--Intelligent Control Approaches for Aircraft Applications

K. KrishnaKumar & Karen Gundy-Burlet
NeuroEngineering Laboratory
NASA Ames Research Center
kkumar@mail.arc.nasa.gov

Abstract

This paper presents an overview of various intelligent control technologies currently being developed and studied under the Intelligent Flight Control (IFC) program at the NASA Ames Research Center. The main objective of the intelligent flight control program is to develop the next generation of flight controllers for the purpose of automatically compensating for a broad spectrum of damaged or malfunctioning aircraft components and to reduce control law development cost and time. The approaches being examined include: (a) direct adaptive dynamic inverse controller and (b) an adaptive critic-based dynamic inverse controller. These approaches can utilize, but do not require, fault detection and isolation information. Piloted simulation studies are performed to examine if the intelligent flight control techniques adequately: 1) Match flying qualities of modern fly-by-wire flight controllers under nominal conditions; 2) Improve performance under failure conditions when sufficient control authority is available; and 3) Achieve consistent handling qualities across the flight envelope and for different aircraft configurations. Results obtained so far demonstrate the potential for improving handling qualities and significantly increasing survivability rates under various simulated failure conditions.

Introduction

In the last 30 years, at least 10 aircraft have experienced major flight control system failures claiming more than 1100 lives [1,2]. The Intelligent Flight Control (IFC) research program began in 1992 to address the need to examine alternate sources of control power to accommodate in-flight control system failures. The major feature of IFC technology is its ability to adapt to unforeseen events through the use of a self-learning neural flight control architecture. These events can include sudden loss of control surfaces, engine thrust, and other causes that may result in the departure of the aircraft from safe flight conditions.

To provide a real-time system capable of compensating for a broad spectrum of failures, NASA researchers began investigating techniques for integrating both flight and propulsion control. The concept was to develop a system capable of utilizing all remaining sources of control power after damage or failures. In order to adapt to varying levels of performance under different control allocation schemes, Propulsion Controlled Aircraft (PCA) technologies [3,4] were incorporated into a neural flight control architecture [5,6]. The resulting Integrated Neural Flight and Propulsion Control System (INFPCS) uses a daisy-chain control allocation technique to ensure that conventional flight control surfaces will be utilized under normal operating conditions. Under damage or failure conditions, the system may allocate flight control surfaces, and incorporate propulsion control, when additional control power is necessary for achieving desired flight control performance.

This research was conducted using a neural flight control architecture that is based upon the augmented model inversion controller, developed by Rysdyk and Calise [7]. This direct adaptive tracking dynamic inverse controller integrates feedback linearization theory with both pre-trained and on-line learning neural networks. Pre-trained neural networks are used to provide estimates of aerodynamic stability and control characteristics required for model inversion. On-line learning neural networks are used to generate command augmentation signals to compensate for errors in the estimates and from the model inversion. The on-line learning neural networks also provide additional potential for adapting to changes in aircraft dynamics due to damage or failure. Reference models are used to filter command inputs in order to specify desired handling qualities. A Lyapunov stability proof guarantees boundedness of the tracking error and network weights [7].

Piloted simulation studies were performed at NASA Ames Research Center on a commercial transport aircraft simulator [6]. Subjects included both NASA test pilots and commercial airline crews. This paper contains a brief overview of the system architecture, and presents simulation results comparing the performance to conventional systems under nominal and simulated failure conditions.

The current research at NASA Ames is focused on improving two aspects of the IFC architecture implemented in the earlier study and to further validate IFC technologies using a C-17 test bed. The improvements and the technologies that enable the improvements are discussed in the body of the paper. In the rest of the paper, first we discuss categorization of intelligent control approaches to give a perspective on where we are in the hierarchy of IFC. In the next section we detail the implementations of the two approaches (past and present) and discuss the results of the past research.

