

Hierarchical Approach to Adaptive Control for Improved Flight Safety

Moshe Idan,^{*} Matthew Johnson,[†] and Anthony J. Calise[‡]
Georgia Institute of Technology, Atlanta, Georgia 30332-0150

Following failures of primary aerodynamic actuators, safe flight can be maintained by introducing alternative actuation systems, such as analytically redundant secondary aerodynamic surfaces and propulsion, for higher-priority stability and control augmentation tasks. An intelligent hierarchical flight control system architecture is presented that is designed using nonlinear adaptive synthesis techniques and online learning neural networks to enhance flight safety. Pseudocontrol hedging is used for proper adaptation in the presence of actuator saturation, rate limits, and failure. The hierarchical structure incorporates nonactive secondary actuation channels that are engaged after a failure of a primary control surface is encountered. The methodology requires only the knowledge that a failure in a specific actuator has occurred. A model of the failed aircraft, the failure type, and the failure size need not to be known: The neural network element of the secondary channel will adapt to the failed actuator effect. The secondary control channels are designed to account for the typically lower authority and degraded performance that can be expected with secondary actuation systems. The proposed hierarchical flight control architecture is attractive, in particular, as a retrofit to existing certified flight control systems for enhanced flight safety. The proposed flight control architecture is evaluated in a nonlinear flight simulation environment, demonstrating its retrofit features.

Introduction

EVER since powered controls were introduced to civil aviation, they have been implemented in a fault-tolerant manner by the use of redundant sources of hydraulic power. Nevertheless, it is still possible that a loss of primary control in one or more axes can occur as a result of a structural failure that either renders critical control linkages inoperable, or that results in a loss of the redundant hydraulic supplies. Landing an aircraft in the presence of actuator failures presents a challenge to even the most experienced pilot. Recent accidents have been caused by a failure of a single actuator or a complete loss of the whole hydraulic actuation system.¹

Conventional flight control requires extensive gain scheduling for a large number of operating points. When such a controller must be extended to account for total loss of a primary actuation system, redesign is required for each potential failure case. Many types of failures can be envisioned, including, but not limited to, hardovers, loss of actuator effectiveness, and free-floating actuators. This leads to a very large scheduling table, making it tedious from design and real-time implementation standpoints. In addition, a truly fault-tolerant control system must also be able to accommodate nonanticipated failures.

Flight control reconfiguration to accommodate actuator failures has been extensively addressed in the last three decades, utilizing a multitude of control design methods and reallocation schemes. Here one finds the pseudoinverse method,^{2,3} model parameter identification coupled with feedback linearization,⁴ online system identification with online receding horizon/model predictive optimal control,⁵ robust servomechanism and control allocation based on quadratic programming formulation,⁶ and quantitative feedback theory (QFT) coupled with reduced-order linear inversion control and adaptive

filtering,⁷ to name but a few. These are model-based methods that either assume knowledge of the aircraft model or that implement an online parameter identification scheme. Most of these methods are limited to linear aircraft models,^{2,3,5–7} and some even require the knowledge of these models in normal and failed conditions.^{2,3,6} In several cases, the actuator failures are modeled as variations in the linear system parameters,^{2,3,5} hence, limiting the type of faults the method can accommodate. Alternatively, some methods assume the knowledge of the failure type and size, thus requiring a fault detection and identification algorithm.^{4,6}

Neural-network- (NN-) based adaptive flight control, within the setting of feedback inversion control, has been shown to require no gain scheduling and to be only minimally model dependent.^{8–10} These flight control systems can accommodate a multitude of unknown actuator failures that act as disturbances on the aircraft. Hence, they provide an attractive candidate flight control architecture to ensure flight safety in the presence of unknown actuator failures. Their application to civil transports requires special attention because redundant actuation is only possible through low-bandwidth and low-authority mechanisms that are usually not intended to be active as a part of the primary flight control system. However, due to their complexity, such adaptive control systems could be used only on aircraft with digital fly-by-wire systems.

