AFRL-VA-WP-TP-2003-330

INTELLIGENT CONTROL OF UNMANNED AIR VEHICLES: PROGRAM SUMMARY AND REPRESENTATIVE RESULTS

D.G. Ward
M. Sharma
N.D. Richards
J. De Luca
M. Mears

SEPTEMBER 2003

Approved for public release; distribution is unlimited.

20031103 053

AIR VEHICLES DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542

REPORT DOCUMENTATION PAGE

1. REPORT DATE (DD-MM-YY)	2. REPORT TYPE	3. DATES COVERED (From - To)
September 2003	Conference Paper Preprint	

4. TITLE AND SUBTITLE
INTELLIGENT CONTROL OF UNMANNED AIR VEHICLES: PROGRAM SUMMARY AND REPRESENTATIVE RESULTS

5a. CONTRACT NUMBER
F33615-01-C-3108

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
62201F

6. AUTHOR(S)
D.G. Ward, M. Sharma, and N.D. Richards (Barron Associates, Inc.)
J. De Luca (Northrop Grumman)
M. Mears (AFRL/VACA)

5d. PROJECT NUMBER
2401

5e. TASK NUMBER
02

5f. WORK UNIT NUMBER
9C

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Barron Associates, Inc.
Charlottesville, VA

AFRL/VACA

Northrop Grumman
El Segundo, CA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Vehicles Directorate
Air Force Research Laboratory
Air Force Materiel Command
Wright-Patterson Air Force Base, OH 45433-7542

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/VACA

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-VA-WP-TP-2003-330

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

14. ABSTRACT
In early 2001, AFRL and NAVAIR issued a PRDA requesting proposals to develop an intelligent controller (IC) for unmanned combat air vehicles. Two key requirements of the IC were (1) a learning approach that could go beyond current adaptive controllers and "remember" what it had learned across flight conditions and (2) a reconfigurable path planner that accounted for changes in the inner-loop behavior and generated near-optimal trajectories in real time. This paper presents a summary of the resulting IC program and some initial technical results.

15. SUBJECT TERMS
Intelligent Control, function approximation, back-stepping, UAV

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT:	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON (Monitor)
a. REPORT	b. ABSTRACT	c. THIS PAGE	SAR	16	Mark J. Mears
Unclassified	Unclassified	Unclassified			19b. TELEPHONE NUMBER (Include Area Code) (937) 255-8685

INTELLIGENT CONTROL OF UNMANNED AIR VEHICLES: PROGRAM SUMMARY AND REPRESENTATIVE RESULTS

D. G. Ward[*], M. Sharma[†], N. D. Richards[‡]
Barron Associates, Inc.
Charlottesville, Virginia

J. De Luca[§]
Northrop Grumman,
El Segundo, CA

M. Mears[¶]
Air Force Research Laboratory
Wright-Patterson AFB, OH

ABSTRACT

In early 2001, AFRL and NAVAIR issued a PRDA requesting proposals to develop an intelligent controller (IC) for unmanned combat air vehicles. Two key requirements of the IC were (1) a learning approach that could go beyond current adaptive controllers and "remember" what it had learned across flight conditions and (2) a reconfigurable path planner that accounted for changes in the inner-loop behavior and generated near-optimal trajectories in real time. This paper presents an summary of the resulting IC program and some initial technical results. Key features of the IC architecture are (a) a direct-adaptive backstepping controller that uses spatially-local models of the vehicle dynamics, (b) a provably-stable approach to learning the structure of the underlying vehicle models online, and (c) a finite-automaton-based path planning approach that computes an near-optimal trajectories using pre-computed maneuver and trim primitives. The IC architecture not only provides on-line inner- and outer-loop reconfiguration for unforeseen failures or damage, but it can also reduce the cost of developing new control systems. To demonstrate this assertion, the IC algorithms were developed using a medium-fidelity UCAV simulation and subsequently evaluated using a high-fidelity nonlinear simulation that was similar in nature but significantly different in detail to the development simulation.

INTRODUCTION

Uninhabited air vehicles (UAVs) and uninhabited combat air vehicles (UCAVs) will play an increasingly important role in future military operations; however, there are a number of significant challenges associated with the development of an advanced control system for these vehicles.

First, because the UAV will be exploited to perform tasks that would otherwise risk the safety of flight crews of manned aircraft, there is an increased probability of damage to the vehicle might resulting from extreme operating conditions, hostile actions, etc. This underscores the need for a reliable system design that can accommodate significant changes in system behavior from a wide variety of sources. The requirement that the UAV must operate in close proximity to humans further emphasizes the need for a reliable system design.

Second, because many UAV systems are expected to cost less than manned systems, it is unlikely that developers will have the resources to collect extensive wind-tunnel and flight-test data of the caliber typically found during manned flight vehicle development. Thus, the model available for UAV development will necessarily contain larger uncertainties, which compels the controls engineer to compromise performance in favor of robustness.

Finally, because there is not always a human in the loop, the controller must be augmented with a very sophisticated autopilot design that not only cruises, climbs, and changes heading, but is capable of performing complex and agile maneuvers, that would normally be performed by a pilot, without the risk of losing control of the vehicle.

In recent years, there have been considerable advances made in developing control methods that enhance fault-tolerance and survivability of fixed-wing manned aircraft.