Intelligent Control Approaches

There are two main aspects to intelligent control: (1) the "intelligence" to analyze the changing environment; and (2) the resources to respond to the changing environment. Intelligence connotes the analytical ability to comprehend and react to the changing environment. Resources connote the physical components of the system that are necessary to react to the environment. In this work, we concentrate on the need to harvest and interpret the information from the network of sensors and to apply it for control such that good performance is maintained under any of the following situations:

- Loss of control due to failure
- Aircraft characteristics change due to damage (center of gravity, inertia, etc.)
- Changing operating conditions (altitude, mach, etc.)
- Environmental effects due to wind and turbulence

Intelligent Control application focus on control problems that otherwise cannot be solved, or cannot be solved in a satisfactory way by traditional control techniques alone. Intelligent control as practiced today encompasses many fields from conventional control such as optimal control, robust control, stochastic control, linear control, and nonlinear control, as well as the more recent fuzzy, genetic, and neuro-control technologies. In the next subsection, Intelligent Control is classified based on the idea of self-improvement as the goal toward higher levels of intelligence.

Levels of Intelligent Control

In a general sense, an intelligent controller design can be stated as the following:

given the dynamic system as:

$$X(t+1) = f(X(t),U(t),t)+\eta;$$ where X are the state variables, U is the control vector, and η is an unknown disturbance

a set of goals generated as a function of time as

$$Xg(t+1) = g(Xg(t),X(t),t)$$

a performance measure as:

$$J(t+1)= \Im(M(Xg(t),X(t),U(t),t)); \text{ where } \Im \text{ is an operator (usually summation over T)}$$

and a planning function as

$$P(t+1) = p(X(t),P(t),t,\nu); \text{ where } \nu \text{ is system faults and emergencies}$$

the intelligent controller needs to arrive at a control, U(t), such that the system (in the order of priority)

- is locally stable (includes handling quality of the system)
- follows closely the desired path (closeness defined by the performance measure)
- constantly optimizes long-term and short-term goals
- reacts to changing environments by properly adapting the planning functionality.

Over the past decade, several innovative control architectures utilizing the intelligent control tools have been proposed. We believe that a practical way to accommodate the above needs is to approach the system as having various levels of capabilities for self-improvement. We emphasize here that self-improvement is an important goal of human intelligence. Self-improvement is quantifiable and measurable in various ways. By defining Intelligent Control with various levels of intelligence, the definition is left 'open ended' such that it will not become obsolete, and it will

accommodate easily the innovations that will inevitably come from the contributions of such fields as cognitive science, computer hardware, sensors and actuators, learning theory, and control architectures.

KrishnaKumar [8,9] has proposed a classification scheme based on the ability of the control architecture for self-improvement (see Table 1). The classification scheme divides the control architectures among levels of intelligent control (LIC). For instance, most of the proposed architectures can be divided among level 0, level 1, level 2, and level 3 intelligent control schemes. Based on this classification scheme, several seemingly differing control architectures can be looked at as achieving similar goals.

Table 1. The Levels of Intelligent Control

Level	Self improvement of	Description
0	Tracking Error (TE)	Robust Feedback Control: Error tends to zero.
1	TE + Control Parameters (CP)	Adaptive Control: Robust feedback control with adaptive control parameters (error tends to zero for non-nominal operations; feedback control is self improving).
2	TE + CP + Performance Measure (PM)	Optimal Control: Robust, adaptive feedback control that minimizes or maximizes a utility function over time.
3	TE+CP+PM+ Planning Function	Planning Control: Level 2 + the ability to plan ahead of time for uncertain situations, simulate, and model uncertainties.

Level 0 Intelligent Control -- A Robust Controller: Self-improvement of Tracking Error (TE) is an important goal of many control techniques. To achieve this, one designs robust feedback controllers with constant gains that improve the error as time goes to infinity. We consider this as "Level 0 Intelligent Control".

Level 1 Intelligent Control -- An Adaptive Controller: Self-improvement of control parameters towards the goal of achieving better tracking error or some error oriented goal, is the next level in intelligent control. We consider this as "Level 1 intelligent Control". This level is essentially a robust feedback controller with adaptive parameters that helps the error tend to zero for non-nominal operations and feedback controller is self improving.

Level 2 Intelligent Control -- An Optimal Controller: Self-improvement of an estimate of the performance error (or some measure of performance over time) towards the goal of minimization or maximization of an utility function over time (error tends to zero and a measure of performance is optimized) is the next level in intelligent control. We consider this as "Level 2 Intelligent Control". This level is essentially a robust feedback controller with adaptive parameters that helps the error tend to zero for non-nominal operations, feedback controller is self improving, and a measure of performance that is self-improving is optimized over time.