The objective of this work is to design an intelligent nonlinear adaptive control architecture that can respond to faults in the system, by utilizing redundancy in the controls. Calise et al.⁹ and Johnson et al.¹⁰ demonstrate that such a system could effectively control an aircraft with major actuator failures. Bull et al.¹¹ and Burken et al.¹² use a nonadaptive, gain-scheduled control design for pure propulsion control to provide stability augmentation for a large transport aircraft without any aerodynamic actuation. Rysdek⁸ demonstrates that similar performance is attainable by employing an adaptive controller without gain scheduling, using a linear model at a single flight condition for feedback inversion. However, those results were limited to examining small command inputs so that position and rate saturation are avoided, to guarantee stability and proper NN adaptation.

In this paper, the problem of continuous control in the presence of both partial and complete loss of a single or multiple actuators is addressed, utilizing all of the remaining control effectors. A hierarchy among these effectors with respect to forces and moments about each axis is developed, leading to a hierarchical adaptive control architecture suitable for enhanced flight safety. This control architecture

Received 18 June 2001; revision received 6 July 2002; accepted for publication 6 July 2002. Copyright © 2002 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/02 \$10.00 in correspondence with the CCC.

^{*}Visiting Research Scholar, School of Aerospace Engineering; currently Senior Lecturer, Faculty of Aerospace Engineering, Technion—Israel Institute of Technology, Haifa 32000, Israel. Associate Fellow AIAA.

[†]Graduate Research Assistant, School of Aerospace Engineering. Student Member AIAA.

[‡]Professor, School of Aerospace Engineering. Fellow AIAA.

can accommodate actuator failures without pilot intervention. The only visible effect on the aircraft is system performance degradation resulting from the alternative secondary actuators having generally lower effectiveness, bandwidth, and authority in the primary flight control tasks. One of the main advantages of the adaptive control design is that it does not require a dynamic model of the failed aircraft or the type and size of the failure: The adaptive, NN-based element will adapt and compensate for the overall effect of the failed actuator.

Available controls considered in this study include the aerodynamic surfaces (ailerons, elevator, and rudder) and the propulsion system. Response to loss of one or several of the aerodynamic control surfaces is investigated. A control system that switches effectors is in greater danger of saturation, which may lead to difficulties in the adaptation process of the proposed nonlinear NN-based control algorithms. A recently developed pseudocontrolhedging (PCH) methodology^{10,13,14} is employed to protect the system from incorrect adaptation in the presence of slow actuation, actuator saturation, and failure. The PCH technique is also used to train the NNs of the secondary control systems online before they are engaged (learning while not in control). This way, the secondary control systems are ready to take over as failures occur in the primary systems.

A central feature of the proposed hierarchical flight control architecture is its capability to enhance the safety of conventional, previously certified flight control systems. The secondary control channels of this methodology can simply be added to an existing system, without intervening with its normal operation. The hierarchical switching between the various secondary channels will be introduced only after a failure occurs and has been identified. These and other features of the proposed hierarchical flight control architecture are illustrated in a numerical flight simulation environment.

Hierarchical Adaptive Control Architecture

The hierarchical adaptive control scheme proposed here addresses sequential utilization of analytically redundant actuation systems, engaged as secondary actuators in response to failures in the primary actuators. The secondary systems can be actuators that normally are not used for flight control, for example, flaps or spoilers, or actuators that normally perform control tasks that are not safety critical. A typical example for the latter is, after encountering an aileron failure, to use the rudder for roll control, while compromising turn coordination. The suggested hierarchical architecture allows for multiple levels of sequentially lower-priority actuation if a higher-level actuation system has been exhausted (saturated) or failed.

After a failure, the use of a secondary actuator for a primary stability and control augmentation task entails an unavoidable degradation in the system performance resulting from the following:

- 1) A control task that may have been performed by the secondary actuator may need to be abandoned.
- 2) The performance (e.g., bandwidth, maximum magnitude and/or rate, or disturbance rejection effectiveness) that can be achieved using this secondary actuator may be considerably lower than the nominal.
- 3) The failed actuator, if frozen at a nonneutral position or if moving unfavorably, may introduce forces and moments that have to be counteracted by the secondary actuator.
- 4) Upsets and transients that occur before engaging the secondary actuation system must be returned to commanded values.