A number of researchers have developed reconfigurable control systems for a variety of flight vehicles with promising results [1,2,3,4,5]. Many of these results have been demonstrated in high-fidelity simulations; the past decade or so has also witnessed four significant flight demonstrations of reconfigurable control. The first of these, the Self-Repairing Flight Control System (SRFCS) [6], culminated in a series of F-15 flight tests that demonstrated the ability of the controller to isolate individual control surface failures and subsequently

---

American Institute of Aeronautics and Astronautics

reconfigure the aircraft. The Self-Designing Controller (SDC), funded by AFOSR and led by BAI was another milestone in reconfigurable control [1]. Here, an on-line control design was used to avoid having to make *a priori* assumptions about the nature of potential failures. To capitalize on the SDC results and further advance reconfigurable controls technology, the Air Force Research Laboratory (AFRL) initiated the *RESTORE* program for tailless fighter aircraft [5,7]. Here, two designs were evaluated; a significant result of [7] – and another landmark in reconfigurable control research – was successful flight testing of a direct-adaptive neural network reconfiguration architecture on X-36. Finally, NASA's F-15 Intelligent Flight Control System (IFCS) is the most recent reconfigurable control program to involve flight tests.

Much of the prior work in reconfigurable control has focused on modifying inner-loop controllers to achieve, to the extent possible, the desired response characteristics. However, there has been some work to address reconfiguration at the level of guidance and trajectory loops. Of specific relevance to current program are efforts that have investigated on-line computation of optimal trajectories and the modification of inner-loop reference commands. The latter is particularly important when inner-loop reconfiguration alone cannot recover nominal performance, and the outer loop (pilot or autopilot) must modify its behavior to ensure safe performance. In [8], the authors leveraged experience in rotorcraft pilot cueing to develop algorithms and methods that provided pilots with physical cues as to the limitations of the inner-loop controller via force feedback on the control inceptor(s) and demonstrated the ability to mitigate PIO using such a system. In [9], the predictive nature of a model-predictive-control algorithm was used to compute a modified reference command to match the performance capabilities of the aircraft. In [10], an adaptive command gradient was used to ensure that the reference command provided to the controller was realizable. Key outer-loop and closed inner-loop characteristics of a failed autonomous reusable launch vehicle (RLV) were identified on-line and used to ensure the stability of the existing guidance loop as well as to generate, in real time, new optimal trajectories that would result in a safe vehicle landing [11]. This work is particularly relevant to the proposed effort because the RLV has minimal inner-loop control redundancy and, so, the guidance and autopilot loops were required to adapt intelligently to handle failures.

The opportunity exists to extend these reconfigurable methods to address the following unique controls challenge posed by UAV systems:

> How does one design an integrated autopilot and inner-loop controller to maximize performance *and* reliability given the uncertain nature of the vehicle models available for controller synthesis?

The answer lies in (a) extending the existing state of the art in trajectory (re-)computation and reconfigurable control designs to incorporate on-line learning that *remembers* what is learned about the vehicle behavior, and (b) developing a modular system that is *robust to adverse interactions* between the autopilot and inner-loop controller.

THE IC PROGRAM
AND SYSTEM ARCHITECTURE

To address these two research issues of inner-loop learning and outer-loop reconfiguration, AFRL and NavAir issued a PRDA in 2001 entitled *Intelligent Control* that sought "a combination of methods which include learning to recognize and remember spatial dependencies, adaptation to address abrupt changes, and optimization to determine optimal trajectories for specific tasks or mission requirements."

In response to this PRDA, Barron Associates, Inc. (BAI) proposed a program with the following technical objectives:

Table 1. Intelligent Control (IC) Technical Objectives	
Long-Term Learning	IC performance will improve over time as it learns based on *observed* behavior of the UAV. This will significantly reduce the need to develop expensive, high-fidelity math models during the controller design process.
Rapid Adaptation	The IC will rapidly adapt to any *sudden* unforeseen change in vehicle dynamics due to failures, stores release, etc.
On-Line Trajectory Reshaping	The IC will *interact* with the mission planner by receiving a request to follow a trajectory or fly to a destination and return a *feasible* trajectory that is as close to the desired or optimal trajectory as possible given the current capabilities of the UAV.
Implement-able and Verifiable	The IC algorithms will be modular, computationally feasible, and have stability proofs that allow them to be implemented, verified, and validated. These algorithms will be transitioned to NGC and other airframers for use in CMUS and future production UAVs.
Reduced Develop-ment Costs	The proposed IC algorithms can be developed with lower cost medium fidelity simulations, *and* they can be reused on new or derivative airframes more readily by relearning the control in a new simulation, thus reducing the amount of analyst involvement required to fine-tune the controller.

BAI believed that its background in indirect-adaptive receding-horizon optimal control [1,12], adaptive backstepping control [13,14], structure-learning modeling [15], reinforcement learning [16], and guidance and trajectory adaptation [11] provided a good foundation for addressing the requirements of the IC PRDA; however, it was deemed essential to augment these skills with those of additional team members from government, industry, and academia. Figure 1 shows the team members and the technology they contributed to this project.
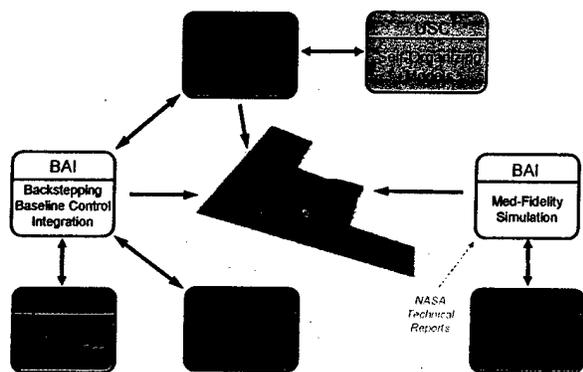

Figure 1: Intelligent Control Team Contributions

Dr. Jay Farrell (UC Riverside) provided methods that augmented adaptation with spatially local learning to give the controller with "memory." Additionally, these methods employ a Lyaponuv-based learning rule that ensures the stability of the overall closed-loop tracking system. Working closely with UCR and BAI, Marios Polycarpou (UC), worked on anti-windup approaches that could be integrated with the learning controller and preserve the stability under effector saturation. Dr. Stefan Schaal (USC) provided expertise in self-organizing approaches that could learn both parameters and the *structure* of a vehicle model online. Dr. Eric Feron (MIT) provided finite-automaton-based trajectory reconfiguration that computed an optimal path using pre-computed maneuver and trim primitives. Table 2 summarizes the benefits of these technologies.