Level 3 Intelligent Control -- A Planning Controller: In addition to level 2 capabilities, Level 3 Intelligent Control includes self- improvement of planning functions. Planning functions include contingency planning, planning for emergencies, planning for faults, etc. These planning functions could be static for Level 2 but needs to be self-improving for Level 3.

Figure 1 presents the implementation of a Level 2 Intelligent Control on a C-17 test bed. The levels as outlined earlier are labeled in the figure. It should be noted that Level 0 is non-adaptive whereas Level 1 is adaptive. Level 1 is non-optimal whereas Level 2 is optimal.

Figure 2 presents a Level 3 Intelligent Control architecture for intelligent aircraft autonomous maneuvering. In this application of IFC, a Level 3 capability for adapting the path (maneuvering) is included.
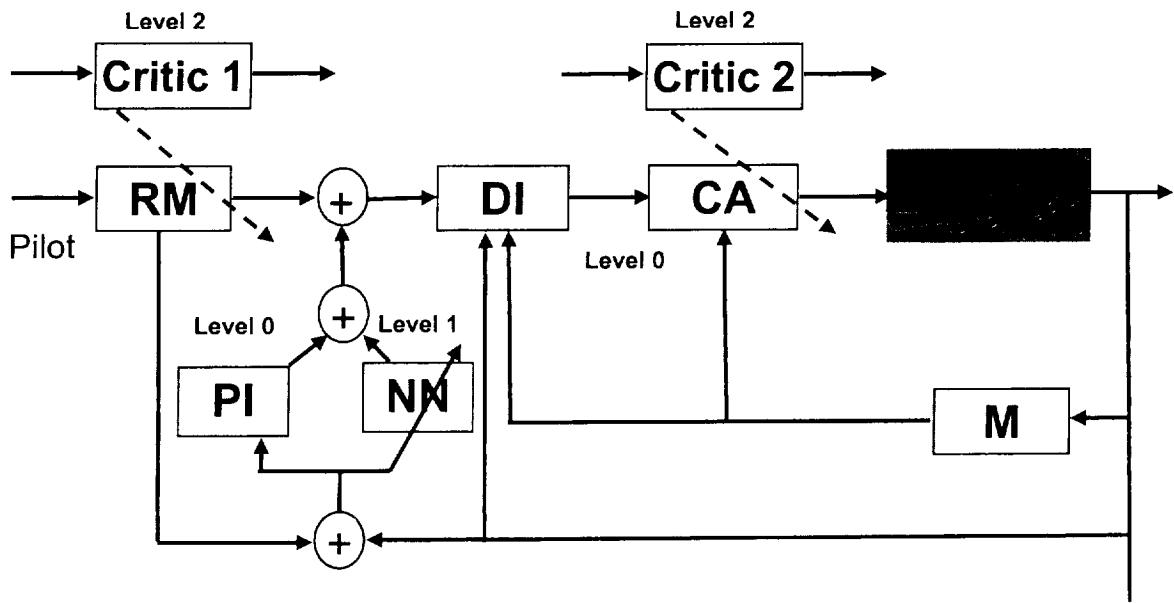
RM: Reference Models: The pilot commands roll rate and aerodynamic normal and lateral accelerations through stick and rudder pedal inputs. These commands are then transformed into body-axis rate commands, which also include turn coordination, level turn compensation, and yaw-dampening terms. First-order reference models are used to filter these commands in order to shape desired handing qualities.

PI Error Controller: Errors in roll rate, pitch rate, and yaw rate responses can be caused by inaccuracies in aerodynamic estimates and model inversion. Unidentified damage or failures can also introduce additional errors. In order to achieve a rate-command-attitude-hold (RCAH) system, a proportional-integral (PI) error controller is used to correct for errors detected from roll rate, pitch rate, and yaw rate (p, q, r) feedback.

NN: On-Line Learning Neural Networks: The on-line learning neural networks work in conjunction with the error controller. By recognizing patterns in the behavior of the error, the neural networks can learn to remove biases through control augmentation commands. These commands prevent the integrators from having to windup to remove error biases. By allowing integrators to operate at nominal levels, the neural networks enable the controller to provide consistent handling qualities.

