Flight Results of the NF-15B Intelligent Flight Control System (IFCS) Aircraft with Adaptation to a Longitudinally Destabilized Plant

John T. Bosworth[*]
NASA Dryden Flight Research Center, Edwards, California, 93523

Adaptive flight control systems have the potential to be resilient to extreme changes in airplane behavior. Extreme changes could be a result of a system failure or of damage to the airplane. The goal for the adaptive system is to provide an increase in survivability in the event that these extreme changes occur. A direct adaptive neural-network-based flight control system was developed for the National Aeronautics and Space Administration NF-15B Intelligent Flight Control System airplane. The adaptive element was incorporated into a dynamic inversion controller with explicit reference model-following. As a test the system was subjected to an abrupt change in plant stability simulating a destabilizing failure. Flight evaluations were performed with and without neural network adaptation. The results of these flight tests are presented. Comparison with simulation predictions and analysis of the performance of the adaptation system are discussed. The performance of the adaptation system is assessed in terms of its ability to stabilize the vehicle and reestablish good onboard reference model-following. Flight evaluation with the simulated destabilizing failure and adaptation engaged showed improvement in the vehicle stability margins. The convergent properties of this initial system warrant additional improvement since continued maneuvering caused continued adaptation change. Compared to the non-adaptive system the adaptive system provided better closed-loop behavior with improved matching of the onboard reference model. A detailed discussion of the flight results is presented.

Nomenclature

A	=	state derivative matrix
B	=	control derivative matrix
B_a	=	neural network basis function vector
CM	=	canard multiplier destabilizing gain, deg canard/deg angle of attack
dt	=	change in time, s
G	=	vector of adaptation gains
FFT	=	fast Fourier transform
IFCS	=	Intelligent Flight Control System
K_{lat}	=	lateral stick command gain, deg/s/in
K_{lon}	=	longitudinal stick command gain, deg/s/in
L	=	vector of neural network error-modification damping terms
L_α	=	apparent lift curve slope, rad/s
MUAD	=	maximum unnoticeable added dynamics
NASA	=	National Aeronautics and Space Administration
NN	=	neural network
PID	=	proportional, integral, and derivative
q	=	pitch rate, deg/s
Sym can	=	symmetric canard deflection, deg

---

[*] Aerospace Engineer, Controls and Dynamics Branch / RC, P.O. Box 273 / MS 4840D, AIAA member.

Sym stab	=	symmetric stabilator deflection, deg
s	=	Laplace transformation variable
U	=	forward path command, deg/s^2
u	=	actuator command vector
W	=	neural network weight vector
WQ1	=	pitch axis forward path command neural network weight
WQ2	=	pitch axis proportional error neural network weight
WQ3	=	pitch axis integral error neural network weight
WQ4	=	pitch axis crossfeed from the roll axis neural network weight
WQ5	=	pitch axis crossfeed from the yaw axis neural network weight
WQ6	=	pitch axis bias term neural network weight
WQ7	=	pitch axis angle of attack weight
$\dot{W}$	=	neural network weight vector derivative
x	=	state vector
$\dot{x}$	=	state vector derivative
α	=	angle of attack, deg
δ_{stk}	=	pilot stick position, in.
ζ_{sp}	=	short-period damping
p	=	roll rate, deg/s
τ_r	=	roll mode time constant, s
ω_{sp}	=	short-period natural frequency, rad/s

Subscripts

ad	=	adaptation
c	=	command
dd	=	feedback
err	=	error compensation
lat	=	lateral
lon	=	longitudinal
ref	=	reference

I. Introduction

Adaptive flight control systems have the potential to be resilient to extreme changes in airplane behavior. Extreme changes could be a result of system failure, damage to the airplane, or a configuration change. One example of a configuration change would be a change in wing-sweep angle. In 1984 an airplane was lost when the wing-sweep angle was changed without a corresponding transfer of fuel;[1] the resulting change in stability was beyond the robustness capability of the control system, causing the airplane to depart from controlled flight. An adaptive flight control system has the potential to adjust to such a sudden change in vehicle stability, to maintain controlled flight, and to provide handling qualities that allow for safe recovery of the airplane. The goal for the adaptive system is to provide an increase in survivability in the event that extreme changes in airplane behavior are encountered.

The National Aeronautics and Space Administration (NASA) Dryden Flight Research Center (DFRC) (Edwards, California) NF-15B Intelligent Flight Control System (IFCS) airplane was equipped with a neural-network-based adaptive flight control system. The system was implemented as a direct adaptive system; in this scheme, the adaptation is driven by the presence of feedback errors and no explicit model identification or prior knowledge of the cause of the errors is required. The adaptive element was incorporated into a dynamic inversion controller with explicit reference model-following. When the overall system is working correctly the feedback errors are minimized and the closed-loop system dynamics match the desired onboard reference model characteristics.

An abrupt change in apparent stability was developed as a test for the IFCS. The NASA NF-15B IFCS airplane, tail number 837 (NASA 837), is shown in Fig. 1. The symmetric canard surfaces on this airplane are programmed to respond to the angle of attack and are scheduled so that the lift carried by the canards provides an apparent plant that is neutrally stable. By changing the gain on the angle-of-attack-to-canard feedback the apparent stability of the system can be altered.

Figure 1. The NASA NF-15B IFCS research airplane, tail number 837.

A range of angle of attack feedback gains were programmed as in-flight simulated destabilizing failures. The adaptation system was engaged and the airplane characteristics were evaluated. This report presents the flight-test results. The performance of the adaptation system is assessed in terms of its ability to stabilize the vehicle and reestablish good onboard reference model-following. Discussions of the convergent characteristics and applicability of time-invariant analysis to a time-varying system are presented.