In addition to the academic input, it was important to have the input of researchers actually involved in the design of UAV systems for real-world missions. Northrop Grumman Corporation (NGC) filled this need by providing expertise in identifying mission scenarios that would challenge the IC architecture, providing key insights into the V&V issues of this type of system, providing a high-fidelity UCAV simulation and an associated baseline (non-reconfigurable) controller, and reviewing and commenting on the control design at each step of the process to ensure that the overall team

received the benefits of NGC's extensive "real-world" experience.

Table 2. IC Technical Approach	
Memory via Locally Weighted Learning (LWL)	Both the inner-loop and outer-loop algorithms incorporate *memory* using a spatially-local nonparametric modeling approach (LWL) originally developed by the machine learning and robotics community. LWL learns the structure and weights of a model of the system dynamics in real time with guaranteed convergence properties. The IC represents the first time LWL has been applied in the context of *direct-adaptive* learning control.
No Adaptation / Learning Tradeoff	Because of the spatially-local nature of LWL, the number of non-zero basis functions (and, hence, coefficients to be learned) is quite small at any specific flight condition. Thus, the models are capable of adapting very rapidly to sudden changes due to failures, stores release, etc. Moreover, once the vehicle has moved to another flight condition, these dynamics are retained for later use.
Novel Outer-Loop Approach that Accounts for Interactions	The IC architecture will employ a trajectory reshaping algorithm based on MIT's hybrid automaton approach. This approach, originally developed for DARPA's SEC program, will be extended by applying it to *fixed-wing* vehicles (as opposed to rotorcraft), accounting for changing vehicle dynamics, and incorporating on-line non-real-time learning of new optimal trajectories that maximally utilize all of the available UAV performance capabilities.
Accounts for Loop Interactions	The IC architecture is designed to minimize adverse interactions between the inner and outer loops by identifying, remembering, and accounting for system dynamics that might give rise to such interactions, including saturation nonlinearities.
Guaranteed Stability	Both inner-loop and outer-loop algorithms have stability guarantees comparable to those associated with classical control design methods.

To compellingly demonstrate the above technical objectives of the IC approach, the team employed a novel strategy in which the algorithms were developed initially using a medium-fidelity MATLAB simulation. Final evaluation, however, was performed using the high-fidelity uninhabited combat air vehicle (UCAV) simulation provided by Northrop Grumman Corporation (NGC). Not only did this simulation have dynamics that are significantly different from those used in the initial design, but it had two additional features of interest: (1) because the NGC UCAV was designed for stealth, it had a reduced effector set and limited inner-loop reconfiguration options. Thus any reconfigurable controller was likely to require outer-

loop learning and adaptation to maintain stability during unforeseen changes in dynamics; (2) Because it contained a respectable baseline controller developed over 1.5 man years using a classical design approach, it provided important performance benchmarks for the missions and maneuvers of interest. Key advantages highlighted by the proposed demonstration approach are summarized in Table 3.

Table 3. IC Demonstration Approach

Reduced Development Cost	Demonstrate that a suitably performing IC system could be developed, initially, using a very low-cost, medium-fidelity simulation.
Performance Comparable to Classical Control for Known Model	Demonstrate that, once the IC has *learned* the dynamics of the high-fidelity UCAV, its performance is comparable to that of a high-quality classically-designed controller.
Improved Performance with New Vehicle Dynamics	Demonstrate that the IC approach, as it learns the vehicle dynamics, can use more aggressive control to achieve performance that approaches the limits of the vehicle's abilities. This is in contrast to a conventional controller that, to deal with large uncertainties, would have to be very robust and conservative.
Improved Performance with Failures	Demonstrate that, compared to the baseline controller, the IC approach can rapidly adapt and maintain stability for significant set of failures.
Integrated Inner- and Outer-Loop Learning	Demonstrate the benefits of an intelligent outer-loop that modifies reference trajectories and inner-loop commands to ensure stability *and* model following even when sufficient control redundancy or authority does not exist to achieve desired inner-loop performance.

MEDIUM-FIDELITY SIMULATION

There are two vehicle models used in the IC program. The first is a high-fidelity model of flying-wing UAV provided by Northrop Grumman Corp. (NGC) and the second is the Barron Associates Nonlinear Tailless Aircraft Model (BANTAM). The latter is a medium-fidelity model that resembles the NGC model in configuration only, but was constructed completely independently using public-release aerodynamics data unrelated to NGC UAV model.

The primary source of aerodynamic data used in BANTAM is NASA TM-4640, which is a wind-tunnel test report on a series of flying wings. Both DATCOM and HASC-95 were used to fill the gaps in TM-4640, and WL-TR-97-3059 (ICE) was used to obtain data on the spoilers and their interactions with the other control effectors.

Tables 3 and 4 cite the sources of the aerodynamic force and moment data used for BANTAM.