DI: Dynamic Inversion: Dynamic inversion is based upon feedback linearization theory. No gain-scheduling is required, since gains are functions of aerodynamic stability and control derivative estimates and sensor feedback. To perform the model inversion, acceleration commands are used to replace the actual accelerations in the quasi-linear model. The model is then inverted to solve for the necessary control surface commands

CA: Control Allocation: A daisy-chain control allocation technique is used to ensure that conventional flight control surfaces will be utilized under normal operating conditions. Unconventional flight control surface allocations are only utilized when the primary flight control surface commands exceed the known limits of deflection. For example, in the longitudinal axis pitch rate control is normally provided through symmetric elevator deflections. If this command should saturate, then the remaining portion of the command is applied to symmetric ailerons. If the symmetric aileron command saturates, then the remaining portion of that command is applied to symmetric thrust. The symmetric aileron command is limited, by the differential aileron command, so that secondary pitch control does not interfere with primary roll control.

M: Pre-Trained Neural Network Model: A Levenberg-Marquardt (LM) multi-layer perceptron [10] is used to provide dynamic estimates for model inversion. The LM network is pre-trained with stability and control derivative data generated by a Rapid Aircraft Modeler, and vortex-lattice code [11]. This block can be replaced by other on-line derivative (parameter) estimation techniques.

Critic 1 and 2: Adaptive Critics are utilized to optimize the control allocation scheme and to shape the reference model dynamics in the event of a failure.

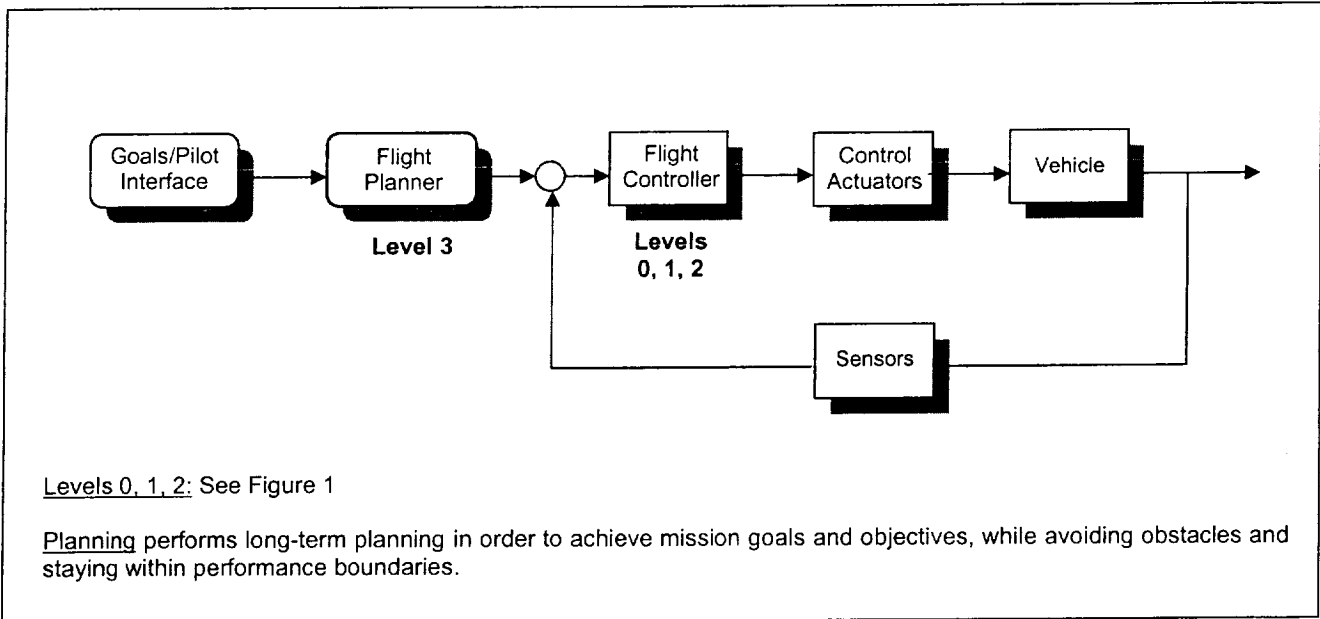Figure 1. Level 2 Intelligent Flight Control System

Levels 0, 1, 2: See Figure 1

Planning performs long-term planning in order to achieve mission goals and objectives, while avoiding obstacles and staying within performance boundaries.