II.   Background

A. Description of the Test Airplane

The test airplane, NASA 837 (Fig. 1), is a highly-modified preproduction F-15B airplane[2] and is not representative of production F-15 aircraft. Modifications to the airplane include two canards mounted on the upper inlet area forward of the wing. The canards are modified F-18 horizontal tail surfaces; their position in flight is scheduled to respond with angle of attack. An additional modification to the airplane is the incorporation of the F100-PW-229 Pratt & Whitney (West Palm Beach, Florida) engines equipped with axisymmetric thrust-vectoring nozzles. The thrust-vectoring capability was not used during the flight phase covered by this report. The airplane is controlled by a quadruplex, digital, fly-by-wire flight control system. All mechanical linkages between the control stick, rudder pedals, and control surfaces have been removed from the airplane. A processor was added to each of the flight computer channels to execute experimental control laws in a redundant safety-critical mode. A more capable single-string research processor was added to execute the adaptation software. Limits and safety monitors exist in the flight-critical control law processor to ensure safe operation.

B. Description of the Intelligent Flight Control System

The F-15 IFCS is a direct adaptive neural-network-based flight control system. The approach is based on the augmented model inversion architecture developed by Calise, et al.,[3] and Rysdyk, et al.[4] The general control scheme consists of a dynamic inversion controller with explicit model-following.[5] An adaptive component is added to accommodate large errors that are outside the normal robustness range of the dynamic inversion controller. The neural network providing the online adaptation within the F-15 IFCS scheme is known as a sigma-pi neural network and was developed at the NASA Ames Research Center, Moffett Field, California.[6] A more detailed description of the implementation of the neural network for flight-testing can be found in Ref. 7. The main components of the controller are illustrated in Fig. 2.
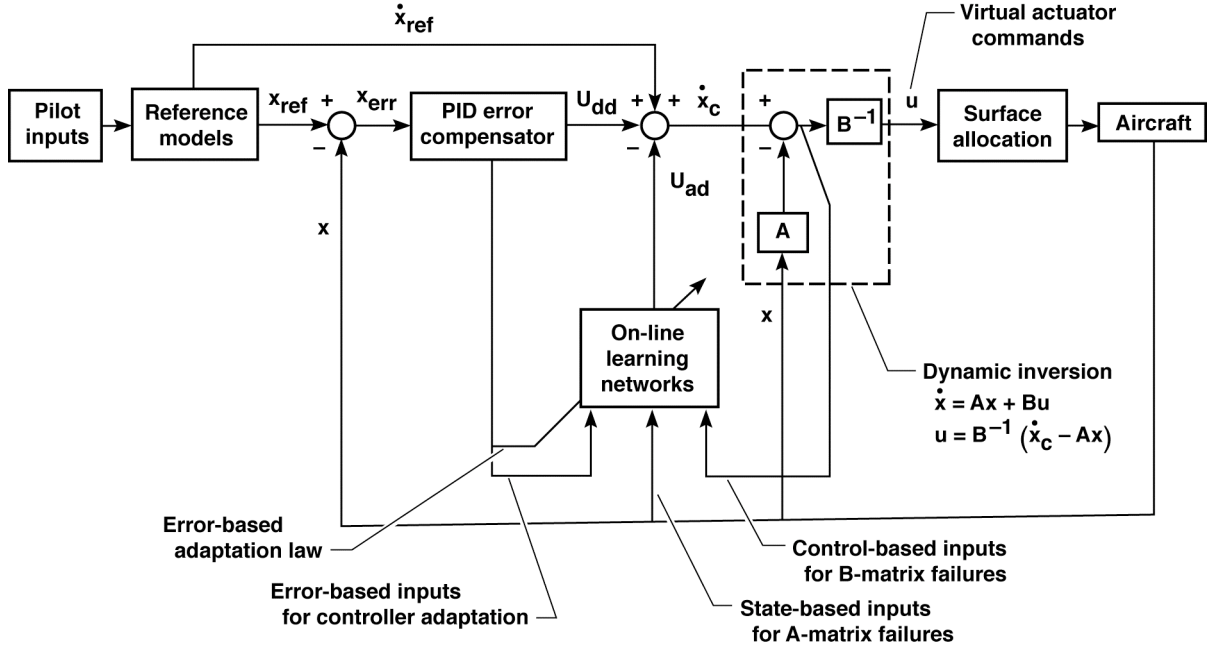
Figure 2. The intelligent flight control system control architecture.

1. Dynamic Inversion

The dynamic inversion portion of the flight control system provides a consistent controlled response for angular acceleration commands ($\dot{x}_c$ in Fig. 2) to the airplane. The primary feedback signals are pitch rate, roll rate, angle of attack, and lateral acceleration. Bank and pitch attitude angle measurements are used for gravity compensation and a sideslip rate (betadot) estimation. A simplified onboard aerodynamic model (contained in the A and B matrices in Fig. 2) is incorporated into the control algorithm. For a given commanded acceleration, simplified equations of motion are used to calculate the required control surface commands. A proportional, integral, and derivative (PID) error feedback compensator is wrapped around the dynamic inversion controller to account for the simplifications used in the dynamic model and to reject disturbances.

2. Explicit Model–Following

An explicit model-following scheme is used to achieve desired handling qualities. Reference models are defined with desired frequency and damping characteristics. The control system attempts to force the response of the airplane to match the reference model. The pitch axis desired reference model is a second-order system and is represented in Eq. (1):

$$\frac{q_{ref}}{\delta_{stk_{lon}}} = \frac{K_{lon}\omega_{sp}^2(s+L_\alpha)}{s^2 + 2\zeta_{sp}\omega_{sp}s + \omega_{sp}^2} \tag{1}$$

The desired short-period natural frequency (ω_{sp}), damping (ζ_{sp}), and apparent lift curve slope (L_α) are selected to achieve Level 1 flying qualities.[8] The longitudinal stick command gain (K_{lon}) is chosen to provide an appropriate stick force per unit normal load factor (g).