Table 4: Sources of Aerodynamic Force Data

Parameter	Source
C_{L_o}	DATCOM
$C_{L_{wing+body}}(\alpha)$	TM 4640
$C_{L_{\delta,flap}}(\alpha, \delta_{spoiler})$	TM 4640
$C_{L_{\delta,spoiler}}(\alpha, \delta_{spoiler})$	WL-TR-97-3059
$C_{L_q}(\alpha)$	HASC-95
C_{D_o}	DATCOM
$C_{D_i}(\alpha, \delta)$	$0.133C_L^2$
	$0.0643 + 0.0217(\alpha - 16°)$
	$C_L \tan \alpha$
$C_{Y_\beta}(\alpha)$	TM 4640
$C_{Y_p}(\alpha)$	HASC-95
$C_{Y_{\delta,flap}}$	TM 4640

Table 5: Sources of Aerodynamic Moment Data

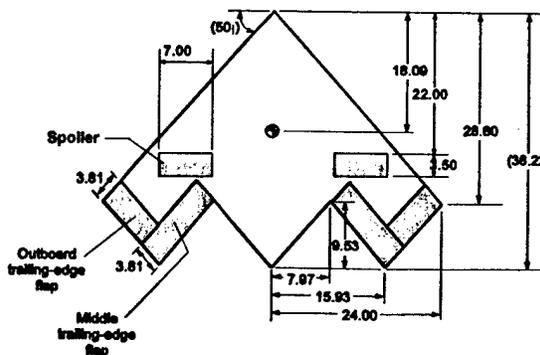| Parameter	Source
$C_{l_p}(\alpha)$	HASC-95
$C_{l_r}(\alpha)$	HASC-95
$C_{l_{\delta,flap}}(\alpha, \delta_{spoiler})$	TM 4640/WL-TR-97-3059
$C_{l_{\delta,spoiler}}$	WL-TR-97-3059
C_{m_o}	DATCOM
$C_{m_{wing+body}}(\alpha)$	TM 4640
$C_{m_{\delta,flap}}(\alpha, \delta_{spoiler})$	TM 4640
$C_{m_{\delta,spoiler}}(\alpha, \delta_{spoiler})$	WL-TR-97-3059
$C_{m_q}(\alpha)$	HASC-95
$C_{n_p}(\alpha)$	HASC-95
$C_{n_r}(\alpha)$	$-C_{D_o}/6$ (strip theory)
$C_{n_{\delta,flap}}(\alpha, \delta_{spoiler})$	TM 4640/WL-TR-97-3059
$C_{n_{\delta,spoiler}}$	WL-TR-97-3059

Figure 2: BANTAM Vehicle Configuration

The BANTAM simulation is based on a flying wing configuration representative of UAVs proposed for

- 4 -

near-term in-flight demonstration (see Figure 2). The all-wing airframe provides many benefits such as stealth, low wing-loading, high fuel volume, and greater aerodynamic efficiency than traditional wing-fuselage configurations, however it does pose several control challenges including (a) low yaw authority due to the airframe configuration (b) a reduced effector set consisting of midboard and outboard body flaps and spoilers and (c) effector interactions due to the fact that are mounted directly upstream of the midboard flaps and cause a significant reduction in midboard flap control power when deployed. The BANTAM simulation used second-order models for actuator dynamics with rate and position limits representative of this type of vehicle.

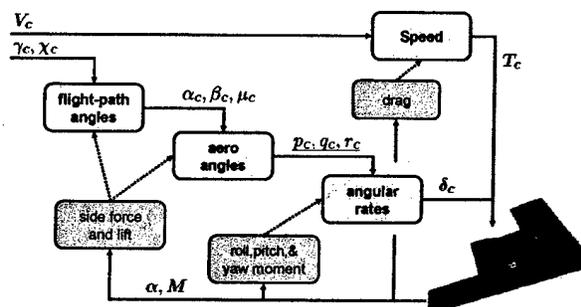have memory properties, and hence *learns* rather simply adapting.

Figure 3: Inner-Loop Architecture

CONTROLLER

Backstepping control takes advantage of the fact that certain states can be used as virtual controls for other states. In effect this results in generating aero-angle commands to meet tracking of flight-path variables, followed by computing required angular rates to follow the aero-angle commands from the previous step, and finally the control surface deflections required for achieving the required angular rates. This is essentially the same as constructing three loops, an inner, a middle, and an outer corresponding to the rate loop, the aero-angle loop, and the flight path loop respectively, as is commonly done in flight control. The advantage of backstepping is that it accounts for the transients in the virtual command, and thus does not require an explicitly time-scale separation assumption. A block diagram of the inner-loop control architecture is given in Figure 3. Table 6 expands upon the inputs and outputs used for each loop. Details of the IC backstepping controller can be found in [17], [18], and [19].

LEARNING

Nonlinear flight-control algorithms such as feedback-linearization and backstepping require accurate knowledge of the plant parameters for successful implementation, and thus some form of adaptation or learning is required to provide robustness to uncertain or altered aerodynamics. Here, learning is distinguished from adaptation in that learning algorithms have memory in the sense that they retain information across multiple flight conditions. This is accomplished through the use of function approximators with local support, i.e. the approximator parameters are adjusted only locally at any given time. The local function approximation does not interfere with the approximation at points outside a closed neighborhood. Thus, the approximator is considered to

Table 6: Inner-Loop Input/Output Signals

Control Loop	Inputs			Outputs
	Commands	Feedback signals	Learned parameters	
Flight path	$\gamma_c, \chi_c,$ V_c, Y_c	$\gamma, \chi,$ V, Y	Lift, Drag, Side-force	$\alpha_c, \beta_c,$ $\mu_c,$ Thrust
Aero angle	α_c, β_c μ_c	(u, v, w) (φ, θ, ψ)		p_c, q_c, r_c
Body rate	p_c, q_c r_c	p, q, r	pitch, roll, yaw moment	pitch, roll, yaw pseudo-controls $(\delta_e, \delta_a, \delta_r)$

B-Splines form the core of the IC function approximators. Some of the advantages afforded by these splines are:

- Local support: splines actually go to zero outside their domain, so only k^2 splines are non-zero during any given time step. (k = spline order; usually, k = 3 is used).