Beyond the physical limitations of the secondary actuators, many of the preceding shortcomings are addressed by the hierarchical control strategy suggested herein. The key elements of this methodology are secondary control loops that are engaged only on demand following failures, but otherwise are continuously trained to reduce transients when engaged.

The secondary control loops are designed using adaptive control techniques that incorporate online training NNs, whereas the primary control loops can be either conventional or adaptive. Adaptive control is introduced in the secondary control loops mainly to address model uncertainty (including the adversary effects of the failed primary actuator) and to reduce gain scheduling. In particular, training of the NNs in the secondary channels during actuator saturation and/or during periods of normal operation (when the secondary actuators they are intended to command are not engaged) is addressed. The switching logic between the various control channels within the hierarchical setup assumes only the knowledge that a particular failure has occurred, without the need to know the specific features (type, magnitude, etc.) of that failure.

The secondary adaptive control channels of the architecture are designed using feedback linearization, linear control design, and NN compensation to address system linearization errors. Training of the NNs during actuator saturation and failure and while not in control is performed using the PCH methodology. Consequently, the secondary actuation channels can be engaged with minimal transients.

The proposed hierarchical controller architecture may vary between applications. This variation depends on the available secondary actuation sources and the choices made for the order in which they are to be employed. As an example, a discussion of longitudinal pitch rate and flight speed control is presented here. The primary actuation devices are elevator and throttle, which control pitch rate and flight speed, respectively. The primary control channels are assumed to be conventional (nonadaptive) control systems. In case of an elevator failure, flight speed control is compromised in favor of the higher-priority propulsion-based pitch rate stability augmentation and control. A block diagram for this longitudinal control problem is presented in Fig. 1. The blocks above the dotted line represent the normal operation mode, whereas the bottom part

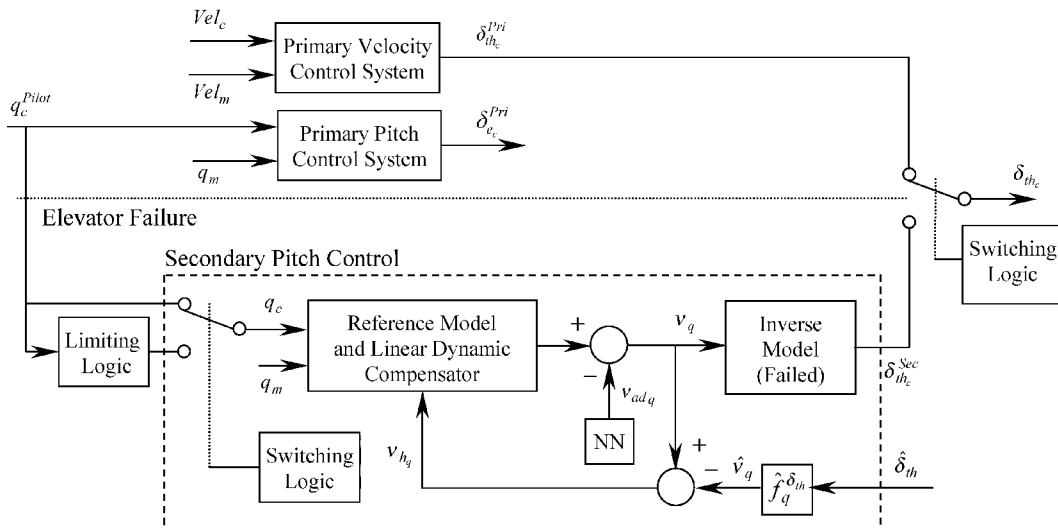


Fig. 1 Hierarchical adaptive control architecture.