The roll axis reference model is first-order and is shown in Eq. (2):

$$\frac{p_{ref}}{\delta_{stk_{lat}}} = \frac{K_{lat}}{\tau_r s + 1} \tag{2}$$

American Institute of Aeronautics and Astronautics

The lateral stick command gain (K_{lat}) is chosen to provide the appropriate amount of roll rate for the given flight conditions, and the roll mode time constant (τ_r) is selected to adjust how fast the roll rate is achieved. Values for these quantities were selected to achieve Level 1 flying qualities.[8]

3. Simulated Destabilizing Failure

A simulated destabilizing failure mode was developed as a test for the adaptive system. This was accomplished by changing the feedback gain from the angle of attack to the symmetric canard command. An extremely simplified block diagram of the F-15 IFCS system is shown in Fig. 3. The sole input to the symmetric canard command is derived from the angle of attack measurement feedback. Pitch axis control augmentation from the flight control system is provided through commands to the symmetric stabilators. Thus the control system "sees" an apparent plant that consists of the bare airframe with the angle-of-attack-to-canard feedback. When the angle of attack feedback gain is changed the stability of this apparent plant can be altered.
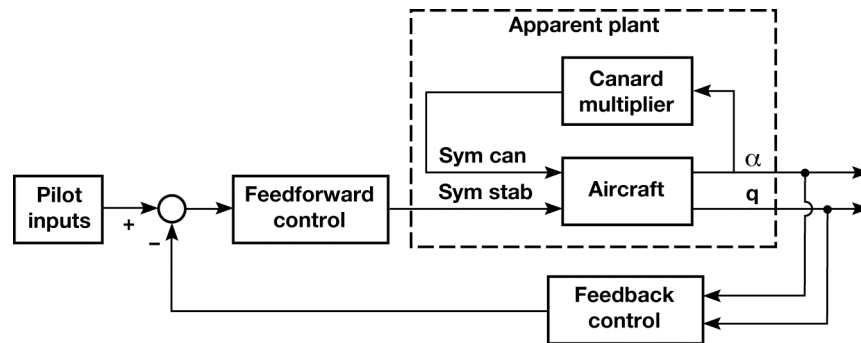

Figure 3. Simplified concept of the pitch axis control loops.

A non-adaptive control system is designed for nominal plant dynamic characteristics. As the apparent plant stability is decreased, eventually the performance of the fixed gain system will decrease in terms of its ability to achieve good reference model-following. An adaptive system has the ability to adjust the feedback gains to the symmetric stabilator to regain stability and better model-following performance.

A selectable canard multiplier gain value was inserted in the angle of attack feedback path. The value was varied from 1.0 (no change) to a value of -1.75. At a value of -1.75 the flight-measured response indicated that, without adaptation, the airplane was essentially neutrally stable with very little margin in the pitch axis.

4. Adaptive Neural Network

The goal of the neural network system is to accommodate large errors that are not anticipated in the nominal control law design. A well-designed flight control system is to some extent robust and maintains stability and controllability for a fairly large range of uncertainty or changes in airplane behavior. As the changes become more extreme the stability and controllability degrade. An adaptive system has the ability to readjust the controller to re-achieve desired stability and controllability or regain robustness about the new point. In the case of a failure (for example, degraded aircraft dynamics or reduced control surface effectiveness) larger-than-expected errors will develop. The adaptive neural networks operate in conjunction with the measured response error of the control system. Weights (gains) on the neural network parameters are dynamically adjusted until the error is reduced. The weights act as adjustments to the proportional, integral, and forward-loop gains (note that the weights do not directly adjust the gains, however, a parallel path is mathematically equivalent to changing these gains). Weights can also provide a control bias, a new feedback to the system, or new crossfeed paths between the control axes. When optimal weights are achieved, the feedback error is minimized and the system achieves better reference model-following and, presumably, better handling qualities.

The F-15 IFCS implementation incorporated a simplified sigma-pi neural network. The cross-multiplication of the input terms was eliminated to limit the number of network weights. This helps by eliminating nonlinear terms and simplifying the sensitivity analysis of each input, as well as increasing learning rates while reducing learning transients. The simplified neural network is essentially the sum of the products of the inputs and a set of weights. Three separate neural networks provided adjustment (U_{ad} in Fig. 2) to the roll, pitch, and yaw forward path

American Institute of Aeronautics and Astronautics

commands. Figure 4 shows the pitch axis neural network. A more complete discussion of the roll and yaw neural networks can be found in Refs. 7 and 9.
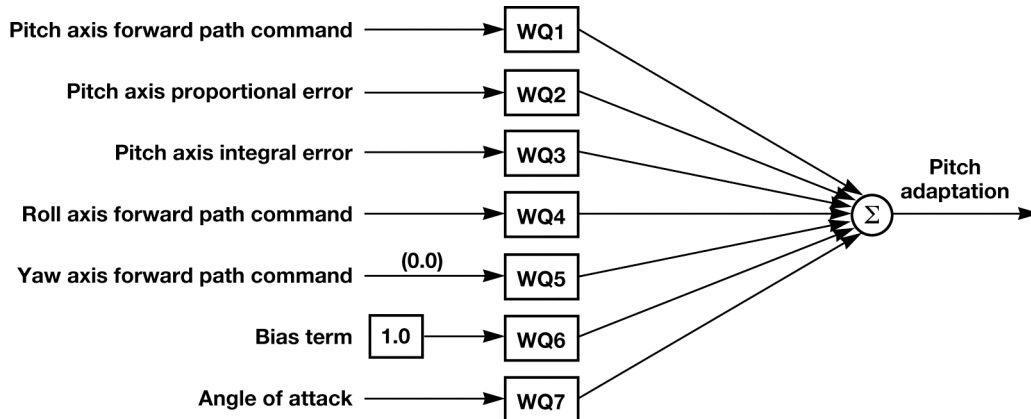

Figure 4. Pitch axis simplified sigma-pi neural network.

Figure 4 shows the specific inputs to the pitch axis neural network. For the simulated destabilizing failure the effect on the aircraft dynamics was primarily in the pitch axis. Specifically, adjustments to WQ1, WQ2, WQ3, and WQ7 provided the ability to account for the failure and regain stability and good reference model-following. WQ1, WQ2, and WQ3 acted to adjust the existing pitch rate PID control loop gains. WQ7 introduced a new angle of attack feedback to the symmetric stabilator. WQ4 and WQ5 provide cross-axis compensation for asymmetric failures[9] so did not contribute significantly to compensating for the symmetric destabilizing failure. WQ6 provides a bias term that, in conjunction with the pitch loop integrator, can retrim the airplane in response to the simulated failure. The bias term has minimal effect on the perturbation closed- and open-loop dynamics. Consequently, this report focuses on the contributions of WQ1, WQ2, WQ3, and WQ7.