- The spline outputs are always positive and normalized, which provides numerical stability

- The algorithms for computing the spline outputs are computationally efficient.

- The number of splines, and their sizes and centers can either be determined *a priori* (e.g., laid out on a grid), or adapted on-line. The latter is known as structure learning.
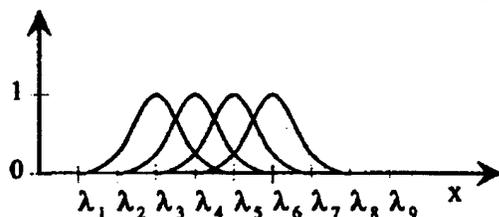

Figure 4: Third-Order B-Splines

The function approximators are linear-in-the-parameters, i.e. the approximated function is expressed as $\hat{f} = \theta^T \phi$ where

- θ is a vector of weights

- ϕ is a regressor vector containing the basis functions. The regressor is predefined by the designer

Figure 5 shows the structure of the function approximator update.
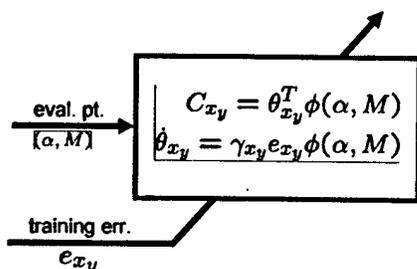
Figure 5: Function Approximator Operation

In the IC architecture, the structure of the splines used to model the stability and control derivatives is fixed as a two-dimensional function of α and *Mach* with a pre-specified number of knots. An exception to this are the $\delta_{mid\,flap}$ control derivatives which are also functions of $\delta_{spoiler}$, i.e.

$$C_{x_{mid\,flap}} = \theta^T_{x_{mid\,flap}} \phi\left(\alpha, M, \delta_{spoiler}\right)$$

One of the advantages of the IC architecture is that unlike most other direct-adaptive control approaches which modify a control gain or compensation parameter directly, the IC function approximators learn the non-dimensionalized stability and control derivatives for all of the aerodynamic forces and moments.

There are three possible signals that can be used to update the function approximators.

- Tracking error – i.e. true direct adaptive control. This method guarantees bounded command-tracking, but function convergence can be very slow.

- Function approximation error – i.e. indirect adaptive control. With this method, function approximation is improved, however command tracking may be quite poor.

- Composite error – a blend of direct and indirect adaptive control. This approach provides both guarantees on command tracking as well as improved functional converges. The ICLAWs use composite error to train the function approximators.

Input saturation must be considered in any adaptive or learning control system since the adaptive and learning elements are essentially integrators that will wind up in the event of input saturation. The ICLAW employs a simple solution to remove the effect of saturation from the learning function approximators' training error, thereby preventing wind-up in the case of magnitude, rate, or bandwidth constraints [20,21].

In many cases, the structure of the underlying aerodynamic model is known relatively well and so determining the structure of the spline approximators ahead of time is not a problem. In some cases, however, it may be desirable to learn the *structure* of the spline approximators as well as the coefficients. To address this, the IC team developed a composite error update rule for structure-learning locally-weighted-linear (LWL) approximators and associated stability proofs [22,23].

PATH PLANNING

The path planner is tasked with generating feasible earth-axis (NED) flight paths on-line. For computational tractability, the planner discretizes the maneuver space into a grid, with travel between grid points performed via interpolation. There are two basic components of the path planner: one that is computationally intensive and generated *off*-line, and one that is used *on*-line to rapidly construct a trajectory using the stored information.

The off-line component itself consists of a maneuver automaton, which forms the core of the path planner, and describes all of the feasible trajectories the aircraft can take. is the automaton is composed of two libraries, one containing feasible trims, which are defined as constant velocity trajectories (including steady-state turn), and the other containing maneuvers, which are finite-time transitions between trims. The automaton is

- 6 -

generated using a closed loop simulation of the aircraft, and therefore accounts for all of the aircraft nonlinear dynamics.

A value-iteration is performed off-line to determine costs associated with traveling from any point in the maneuver-space grid to the origin in a trim (for every maneuver, the reference grid is defined such that the goal state or waypoint is considered the origin). On-line trajectory generation is akin to solving a dynamic programming problem, and the resulting optimal trajectory is simply a sequence of trims and maneuvers stored in the automaton. Furthermore, the trajectory is guaranteed to be feasible since the maneuver automaton and the value function are generated using the nonlinear vehicle dynamics.

Figure 7: Reconfigured Automaton

Figure 6: Maneuver Automaton

The inputs that the planner expects from the mission planner are

1) Waypoints in (NED) coordinates (x,y,z)
2) Target trims in which to reach the destination. $(V, \gamma, \chi, \dot{\chi})$