Figure 2. Level 3 IFC Architecture

Intelligent Aircraft Control Research at NASA Ames

An aircraft system is a non-linear distributed system with complex interactions between the pilot, the aircraft, and the environment. Complexity here implies that a simple failure can lead to a catastrophic aircraft response. Intelligent control provides a means to adapt the control law to changes in the system and to gracefully degrade the performance in instances of catastrophic failures. Due to this obvious benefit along with other benefits that include low developmental costs, a great deal of interest has been shown towards intelligent control applications to aircraft and spacecraft operations.

Past Research

This section describes an integrated neural flight and propulsion control system that was extensively studied using piloted simulations with simulated failure conditions. Under normal operating conditions, the system utilizes conventional flight control surfaces. Neural networks are used to provide consistent handling qualities across flight conditions and for different aircraft configurations. Under damage or failure conditions, the system may utilize unconventional flight control surface allocations, along with integrated propulsion control, when additional control power is necessary for achieving desired flight control performance. Piloted simulation studies were performed on a commercial transport aircraft simulator. Subjects included both NASA test pilots and commercial airline crews. Results demonstrate the potential for improving handing qualities and significantly increasing survivability rates under various simulated failure conditions [6].

System Architecture

The neural flight control architecture is based upon the augmented model inversion controller, developed by Rysdyk and Calise [7]. This direct adaptive tracking dynamic inverse controller, shown in Figure 1 (without the Level 2 components), integrates feedback linearization theory with both pre-trained and on-line learning neural networks. Pre-trained neural networks are used to provide estimates of aerodynamic stability and control characteristics required for model inversion. On-line learning neural networks are used to generate command augmentation signals to compensate for errors in the estimates and from the model inversion. The on-line learning neural networks also provide additional potential for adapting to changes in aircraft dynamics due to damage or failure. Reference models are used to filter command inputs in order to specify desired handling qualities. A Lyapunov stability proof guarantees boundedness of the tracking error and network weights.

A daisy-chain control allocation technique is used to ensure that conventional flight control surfaces will be utilized under normal operating conditions. Under damage or failure conditions, the system may utilize unconventional flight control surface allocations, along with integrated propulsion control, when additional control power is necessary for achieving desired flight control performance. Of significant importance here is the fact that this system can operate without emergency or backup flight control mode operations. An additional advantage is that this system can utilize, but does not require, fault detection and isolation information or explicit parameter identification.

5

Simulator Description

The INFPCS system was evaluated on the Advanced Concepts Flight Simulator (ACFS) at NASA Ames Research Center. The simulator is equipped with a six degree-of-freedom motion system, programmable flight displays, digital sound and aural cueing system, and a 180-degree field of view visual system [12].

The simulated aircraft is representative of a mid-size two-engine jet transport with general characteristics of a wide-body, T-tail, low wing airplane with twin turbofan engines located under the wings. The aircraft is based on a Lockheed Georgia Company Commercial transport concept developed in 1983. The physical dimensions are similar to a Boeing 757 aircraft, with flight characteristics representative of a mid-sized jet transport. This particular transport aircraft, designed as a platform for testing advanced concepts, is equipped with active flight controls representative of an advanced fly-by-wire control system.

The Dryden turbulence model provides turbulence RMS and bandwidth values which are representative of those specified in Military Specifications Mil-Spec-8785 D of April 1989. The Earth atmosphere is based on a 1976 standard atmosphere model.

Test Description

The INFPCS system was evaluated by a total of 12 pilots, 6 NASA test pilots and 6 commercial airline pilots (3 commercial airline crews). Test pilots performed select maneuvers and approach and landing scenarios under nominal and simulated failure conditions. Handling characteristics were compared to those of a standard "Open-Loop" cable-driven controller, and a "Conventional" rate-command-attitude-hold (RCAH) fly-by-wire (FBW) controller. Commercial pilots performed full-mission flights, from takeoff to landing, in order to evaluate the effectiveness under select operational scenarios. The evaluation criterion was based upon performance measurements, Cooper-Harper Ratings (CHR), and pilot comments.