The neural network weights were adjusted using the learning rule[6] shown in Eq. (3):

$$\dot{W} = -G(U_{err}B_a + L \ |U_{err}|W) \tag{3}$$

where W is the neural network weight vector, G is the adaptation gain vector, L is the error-modification term vector, U_{err} is the error compensation vector, B_a is the basis function vector, and dt is the iteration time step (0.0125 s). Dead zones were applied to the inputs to the neural network learning. These dead zones were used to prevent the neural networks from constantly adapting to small errors. The sizes of the dead zones were determined by using the six-degree-of-freedom software simulation. With no failures, moderately aggressive maneuvers were flown and the dead zone values were set to 120 percent of the observed errors. With a failure or change in vehicle behavior, when the elements of the error vector (U_{err}) exceed a dead zone threshold, the corresponding weight is adjusted until the error is reduced below the threshold. Limits were also placed on the weight magnitudes. These weight limits helped provide a limited authority system for initial flight-test purposes.

III.   Technical Performance Metrics

In-flight measurements of the dynamic characteristics of the overall system were taken and pilot-generated frequency sweeps were performed. The fast Fourier transformation (FFT) of the frequency sweep data allows for stability and closed-loop dynamic characteristics to be evaluated, enabling assessment of the impact of the simulated failure and the adaptive system on the overall system dynamic characteristics.

A. Stability Metric

The first objective of an adaptive system after encountering a destabilizing failure is to reestablish a stable airplane. Control law designers typically use stability margins to measure stability robustness of a system. Stability margins can be calculated from measurements taken during pilot-generated frequency sweeps. For the F-15 IFCS there is no explicit measurement or feedback of stability margin to the adaptive system. The adaptive portion of the

controller attempts to null out any excessive feedback error. This will directly increase performance and drive the system to a more stable damped state. In some cases this will indirectly increase stability margin, however as soon as sufficient performance is achieved the current system has no mechanism to maximize stability margin. Even though the F-15 IFCS system is not designed to maximize stability margin, it does provide increased margin when compared to the system without adaptation.

1. Open–Loop Transfer Function Reconstruction

In-flight measurement of the open-loop frequency response has been successfully accomplished on many airplanes. The frequency response can be used to calculate the stability margin of the system. Ideally this is accomplished with a programmed excitation injected directly to the surface command. The F-15 IFCS did not have the capability for programmed frequency sweeps, so pilot-generated sweeps were used. With pilot-generated excitation, direct measurement of the loop transfer function is complicated due to the presence of a feedforward command in the flight controller ($\dot{X}_{ref}$ in Fig. 2). To overcome this limitation the open-loop frequency response was reconstructed using a combination of flight measurements and transfer functions from a detailed model of the control system.[10]

The apparent plant portion of the flight system can be measured directly since this is a single-input system driven by the symmetric stabilator command. The apparent plant also encapsulates the real plant (the airplane) so the portion of the dynamic system with the largest uncertainty is directly measured. The apparent plant transfer function can be combined with mathematical models of the control system to produce a reconstructed open-loop transfer function. Similarly, the closed-loop frequency response can be reconstructed. This provides a check on the validity of the mathematical model of the controller. Figure 5 shows a comparison of a direct measurement of the closed-loop frequency response and a reconstructed closed-loop frequency response from flight data. The excellent match shown in Fig. 5 validates the mathematical model of the control law and the reconstruction method.


Figure 5. Closed-loop check to validate the frequency response reconstruction process.

B. Closed–Loop Metric

A secondary objective of an adaptive system is to reestablish controllability. The F-15 IFCS achieves this objective by forcing the system to track the previously-described onboard reference model. When good tracking is achieved, the pilot perceives a system with the frequency, damping, and time-delay characteristics of the onboard reference model. The relationship between control stick input and airplane response can be directly computed using FFT techniques from the pilot-generated frequency sweeps. The handling qualities community has defined a region of maximum unnoticeable added dynamics (MUAD).[11–13] The MUAD envelope is computed for a given lower-order system (in our case the onboard reference model). Flight-testing has shown that if the system response is within the MUAD envelope the response is indistinguishable from the lower-order system. Thus, one metric used to assess the effectiveness of the adaptation system was to show how well the system with adaptation compared to the MUAD envelope.

IV.   Flight Results

A careful build-up approach was used to test the simulated destabilization failure in flight. The canard multiplier destabilizing gain amplitude was increased (more negative), and controllability and maneuverability checks were accomplished. The system was tested first without adaptation engaged and then, after clearance, with adaptation engaged. At the time of the writing of this report, canard multiplier gains of up to -1.75 had been flown without adaptation and up to -1.5 with adaptation. Maneuvering clearance was completed but formation flight and air-to-air tracking handling qualities tasks were not yet complete.