and the outputs are: (V, γ, χ) commands

A critical piece in the integration of the path planner and the inner loop is their interface. During nominal operating conditions, the path planner simply provides inner-loop commands. However, in the event of unanticipated actuator saturation or damage, the closed-loop capabilities of the airframe and/or its dynamics change. This change in dynamics is represented in the automaton as a set of discrete trims and maneuvers that are no longer feasible (Figure 7). In this case the trajectory generation algorithm is no longer "optimal;" however the finite automaton will still generate and regenerate feasible trajectories quite rapidly while, in the background, the costs-to-go associated with each trim point in the automaton are updated. Figure 8 gives a summary of the entire reconfiguration process.
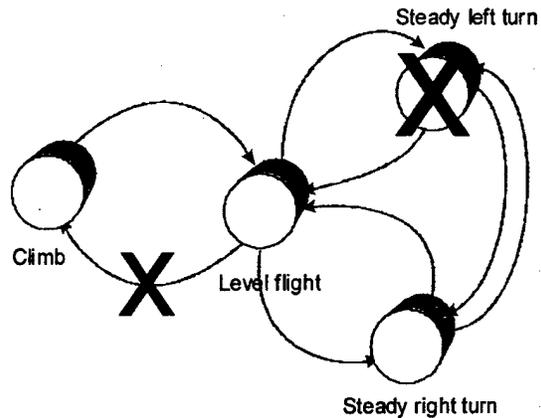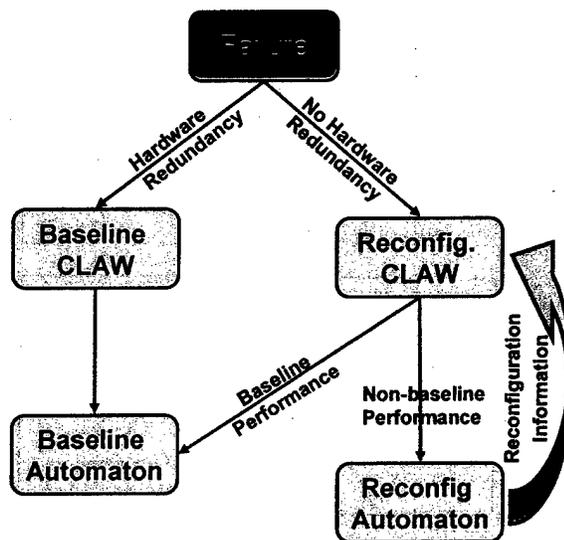

Figure 8: Inner/Outer Loop Reconfiguration

REPRESENTATIVE SIMULATION RESULTS

At present, the IC control software has been developed, implemented, evaluated in BANTAM, and is being ported to the high-fidelity simulation environment. This section presents some representative results.

To test the inner-loop reconfiguration, the authors drove it with aggressive command sequences that are representative of the kinds of commands that are required for tasks such as missile evasion or NOE flight. Figure 9 shows the inner-loop response to the externally-generated flight-path-angle commands. As shown in Figure 3, the backstepping controller uses aerodynamic angle commands to track the flight-path angle commands, angular rate commands to track the

- 7 -

aerodynamic angle commands, and uses effector commands to track the angular rate commands. Figure 10 and Figure 11 show commands generated and tracking achieved by these inner backstepping loops.
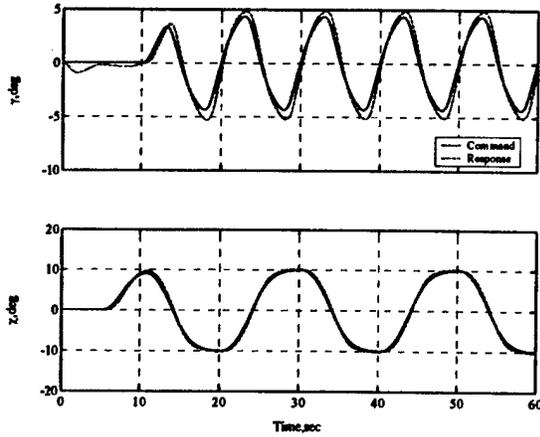

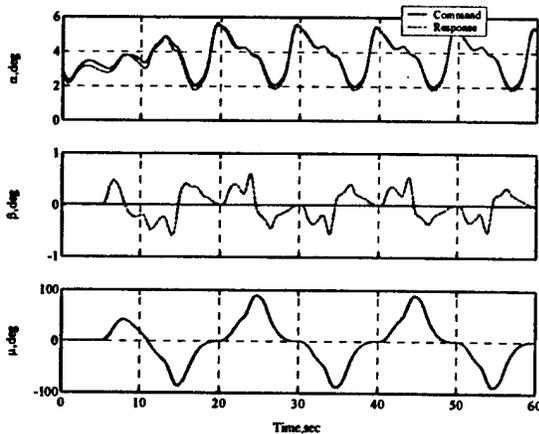**Figure 9: Flight Path Tracking (Unfailed)**

Figure 10: Aerodynamic Angle Tracking

Figure 12 shows the results of the learning on the baseline pitching moment coefficient. The true coefficient is approximately linear in angle of attack; however, the B-Spline function approximators have been initialized at zero. The shaded region represents the range of angle-of-attack for the maneuver.

It can be seen that after this brief maneuver, the baseline moment coefficient has converged *in this region*, but has not changed in regions outside the envelope of the maneuver. This nondestructive, spatially-local property of the learning enables the IC to "remember" what it has learned in one flight condition while being updated in another (in this case, it hasn't yet learned anything about the other flight conditions so they are held at the nominal/initial values).
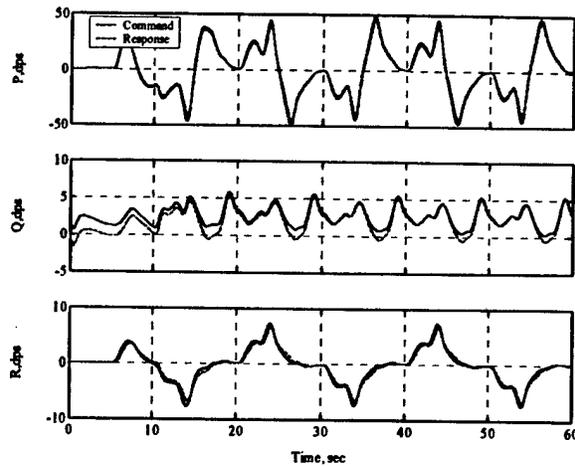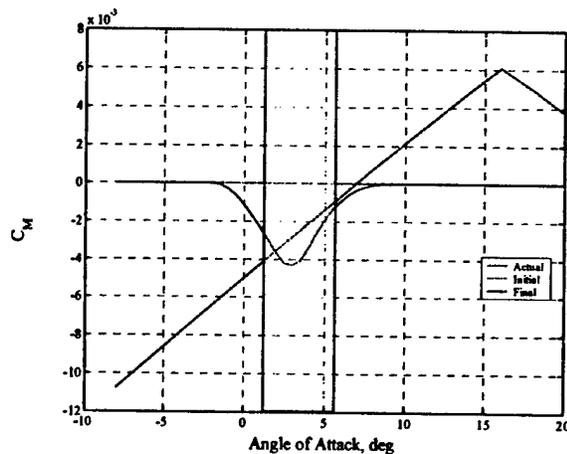

