landers, "Self-Designing Controller: Design, Simulation, and Flight-Test evaluation," Final Tech. Rep. for Air Force Office of Scientific Research, Contract F49620-94-C-0087, Barron Associates, Inc., Nov. 1996.
- [2] T. Söderström and P. Stoica, *System Identification*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- [3] V. Cherkassky, J. Friedman, and H. Wechsler, *From Statistics to Neural Networks, Theory and Pattern Recognition Applications*. Berlin: Springer-Verlag, 1993.
- [4] L. Ljung and T. Söderström, *Theory and Practice of Recursive Identification*. Cambridge, MA: MIT Press, 1983.
- [5] G. Seber and C. Lee, *Advances in Neural Information Processing Systems*. New York, NY: Wiley and Sons, 1989.
- [6] M. Bodson, "An adaptive algorithm with information-dependent data forgetting," in *Proc. 1995 American Control Conf.*, (Seattle), June 1995.
- [7] C. Garcia and D. Prett, "Model predictive control: Theory and practice - a survey," *Automatica*, vol. 25, pp. 335–348, 1989.
- [8] H. Michalska and D. Mayne, "Robust receding horizon control of constrained nonlinear systems," in *Trans. on Automatic Control, IEEE*, vol. 38, pp. 1623–1633, Nov. 1993.
- [9] R. Bitmead, M. Gevers, and V. Wertz, *Adaptive Optimal Control*. Englewood Cliffs, NJ: Prentice Hall, Inc., 1990.
- [10] B. Anderson and J. Moore, *Optimal Control: Linear Quadratic Methods*. Englewood Cliffs, NJ: Prentice Hall International, 1989.
- [11] P. Chandler, M. Pachter, and M. Mears, "System identification for adaptive and reconfigurable control," *Guidance, Control, and Dynamics*, vol. 18, pp. 516–24, 1995.
- [12] U.S. Dept. of Defense, "MIL-STD-1797A, *Military Standard: Flying Qualities of Piloted Aircraft*," Jan. 30, 1990.
- [13] D. Ward, R. Barron, J. Monaco, and R. Bird, "Neural Network Flight Control System for High-Agility Air Combat," Final Tech. Rep. for Flight Dynamics Directorate, Wright Laboratory, Contract F33615-93-C-3612, Barron Associates, Inc., Feb. 1996.
- [14] P. Chandler, M. Pachter, and M. Mears, "Constrained linear regression for flight control system failure identification," in *Proc. Amer. Control Conf.*, (Seattle, WA), June 1993.