A. Simulated Failures without Adaptation

The primary effect of the simulated destabilizing failure is to reduce the gain at low frequency in the symmetric stabilator control loop. Figure 6 shows the symmetric stabilator open-loop frequency response reconstructed from flight data with different destabilizing failure gain (CM) changes. The X marks on the plot indicate the crossover points where the low-frequency gain and phase margins are measured. This figure shows that as the canard multiplier gain absolute magnitude is increased the stability margins are significantly reduced. Table 1 shows the margin values obtained. At a canard multiplier of -1.75 the overall vehicle is just on the stable side of the neutral stability point.
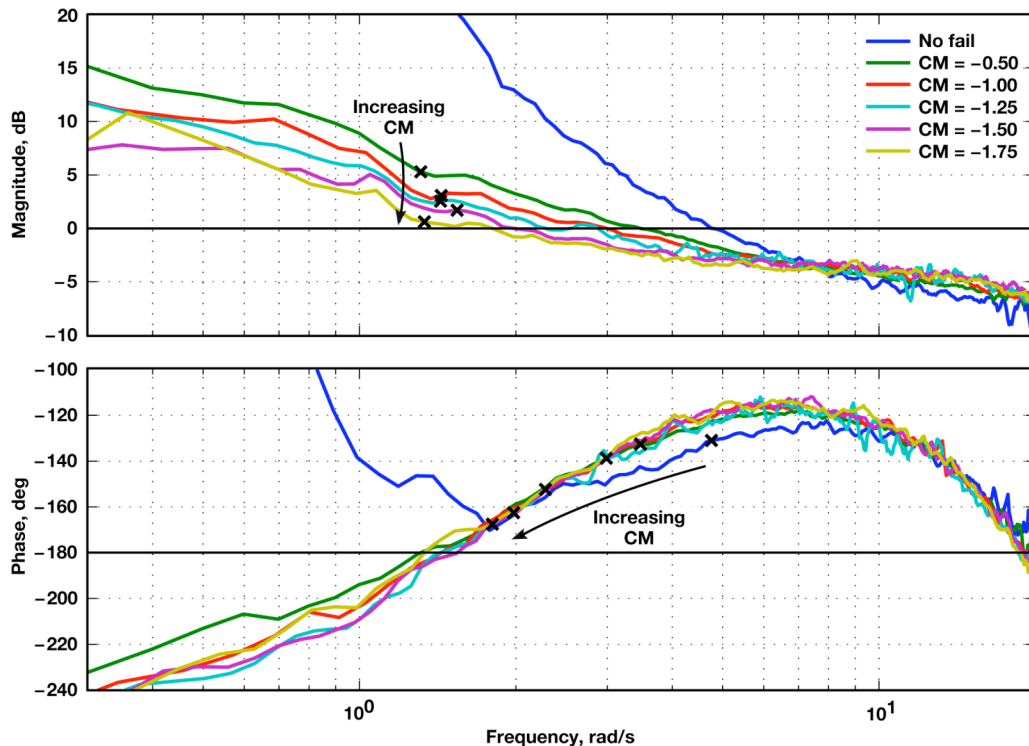

Figure 6. Effect of the destabilization gain on the symmetric stabilator open-loop frequency response with no adaptation.

American Institute of Aeronautics and Astronautics

Table 1. Flight-measured stability margins for the symmetric stabilator loop with simulated destabilizing failure and no adaptation.

Canard multiplier	Gain margin, dB	Phase margin, deg	Gain crossover, rad/s	Phase crossover, rad/s
-0.50	-5.3	47.3	3.47	1.31
-1.00	-3.1	41.3	2.99	1.44
-1.25	-2.5	27.6	2.28	1.43
-1.50	-1.7	17.4	1.98	1.54
-1.75	-0.6	12.4	1.80	1.33

The closed-loop frequency responses are shown in Fig. 7. As the destabilizing gain negative amplitude is increased the response diverges from the desired reference model transfer function. The MUAD region is also shown in this figure. The effect of the destabilizing gain is to move the closed-loop frequency response outside the desired MUAD region.

Figure 7. Effect of the destabilization gain on the closed-loop frequency response with no adaptation.

B. Simulated Failures with Adaptation

To show the behavior trend with changing canard multiplier, both open- and closed-loop frequency response data were processed. The data represent the state of the system during the pilot-generated frequency sweep after a series of training maneuvers had been performed. These results represent a non-time-varying approximation of the system after a fairly consistent set of training inputs.

1. Neural Network Training

For each canard multiplier a series of standard flight-test maneuvers were flown to assess the handling characteristics of the airplane (Fig. 8). The maneuvers were performed in the following order:

1. pitch, roll, and yaw raps for aeroservoelastic damping assessment;
2. pitch, roll, and yaw doublets for aerodynamic response assessment;
3. pitch and bank angle captures for handling qualities assessment;

9

American Institute of Aeronautics and Astronautics

4. left and right half-stick 360-degree rolls for loads maneuvering clearance;
5. left and right 3g wind-up turns for loads maneuvering clearance;
6. left and right 4g wind-up turns for loads maneuvering clearance;
7. pitch axis frequency sweep for dynamic stability and control assessment.

This series of maneuvers also provided training input for the neural network weights.

Figure 8. Representative time history of adaptive system training sequence.

2. Open–Loop Behavior with Adaptation

Figure 9 shows the symmetric stabilator open-loop frequency response with the adaptation system engaged. The open-loop frequency response was reconstructed using the previously-described process. The neural network weights were fixed at the flight-measured average value over the frequency sweep. Figure 9 shows that the adaptation system reshapes the open-loop frequency response to maintain a consistent gain crossover at approximately 3-4 rad/s regardless of the level of destabilizing gain. The resulting stability margins (Table 2) are significantly improved over the system without adaptation. Figure 10 shows a comparison of the stability margins with and without adaptation.

American Institute of Aeronautics and Astronautics

Figure 9. Effect of the destabilization gain on the symmetric stabilator open-loop frequency response with adaptation.


Figure 10. Effect of the destabilization gain on the symmetric stabilator loop stability margin.

11

Table 2. Flight-measured stability margins for the symmetric stabilator loop with simulated destabilizing failure and adaptation.

Canard multiplier	Gain margin, dB	Phase margin, deg	Gain crossover, rad/s	Phase crossover, rad/s
-0.50	-6.3	47.1	3.94	1.46
-1.00	-6.5	47.3	4.05	1.54
-1.25	-4.6	41.3	3.30	1.62
-1.50	-5.0	46.3	3.91	1.58

3. Closed–Loop Behavior with Adaptation

Figure 11 shows the closed-loop frequency response with adaptation enabled. For these data the frequency response was obtained directly using FFT of flight-measured inputs. The FFT averages the effect of the time-varying weights. With adaptation, the closed-loop response is closer to the desired frequency response of the reference model as compared to the non-adaptive system (Fig. 7). Although the objective of matching within the MUAD envelope was not achieved, the adaptive system did achieve more consistent performance over a range of destabilizing gains.


Figure 11. Effect of the destabilization gain on the closed-loop frequency response with adaptation.