Failure conditions consisted of "frozen" flight control surfaces at neutral or offset positions, shifts in the center of gravity, and special controller failures to measure levels of adaptation. The failures used in this evaluation were selected in order to investigate specific control issues, and to represent realistic scenarios encountered in aircraft accidents and incidents involving the loss of some or all flight control surfaces.

Figure 3 presents a comparison of INFPCS with and without Level 1 adaptation for a normal and a failed condition. It is clearly seen that the INFPCS with Level 1 adaptation is significantly better for the failed condition. Similar data were obtained for several other failure conditions. These test cases included (see Reference 6 for more details).

- Approach and Landing Tests
- Tail Failure Tests
- Control Allocation Transition Tests
- Control Saturation Tests
- Operational Scenario Tests
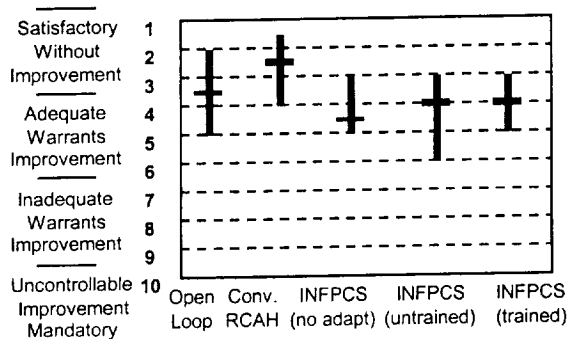- Dead-Band Adaptation Tests
- Out-of-trim Tests

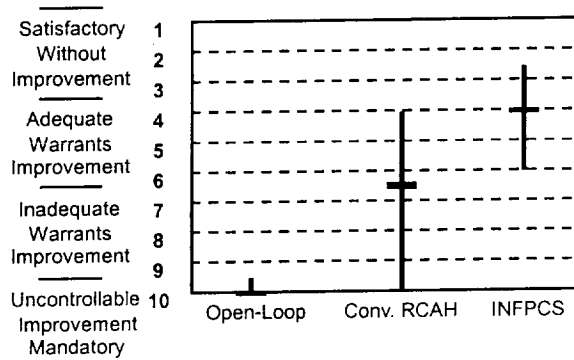Figure 3a. Lateral Cooper-Harper Ratings for Nominal condition

Figure 3b. Longitudinal Cooper-Harper Ratings (Runaway Stabilizer Trim)

During the Controller Failure tests, the effect of adaptation became more apparent. Figure 4 displays roll rate and pitch rate errors respectively, during a sample test case when performing sequential sets of maneuvers. The first set of maneuvers, performed without adaptation, contains the largest errors. Adaptation is then used during the second set of maneuvers, starting with untrained neural networks. The magnitude of the error continues to reduce as the maneuver is repeated a third time, representing trained neural networks. Pilots commented that there was a "big difference" in performance, as lateral ratings improved from 6.7 to 3.8 CHR, and longitudinal ratings improved from 6.6 to 4.2 CHR.
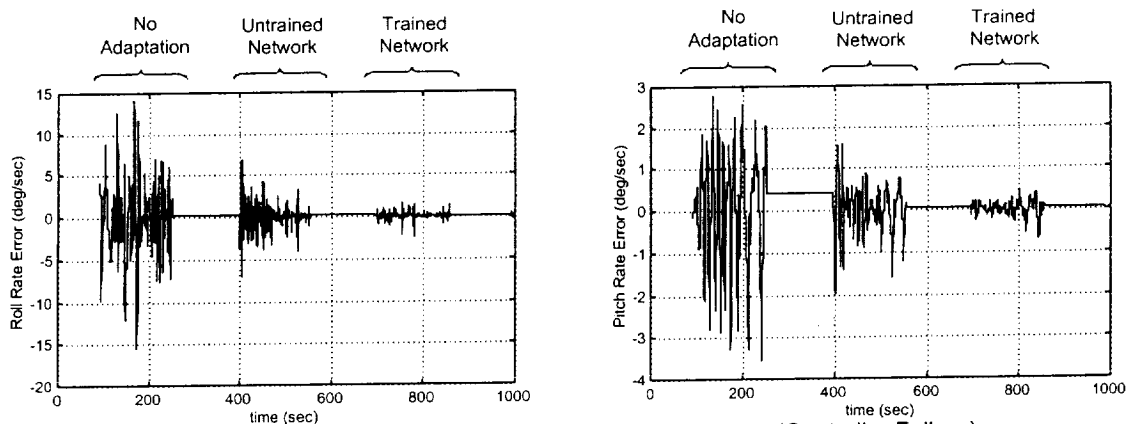