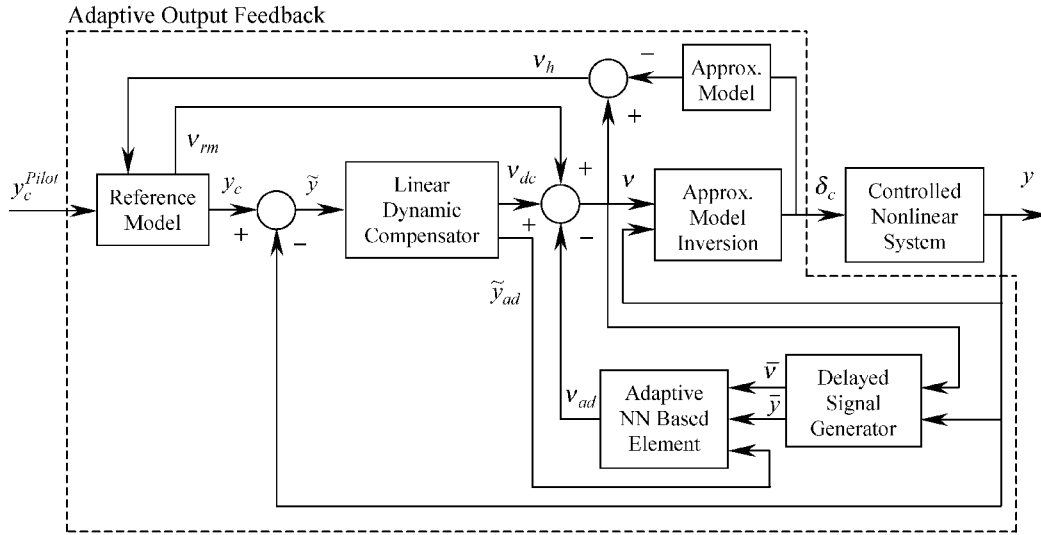


Fig. 2 Model reference adaptive control setup, including approximate dynamic inversion and PCH compensation.

of Fig. 1 depicts the elevator-failed case. The transition between the two modes is controlled by a switching logic. A detailed description of an adaptive control channel design including PCH is delayed to the next section.

In nominal fail-free operation, two control channels are activated: flight speed control using the throttle and pitch rate control using the elevator. Vel_c^{pilot} and q_c^{pilot} are the pilot stick commands, and Vel_m and q_m are, respectively, the measured flight speed and pitch rate used in the primary control logic. The outputs of the controllers are the throttle and elevator servo commands δ_{thc} and δ_{ec} , respectively. Here we assume that these two control loops were designed using conventional control techniques. In parallel, an alternative channel for pitch rate control using the throttle is executed. It has no effect on the aircraft control until an elevator failure is detected or measured. After an elevator failure is encountered, the switching logic disengages the flight speed control circuit and engages the secondary channel of pitch rate control using the throttle (the two switches in Fig. 1 are moved down). The flight speed control loop continues to function with no effect on the actual throttle commands.

The secondary pitch rate control channel using the throttle (propulsion control) is designed to incorporate the following important features:

- 1) An approximate inverse model is utilized that assumes pitch rate is controlled using the throttle only.
- 2) The reference model and the linear compensator are designed for the anticipated lower performance of this degraded mode, caused by the lower effectiveness and lower bandwidth of the propulsion actuation for pitch rate control. This is also addressed by the command limiting logic acting on the pilot command, represented by the “Limiting Logic” block in Fig. 1.
- 3) The relative degree of the throttle to pitch rate transmission, which is different from the case of elevator control, is accounted for in the design of the compensator and the reference model.
- 4) The PCH signal is based on a measured or estimated throttle position. In the normal operation mode, δ_{th} is dictated by the speed control channel (top of Fig. 1), thus ensuring correct training of the NN, even though its output does not affect the controlled aircraft. This way, the NN is properly trained and ready to be engaged with minimal transients in the event of an elevator failure. Once this control loop takes over, PCH is carried out in a standard manner based on the estimated throttle position commanded by this channel.

The case of partial elevator failure, such as partial mechanical damage or mechanical blocking, can also be addressed using this architecture and switching logic. If the damage is significant, requiring additional pitch rate authority for flight safety, the propulsion channel can be engaged in parallel with the damaged elevator channel. The two control loops will “cooperate” through the measured pitch rate feedback.

The hierarchical controller architecture just presented can be applied to any longitudinal or lateral control channel of an aircraft. Moreover, this hierarchical structure can be further extended to use additional control effectors in a particular control channel. For example, roll rate control, normally performed with the ailerons, can be diverted to the rudder after the ailerons have failed and then further redirected to propulsion control for additional authority or due to rudder failure.