4. Adaptation Characteristics

The dynamic characteristics of the adaptive system depend on the weights of the neural network. The weights are adjusted when errors are encountered as the vehicle is maneuvered. If these weights converge to steady-state values the adaptive portion of the control system becomes linear and time-invariant. The overall system should have lower errors and consequently better dynamic characteristics. The data shown in the previous section present results representing the system as it was after the series of training inputs. Ideally, the system should converge. Further training should not cause additional change to the system weights. The following section discusses details of the adaptive system training and convergence characteristics.

Figure 12 presents a time history of the neural network weights during this training maneuver sequence. As Fig. 12 shows, the initial transient caused by the failure insertion causes an immediate change in all the weights. Each of the maneuvers performed causes some change in the weights when the control feedback errors exceed the deadbands. A large weight change is caused by the frequency sweep at the end of the maneuver sequence. During

12

the high-frequency portion of the frequency sweep the airplane cannot (nor is it expected to) keep up with the command. Even though the onboard reference model rolls off at higher frequencies there is enough error generated to cause neural network weight adjustment.
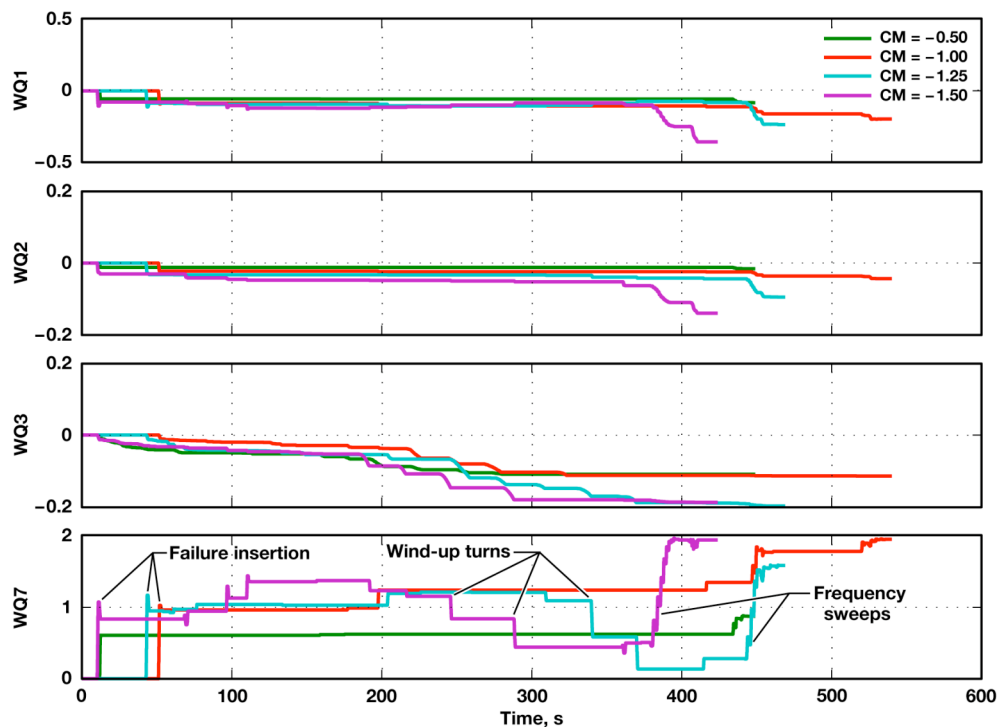
Figure 12. Effect of training on the adaptive system weights.

Of the four weights shown in Fig. 12, WQ2 shows the most consistent change in proportion to the size of the simulated failure. A desired behavior of a consistent and predictable adaptive system is that the weight change would be correlated with the severity of the failure. In the case of WQ2 this behavior is clearly shown. For the other weights the change is not as clear.

Ideally the size of the weight adjustments should decrease as the system adjusts to better follow the onboard reference model. With the exception of the large change due to the frequency sweep at the end, WQ1 and WQ2 exhibit that behavior. For the smaller failures, WQ3 and WQ7 show a tendency to converge. For the larger failures it is difficult to say that WQ3 and WQ7 have converged.

WQ3 shows more continuous adjustment when compared to the other weights. The input to the learning function for this weight is more heavily dependant on integral error versus proportional error than for other weights. In this design the bias term (WQ6) provides an integrator-like function that unloads the control loop integrator value. In this way the neural network can minimize both the proportional and integral errors. Upon failure insertion, the integral error tends to move away from zero and it takes some time for the neural network bias term to reduce this error back to zero. Because the input to the learning function for WQ3 remains outside the deadzone the weight shows more continuous adjustment.

The majority of the weight adjustments are in a consistent direction, the exception being most apparent in WQ7. The pushover at the conclusion of the wind-up turns causes this weight to reduce in magnitude. Other maneuvers such as the frequency sweep tend to increase the magnitude of this weight.

An undesired characteristic of this neural network is that it did not converge to a clear steady-state solution. This is especially true for the larger destabilizing failures. The frequency sweep was repeated for a destabilizing gain of -1.5 without resetting the adaptation system weights. Figure 13 shows the how the weights continued to change. WQ1 actually reached its lower limit of -0.5 during the second frequency sweep. It should be noted that the training sequence consists of "normal" maneuvers that might be performed. The exception is the frequency sweep where the high-frequency portion would not occur in normal flight; however, abrupt maneuvering in response to a failure, evasive collision avoidance maneuvers, or gust encounters may have a similar effect on the learning system.