Figure 11: Angular Rate Tracking

Figure 12: Learned Pitching Moment Coefficient

Figure 13 shows an offset-landing maneuver during a failure in which the left outboard flap goes hard over shortly after the initiation of the maneuver (See Figure 14). With a non-reconfigurable controller this failure results in an immediate departure of the vehicle and an obvious inability to complete the task; however, it can be seen that the IC is able to reconfigure and complete the task with approximately 8 feet of downrange error and 4 feet of crossrange error. Figure 15 shows the performance of one of the inner backstepping loops during this maneuver.

American Institute of Aeronautics and Astronautics
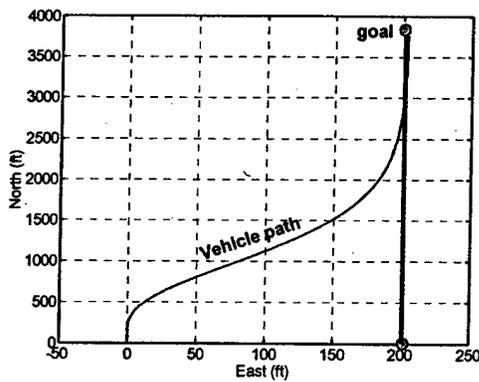
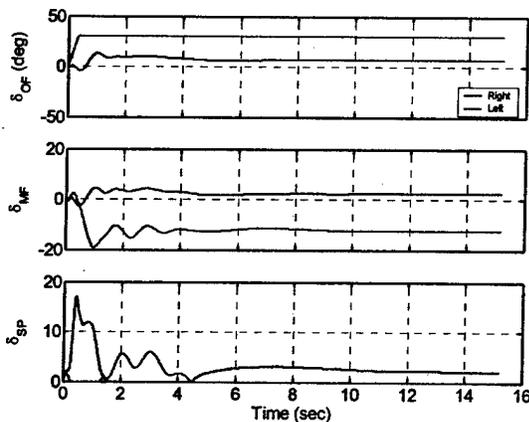Figure 13: Offset Landing with Hardover Outboard Flap

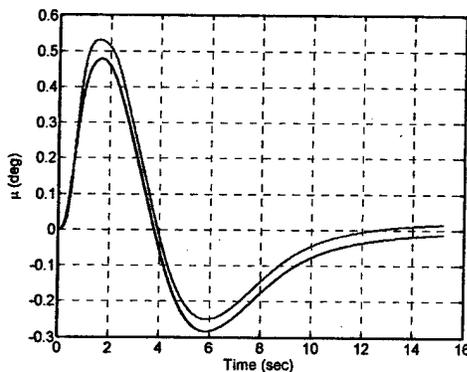Figure 14: Actuator Positions During Offset Landing

Figure 15: Aero Angle Tracking During Offset Landing

In all of the examples above, the IC was able to reconfigure for significant failures and still recover the nominal dynamics of the vehicle and outer-loop reconfiguration was not required. Figure 16 shows a case where trajectory reconfiguration is required. Here the right spoiler locked at 45 deg. The original offset landing path requires approximately 7.5 deg./sec. turn rates to correct for the offset; however, with the

hardover spoiler, the inner loop is only capable of achieving 2.5 deg./sec. and the UCAV cannot line up without going around in a much more gentle turn.
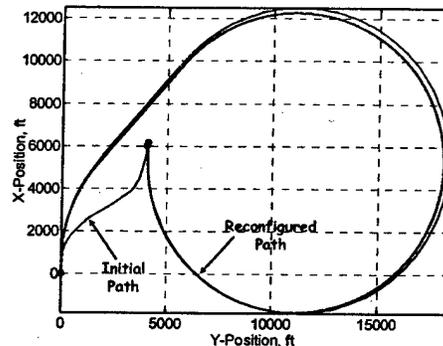

Figure 16: Offset Landing with Hardover Spoiler

SUMMARY AND CONCLUSIONS

Over the past two years the IC program has achieved the following:

- Backstepping flight control that can adapt rapidly to sudden changes and yet can learn a global model of vehicle behavior over time.

- Methods for learning the structure of the underlying aerodynamic model in flight.

- Provably-stable learning rules for the adaptive system with built-in anti-windup algorithms to allow learning even under actuator and state constraints.

- Rapid path planning that accounts for all of the underlying nonlinear vehicle dynamics.

These achievements serve as enabling technology that can provide UCAVs with robustness to unforeseen failures and autopilots capable of aggressive maneuvering when required by certain mission scenarios (weapon delivery, NOE flight, missile evasion, etc.). The learning control approach also allows controllers for new or derivative vehicles can to be developed rapidly in high-fidelity simulations wherein learning would significantly reduce the need for manual tuning of the control law. Finally, the global nature of the learned models as well as the fact that they make physical sense allow any models updated in flight to be replicated on other UCAVs as well as used to update batch simulations.