Nonlinear Adaptive Flight Control Channel

Each secondary control channel of the proposed hierarchical flight control system is constructed in a model reference adaptive control (MRAC) scheme while operating with possibly saturated or failed actuators, or while not in control of the aircraft. Figure 2 presents the conceptual layout, which incorporates an approximate dynamic inversion block, a linear compensator, and an online adaptive NN element. The system is driven by the outputs of a reference model block, which has an input from the PCH element.

The main advantage of the proposed control setup is in its minimal dependence on a specific aircraft model. The adaptive NN is used to compensate for a wide range of modeling (inversion) errors, which may include the effect of failed actuators. The compensator design is straightforward and relies mainly on linear control theory. In this work, a recently developed output feedback adaptive control design is incorporated, utilizing only actually available measured signals while avoiding state estimation.¹⁵ The NN adaptation rule results from nonlinear stability analysis, which ensures that the error signals and network weights are bounded. In this section, the various elements of this controller setup are discussed.

Approximate System Linearization

One of the common methods for controlling nonlinear dynamic systems is based on approximate feedback linearization.¹⁶ We illustrate this approach for a single-input/single-output dynamic system described by the state-space and measurement equations,

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \delta), \quad y = h(\mathbf{x}) \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the system state vector, δ is the scalar control input, y is the scalar measurement and regulated output, and $\mathbf{f}(\cdot)$ and $h(\cdot)$ are sufficiently smooth functions. Additional outputs, which are not regulated, may be incorporated into the design approach. The only modeling assumption is that the relative degree ($r \leq n$) of the output is known. Thus, the r th derivative of the output is the first derivative of the output that is “strongly” affected by the control, that is,

$$y^{(r)} = h_r(\mathbf{x}, \delta) \quad (2)$$

where $h_r(\mathbf{x}, \delta)$ is the r th Lee derivative of $h(\mathbf{x})$ along $\mathbf{f}(\mathbf{x}, \delta)$. Approximate feedback linearization is performed by introducing the transformation

$$v = \hat{h}_r(y, \delta) \quad (3)$$

where $\hat{h}_r(y, \delta)$ is the best available approximation of $h_r(\mathbf{x}, \delta)$ that is invertible with respect to δ and that uses only the available measurements. The variable v is commonly referred to as pseudocontrol. The actuator commanded input is defined based on this approximate model and is given by

$$\delta_{\text{cmd}} = \hat{h}_r^{-1}(y, v) \quad (4)$$

In the case of perfect actuation ($\delta = \delta_{\text{cmd}}$), this approximation leads to an inversion error

$$\Delta = h_r(\mathbf{x}, \delta) - \hat{h}_r(y, \delta) = h_r(\mathbf{x}, \delta) - v \quad (5)$$

present in the $y^{(r)}$ dynamics, which is now expressed as

$$y^{(r)} = v + \Delta \quad (6)$$

When the system output y is required to track a known bounded input y_c , the pseudocontrol is chosen to have the form

$$v = v_{\text{rm}} + v_{\text{dc}} - v_{\text{ad}} \quad (7)$$

where v_{dc} is the output of a stable linear dynamic compensator, having only the tracking error as an input; v_{rm} is the pseudocontrol signal generated by a stable reference model that accounts for imperfect actuation (PCH); and v_{ad} is the adaptive control signal (NN output) designed to cancel the inversion error. In the case of perfect actuation ($\delta = \delta_{\text{cmd}}$), there will be no PCH and $v_{\text{rm}} = y_c^{(r)}$, which is the r th derivative of the commanded trajectory. Hence, the commanded trajectory has to be generated using an at least r th-order reference model, for example,

$$y_c^{(r)} = h_{\text{rm}}(y_c, \dot{y}_c, \dots, y_c^{(r-1)}, y_c^{\text{pilot}}) \quad (8)$$

where y_c^{pilot} is an external command signal. The details on PCH will be presented at the end of this section.

Let the tracking error \tilde{y} be defined as

$$\tilde{y} = y_c - y \quad (9)$$

When Eqs. (6–8) are applied, the r th derivative of \tilde{y} can be expressed as

$$\tilde{y}_r = -v_{\text{dc}} - (\Delta - v_{\text{ad}}) \quad (10)$$

When it is assumed that the adaptive control signal v_{ad} exactly cancels the model inversion error Δ , the linear tracking error Eq. (10) represents a linear system with r poles at the origin.