American Institute of Aeronautics and Astronautics

The significance of the adaptation due to the second frequency sweep can be seen in Fig. 14. The open-loop transfer function was reconstructed using flight-measured weights at different time points. The adaptation continues to raise the loop gain; this causes a reduction in high-frequency gain margin and could eventually cause an adverse structural mode interaction.
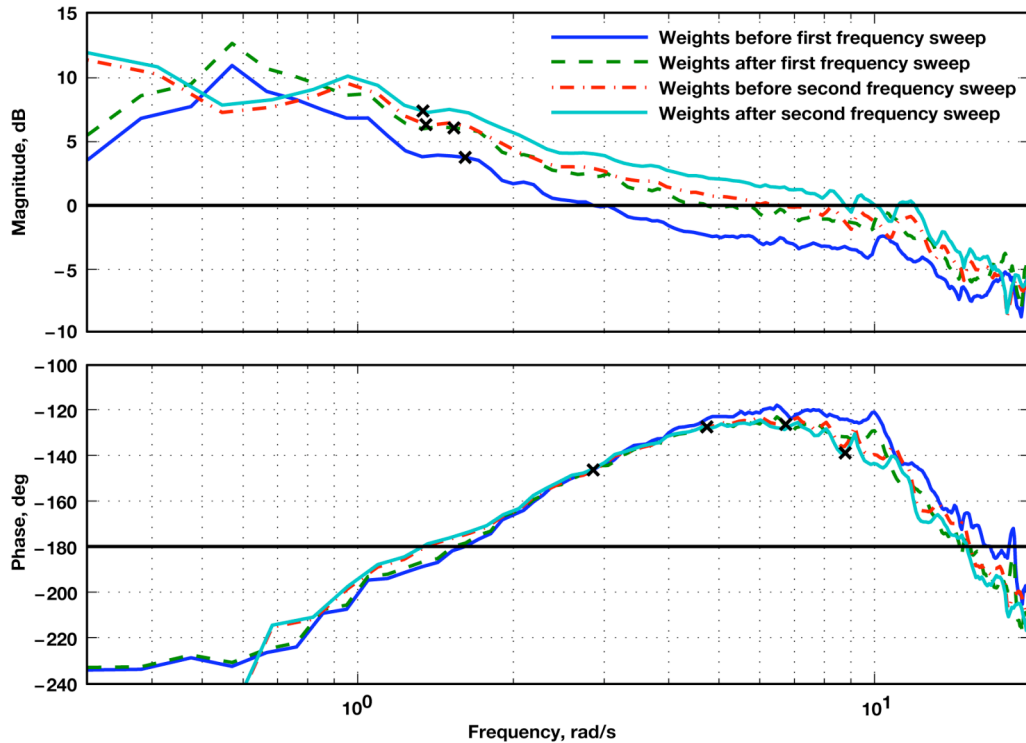

Figure 13. Effect of training on the adaptive system weights for a destabilizing gain of -1.5 with a second frequency sweep.

American Institute of Aeronautics and Astronautics

Figure 14. Effect of adaptive system weights on the open-loop frequency response for a destabilizing gain of -1.5.

Figure 15 shows the corresponding effect on the closed-loop frequency response. The figure shows that the after the second frequency sweep a better match to the onboard reference model is achieved. The data also show a resonant peak tendency at approximately 15 rad/s. These data are reconstructed using constant weights at different time periods. Figure 16 shows a comparison of reconstructed data with directly-measured closed-loop frequency response. The reconstructed data bracket the directly-measured data and add some credence to the validity of the reconstructed data.

The response change due to weight adaptation illustrates the inherent conflict between achieving the desired closed-loop response while maintaining adequate stability margins. Since the F-15 IFCS system does not have specific measurement of stability margin, reduction in margin is not an explicit input to the adaptation. Only when the margin is reduced to the point where oscillations begin to occur will the adaptation begin to respond. Another potential solution would be to adjust the desired onboard reference model so that the closed-loop objective could be achieved while still maintaining some stability margin. Planned follow-on work is attempting to provide improvements to the adaptive system to provide better convergence characteristics.

Figure 15. Effect of adaptive system weights on the closed-loop frequency response for a destabilizing gain of -1.5.
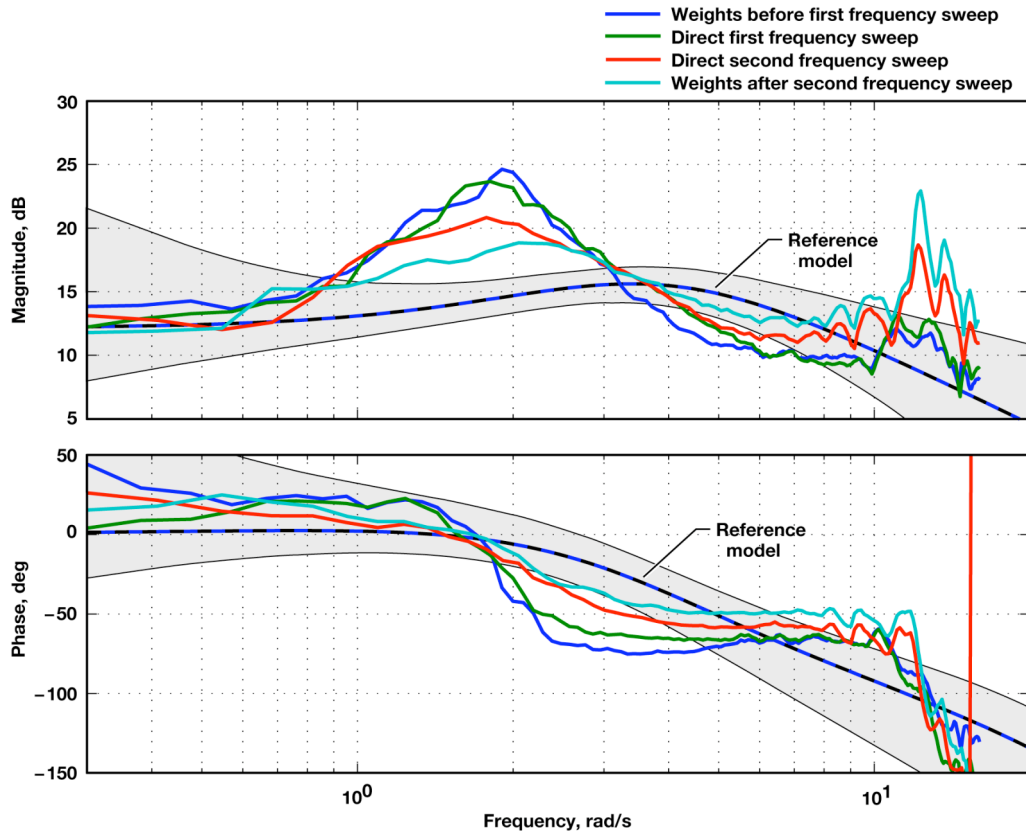
Figure 16. Comparison of reconstructed and direct calculation of the closed-loop frequency response for a destabilizing gain of -1.5.