ACKNOWLEDGMENTS

REFERENCES

[1] Ward, D., J. Monaco, R. Barron, R. Bird, J. Virnig, and T. Landers, Self-Designing Controller: Design, Simulation, and Flight Test Evaluation, Barron Associates, Inc. Final Technical Report for AFOSR, Contract F49620-94-C-0087, Nov. 1996.

[2] Pachter, M., P. Chandler, and M. Mears, "Reconfigurable tracking control with saturation," Journal of Guidance, Control, and Dynamics, Vol. 18, No. 5, Sept. – Oct. 1995, pp. 1016 – 1022.

[3] Rysdyk, R. and A. Calise, "Nonlinear adaptive flight control using neural networks," IEEE Controls Systems Magazine, Vol. 18, No. 6, Dec. 1998.

[4] Steinberg, M., "A comparison of intelligent, adaptive, and nonlinear flight control laws," Proc. AIAA Guidance, Navigation, and Control Conf., Portland, OR, Aug. 1999.

[5] Eberhardt, R., M. Niestroy, G. Tallant, J. Monaco, and D. Ward, Reconfigurable Systems for Tailless Fighter Aircraft, Lockheed Martin Tactical Aircraft Systems and Barron Associates, Inc. Final Technical Rept. for Air Force Research Laboratory, AFRL-VA-WP-TR-1999-3078, Nov. 1999.

[6] Urnes, J., R. Yeager, and J. Stewart, "Flight demonstration of the self-repairing flight control system in a NASA F-15 aircraft," National Aerospace Electronics Conf., Rept. 90CH2881-1, Dayton, OH, May 1990.

[7] Brinker, J. and K. Wise, Reconfigurable Systems for Tailless Fighter Aircraft -RESTORE, Boeing Phantom Works Final Technical Rept. for Air Force Research Laboratory, AFRL-VA-WP-TR-1999-3067, Sept. 1999.

[8] Bateman, A., D. Ward, J. Monaco, and F. Koenig, Tactile Cueing for PIO Avoidance in Manned and Unmanned Aircraft, Barron Associates, Inc. Final Technical Rept. for Air Force Research Laboratory, Contract F41624-98-C-5034, Mar. 1999.

[9] Pachter, M., P. Chander, and M. Mears, "Reconfigurable tracking control with saturation," submitted to AIAA J. Guidance, Control, and Dynamics, 1994.

[10] Monaco, J. and D. Ward, Retrofit Reconfigurable Control System, Barron Associates, Inc. Final Technical Rept. for Naval Air Systems Command, Contract N68335-00-C-0127, May 2000.

[11] Schierman, J., D. Ward, J. Monaco, J. Hull, and M. Ruth, "A Reconfigurable Guidance Approach for Reusable Launch Vehicles," submitted to AIAA Guidance, Navigation, and Control Conf., Aug. 2001.

[12] Ward, D.G., J.F. Monaco, and M. Bodson, "Development and Flight Testing of a Parameter Identification Algorithm for Reconfigurable Control," AIAA J. Guidance, Control, and Dynamics, Vol. 21, No. 6, Nov. - Dec. 1998, pp. 948-956.

[13] Sharma, M. and Calise, A. J., "Adaptive Backstepping Control for a Class of Nonlinear Systems via Multilayered Neural Networks," Proceedings of the American Control Conference, May 2002.

[14] M. Sharma and D. G. Ward, "Flight-Path Angle Control via Neuro-Adaptive Backstepping," Proceedings of the AIAA Guidance, Navigation, and Control Coference, AIAA-2002-4451.

[15] Ward, D.G, "Generalized networks for complex function modeling," Proc. IEEE Systems, Man, & Cybernetics (SMC-94) Conf., Oct. 2-5, 1994.

[16] Monaco, J.F., D.G. Ward, and A.G. Barto, "Automated aircraft recovery via reinforcement learning: Initial experiments," Advances in Neural Information Processing Systems 10, (M.I. Jordan, M.J. Kearns, and S.A. Solla, Eds.), MIT Press, 1998, pp. 1022-1028.

[17] Farrell, J. A., M. Sharma, and M. Polycarpou, "Longitudinal Flight-path Control using On-Line Function Approximation," AIAA Journal of Guidance, Control, and Dynamics (to appear).

[18] Farrell, J. A., M. Sharma, and M. Polycarpou, "On-line approximation based aircraft longitudinal flight control," In Proceedings of the 2003 American Control Conference, 2003.

[19] Sharma, J. A. Farrell, M. Polycarpou, N. D. Richards, and D. G. Ward, "Backstepping Flight Control using On-Line Function Approximation," Proceedings of the AIAA Guidance, Navigation, and Control Coference, 2003.

[20] J. A. Farrell, M. Polycarpou, and M. Sharma, "Adaptive Backstepping with Magnitude, Rate, and Bandwidth Constraints: Aircraft Longitudindal Control," In Proceedings of the 2003 American Control Conference, 2003.

[21] M. Polycarpou, J. A. Farrell, and M. Sharma, "On-Line Approximation Control of Uncertain Nonlinear Systems: Issues with Control Input Saturation," In Proceedings of the 2003 American Control Conference, 2003.

[22] Hastie, T. and R. Tibshirani, Generalized additive models. London: Chapman and Hall, 1990.

[23] Schaal, S., C. Atkeson, and S. Vijayakumar, Scalable techniques from nonparameteric statistics for real-time robot learning. Applied Intelligence, 2000.