Figure 4. INFPCS Roll-Rate and Pitch-Rate Errors (Controller Failure)

Current Research

The research highlighted above was successful in demonstrating the benefits of intelligent adaptive control. One of the improvements suggested by the pilots included a faster control allocation technique. The Level 1 architecture did not support an optimal allocation scheme. The scheme used an ad-hoc approach as outlined in Table 2 below.

Table 2. Daisy-Chain Control Allocation

	Lat.	Dir.	Long.
Primary Allocation	δ_{dail}	δ_{drud}	δ_{elev}
Secondary Allocation	yaw-based roll control	δ_{dthr}	δ_{ail}
Tertiary Allocation	- - -	- - -	δ_{thr}

The above approach cannot be easily scaled to aircraft with several redundant control surfaces, such as the C-17. It is highly desirable to "optimally" allocate the control surfaces in the event of failure and it is desirable to arrive at a technique that is independent of the aircraft. In addition, in the event of a severe degradation in performance, pilot handling qualities as dictated by the reference model cannot be maintained. Once again, it will be desirable to

"optimally" modify the dynamics of the reference model to suit the situation in hand. Towards these goals an adaptive critic augmentation that enables a Level 2 IFC is being developed (see Figure 1).

Adaptive Critic Technology

Adaptive critic designs have been defined as designs that attempt to approximate dynamic programming. Paul Werbos [22], who is one of the foremost proponents of adaptive critics, has stated that "adaptive critics are the only design approach that show serious promise of duplicating critical aspects of human intelligence: the ability to cope with large number of variables in parallel, in real time, in a noisy nonlinear environment".

The three inputs required for designing an adaptive critic design are:
- The cost function.
- A parameterized representation of the critic.
- A method for adapting the parameters of the critic.

The choice of the cost function that the critic approximates comes from the problem at hand. The cost could be distributed over the entire length of time or be defined at the end of the process. Typical examples of the two types are minimizing the fuel spent for a certain flight mission or intercepting a projectile where the utility depends only on the final error.

Parameterization of the critic is achieved by the use of the neural network formulation. Neural networks provide nonlinearity and ease of implementation.

The choice of the method for training the critic depends on the problem at hand. Some of the methods proposed for training the adaptive critics based on dynamic programming are the heuristic dynamic programming (HDP) approach, the dual heuristic programming (DHP) approach, the global dual heuristic dynamic (GDHP) programming approach. References 21 and 22 provide a more detailed discussion on the subject. Adaptive critic methods have been successfully applied to aerospace problems in references 13 and 14.

Test Article

The adaptive critic architecture and its control reallocation characteristics will initially be assessed in piloted full-motion simulations using a C-17 model currently being integrated with the Advanced Concepts Flight Simulator (ACFS). The high-fidelity C-17 model has substantially more control effectors than the baseline generic transport ACFS model, providing a chance to utilize multiple sources of control redundancy in each axis. After the software is validated in the simulator, it will be integrated with a C-17 airframe and flight tested at Dryden Flight Research Center in conjunction with the USAF and Boeing.

Conclusions

In the event of damage or failure, an aircraft's response can change to the point of essentially becoming a new aircraft. Since conventional transport aircraft have limited additional control power available, alternate control allocation schemes may be necessary. The Intelligent Flight Control research demonstrates the effectiveness of using neural networks to not only adapt to changes in aircraft dynamics, but also to adapt to different control allocation schemes.

The Intelligent Flight Control technology potentially can be developed and installed in commercial and military aircraft to provide information and assistance to the flight crew in situations of major damage or component failure of its control systems and has been flown in initial stages on the X-36 and F-15 ACTIVE. While these systems will increase the safety and survivability of the vehicle, they must also be produced in shorter time spans, with lower costs, and ultimately increase profitability in order for the US industry recipients to remain competitive. The use of neural networks also leads to the concept of Intelligent Vehicle Systems that takes advantage of advances in information technology that enable aircraft on-board flight systems to detect characteristic changes, predict future systemic events and independently regain its original state or achieve an optimal state through reconfiguration or re-planning.