V. Conclusion

A direct adaptive neural-network-based flight control system was developed to provide increased resiliency to extreme changes in airplane behavior. A simulated destabilizing failure was used to challenge the adaptive system. Typical flight-test maneuvers were flown to train the system. Pilot-generated frequency sweeps provided data to analyze the resulting dynamic performance, and comparisons were made to the non-adaptive system.

The simulated destabilizing failure produced a significant change in the pitch axis dynamic characteristics of the airplane. Symmetric stabilator loop stability margin measurements were reconstructed from flight data. Direct measurement of the stability margins was not possible since a programmed surface sweep was not available; instead, the open-loop transfer function was reconstructed using flight measurement of the apparent plant and a high-fidelity model of the control system. This reconstruction process was validated by reconstructing the closed-loop frequency response and comparing it to a directly-measured closed-loop frequency response.

Open-loop analysis of the flight-test results showed that a destabilizing angle-of-attack-to-canard gain multiplier of -1.75 produced essentially a neutrally stable vehicle without adaptation. The failure produced a lower loop gain and consequently reduced the low-frequency gain margin. When the adaptation system was engaged, the adaptive system weights adjusted to increase the symmetric stabilator loop gain, providing a more consistent gain crossover and higher stability margins regardless of the size of the simulated failure.

The closed-loop performance of the system was assessed in terms of its ability to reestablish good pitch axis model-following. The desire is to adapt so that the resulting closed-loop frequency response falls within the maximum unnoticeable added dynamics (MUAD) envelope defined by the desired onboard reference model. Without adaptation the closed-loop frequency response showed significant changes from the desired response. The response departed from a typical second-order response and strayed from the MUAD envelope. The adaptation system tended to shape the response to more closely follow the desired onboard reference model. To achieve better reference model-following a higher loop gain would be required. Eventually, problems with the high-frequency stability margin or structural mode interaction could result. A suggested improvement to the system is to add adaptation to the onboard reference model parameters so that model tracking could be more easily achieved while

American Institute of Aeronautics and Astronautics

maintaining some stability margin. Additionally, it would be beneficial to directly incorporate the stability margin measurement into the adaptation algorithm.

The flight-test results showed that the neural network weights did not necessarily converge to a steady-state value. Especially for the larger destabilizing failures, continued maneuvering caused the weights to continue to change. For most of the analysis described in this report the weights were, somewhat arbitrarily, fixed to the average value over the piloted frequency sweep. Using frequency response analysis necessitates making a time-invariant approximation to what is in actuality a time-varying system. Comparisons made with weights obtained at various time points showed that this approximation was reasonable. For future applications better convergence properties are desired.

The F-15 Intelligent Flight Control System provided an initial flight evaluation of a direct adaptive neural-network-based flight control system. The system provided increased stability margins in the presence of large destabilizing failures. The convergent properties of this initial system warrant additional improvement since continued maneuvering caused continued adaptation change. The adaptive system provided better closed-loop behavior with improved matching of the onboard reference model. These flight data provide a basis for analysis and understanding of this type of adaptive system. This flight-test experience and continued testing will help to advance adaptive controls technology as an option for future aerospace vehicles.

References

[1] USAF Safety Investigation Report, 29 AUG. 84, B-1A 740159.

[2] Smolka, J.W., Walker, L.A., Johnson, Major G.H., Schkolnik, G.S., Berger, C.W., Conners, T.R., Orme, J.S., Shy, K.S., and Wood, C.B., "F-15 ACTIVE Flight Research Program," *1996 Report to the Aerospace Profession Fortieth Symposium Proceedings*, 1996, pp. 112-145.

[3] Calise, A.J., Lee, S., and Sharma, M., "Direct Adaptive Reconfigurable Control of a Tailless Fighter Aircraft," AIAA-98-4108, Aug. 1998.

[4] Rysdyk, R.T., and Calise, A.J., "Fault Tolerant Flight Control Via Adaptive Neural Network Augmentation," AIAA-98-4483, Aug. 1998.

[5] Williams-Hayes, P.S., "Flight Test Implementation of a Second Generation Intelligent Flight Control System," NASA/TM-2005-213669, Nov. 2005.

[6] Kaneshige, J., Bull, J., and Totah, J.J., "Generic Neural Flight Control and Autopilot System," AIAA-2000-4281, Aug. 2000.

[7] Burken, J.J., Williams-Hayes, P., Kaneshige, J.T., and Stachowiak, S.J., "Reconfigurable Control with Neural Network Augmentation for a Modified F-15 Aircraft," NASA/TM-2006-213678, Apr. 2006.

[8] *Flying Qualities of Piloted Vehicles*, U.S. Department of Defense, MIL-STD-1797, March 31, 1987.

[9] Bosworth, J.T., Williams-Hayes, P.S., "Flight Test Results from the NF-15B Intelligent Flight Control System (IFCS) Project with Adaptation to a Simulated Stabilator Failure," NASA/TM-2007-214629, Dec. 2007.

[10] Tischler, M.B., Remple, R.K. "Aircraft and Rotorcraft System Identification Engineering Methods with Flight Test Examples," AIAA Education Series, 2006.

[11] Wood, J.R., and Hodgkinson, J., "Definition of Acceptable Levels of Mismatch for Equivalent Systems of Augmented CTOL (Conventional Take-off and Landing) Aircraft," McDonnell Aircraft Corporation, Report number MDC A6792, St. Louis, MO, Dec. 1980.

[12] Hoh, R.H., Mitchell, D.G., Ashkenas, I.L, Klein, R.H., Heffley, R.K., and Hodgkinson, J., "Proposed MIL Standard and Handbook – Flying Qualities of Air Vehicles, Volume II: Proposed MIL Handbook," AFWAL-TR-82-3081 Volume II, Wright-Patterson Air Force Base, OH, Nov. 1982.

[13] Mitchell, D.G., Hoh, R.H., He, C., and Strope, K., "Determination of Maximum Unnoticeable Added Dynamics," AIAA-2006-6492, Aug. 2006.