References

1. Burcham, Frank W., Jr., Trindel A. Maine, C. Gordon Fullerton, and Lannie Dean Webb, Development and Flight Evaluation of an Emergency Digital Flight Control System Using Only Engine Thrust on an F-15 Airplane, NASA TP-3627, Sept. 1996.
2. National Transportation Safety Board, Aircraft Accident Report, PB90-910406, NTSB/ARR-90/06, United Airlines Flight 232, McDonnell Douglas DC-10, Sioux Gateway Airport, Sioux City, Iowa, July 1989.
3. Burcham, Frank W., Jr., John J. Burken, Trindel A. Maine, and C. Gordon Fullerton, Development and Flight Test of an Emergency Flight Control System Using Only Engine Thrust on an MD-11 Transport Airplane, NASA TP-97-206217, Oct. 1997.
4. Kaneshige, John, John Bull, Edward Kudzia, and Frank W. Burcham, Propulsion Control with Flight Director Guidance as an Emergency Flight Control System, AIAA 99-3962, August 1999.
5. Kaneshige, John, John Bull, and Joseph J. Totah, Generic Neural Flight Control and Autopilot System, AIAA 2000-4281, August 2000.
6. Kaneshige, John and Gundy-Burlet, Karen, Integrated Neural Flight and Propulsion Control System, AIAA-2001-4386, August 2001.
7. Rysdyk, Rolf T., and Anthony J. Calise, Fault Tolerant Flight Control via Adaptive Neural Network Augmentation, AIAA 98-4483, August 1998.
8. Krishnakumar, K., Levels of Intelligent Control: A Tutorial, New Orleans, LA, 1997.
9. Krishnakumar, K., Kulkarni, N., "Inverse Adaptive Neuro-Control for the control of a turbofan engine", Proceedings of AIAA conference on Guidance, Navigation and Control, Portland, OR, 1999.
10. Norgaard, M., Jorgensen, C., and Ross, J., Neural Network Prediction of New Aircraft Design Coefficients, NASA TM-112197, May 1997.
11. Totah, Joseph J., David J. Kinney, John T. Kaneshige, and Shane Agabon, An Integrated Vehicle Modeling Environment, AIAA 99-4106, August 1999.
12. Blake, Matthew W., The NASA Advanced Concepts Flight Simulator: A Unique Transport Aircraft Research Environment, AIAA-96-3518-CP.
13. Balakrishnan, S.N., Biega, V., Adaptive Critic Based Neural Networks for Aircraft Optimal Control, Journal of Guidance, Control and Dynamics, Vol. 19, No. 4, July-August, 1996.
14. Kulkarni, N, Krishnakumar, K, Intelligent Engine Control Using an Adaptive Critic, IEEE Transactions on Control System Technology (To Appear), 2002.
15. Bellman, R., Dynamic Programming, Princeton University Press, 1957.
16. De Silva, C.W., Shoureshi, R., (eds), Intelligent Control Systems 1993, presented at the ASME Winter Annual Meeting, New Orleans, LA, 1993.
17. Harris, C. J., (ed), Advances in Intelligent Control, Taylor and Francis, 1994.
18. Lightbody, G., Irwin, G., Direct Neural Model Reference Adaptive Control, IEEE Proceedings on Control Theory Applications, vol. 142, no.1, 1995.
19. Moore, K. L., Naidu, D. S., Siddaiah, M., A real time adaptive linear quadratic regulator using neural networks, ECC '93, 1993.
20. Nguyen, D. H., Widrow, B., "Neural Networks for self-learning control systems",
21. Prokhorov, D., Wunsch, D. C., "Adaptive Critic Designs", IEEE transactions on Neural Networks, Vol. 8, No. 5, September, 1997.
22. Werbos, P. J., Neuro Control and Supervised Learning: An Overview and Evaluation", Handbook of Intelligent Control, White D. A., Sofge, D. A., (eds), Van Nostrand Reinhold, NY, 1992.
23. White, D. A., Sofge, D. A., (eds), Handbook of Intelligent Control, Van Nostrand Reinhold, NY, 1992.