From Eqs. (4), (5), and (7), it is evident that the modeling error Δ is an implicit function of the adaptive control signal v_{ad} . To guarantee the existence and uniqueness of v_{ad} that can at least theoretically cancel the model inversion error Δ in Eq. (10), it is assumed that the map from v_{ad} to Δ is a contraction. It can be shown that because $\Delta = h_r - \hat{h}_r$, this assumption is equivalent to the assumption that the following control effectiveness sign and magnitude conditions are satisfied¹⁵:

$$\text{sign}\left(\frac{\partial \hat{h}_r}{\partial \delta}\right) = \text{sign}\left(\frac{\partial h_r}{\partial \delta}\right) \quad (11)$$

$$\infty > \left| \frac{\partial \hat{h}_r}{\partial \delta} \right| > \left| \frac{\partial h_r}{\partial \delta} \right| / 2 > 0 \quad (12)$$

Linear Compensator Design

As presented in Eq. (10), the ideally linearized plant consists of r poles at the origin. A stable linear dynamic compensator that acts on the tracking error signal \tilde{y} is designed so that these r poles are stabilized. Classically, this is achieved using standard lead-lag compensators, although additional integral action can be incorporated to address steady-state performance. The latter could potentially slow down the closed-loop system response in return for improved steady-state characteristics. In general, the linear compensator can be designed using any linear control design technique as long as the linearized closed-loop system is stable.

The linear compensator within the adaptive output feedback setting (see Fig. 2) has an additional nonstandard output \tilde{y}_{ad} that is a linear combination of the compensator states and its input. This output is designed so that the transfer function between a low-pass filtered model inversion error Δ_f and \tilde{y}_{ad} is strictly positive real (SPR).¹⁵ This second compensator output is required to construct NN training laws that depend only on measurable quantities.

NN for Inversion Error Compensation

The dynamic model inversion and, thus, the nonlinear system linearization, are, in general, not exact, mainly because the exact nonlinear model is not known, full-state information is not available, or inversion is too complex to be implemented in its exact form. Clearly, simplified inversion functions are advantageous from a real-time implementation perspective and, thus, are often adopted when adequate feedback linearization error compensation is incorporated in the controller design.

The theory of Ref. 15 for output feedback NN-based adaptive control supports only linear-in-parameters NNs, where

$$v_{\text{ad}} = \hat{\mathbf{W}}^T \phi(\boldsymbol{\eta}) \quad (13)$$

Here, $\hat{\mathbf{W}}$ are the NN weights that are adjustable online, and $\phi(\boldsymbol{\eta})$ are the basis functions, typically of the radial basis function (RBF) type, which are known to possess the universal approximation property of linear-in-parameters NNs.^{17,18} The NN input vector $\boldsymbol{\eta}$ is constructed from present and past measurements of the tracking error and the pseudocontrol signal, generated by the delayed signal generation block of Fig. 2. When observability of the controlled aircraft is assumed, the universal approximation property of the NN guarantees that the unknown model inversion error can be approximated using Eq. (13).

To ensure that NN online training is performed based on only the available tracking error signal, a low-pass filter is incorporated in the NN training loop. This filter is introduced to generate an SPR transfer function (discussed earlier) and is used to derive an output-only based NN training law within a Lyapunov stability analysis (see Ref. 15). The analysis guarantees uniform ultimate boundedness of the closed-loop system tracking errors and NN weight errors. The NN training law resembles a classical backpropagation algorithm with a σ -modification term:

$$\dot{\hat{\mathbf{W}}} = -F[\tilde{y}_{\text{ad}}\phi_f + \sigma\hat{\mathbf{W}}] \quad (14)$$

where F is a gain matrix, defining the “learning rate”; \tilde{y}_{ad} is the second output of the dynamic compensator described earlier; ϕ_f are the RBF outputs filtered through the low-pass SPR filter; and $\sigma\hat{\mathbf{W}}$ is the σ -modification term.

Examining the model and tracking error, Eqs. (5) and (10), reveals that any failure modeled as a (not necessarily known) function of the system states can be addressed using the adaptive control scheme presented in the preceding section. This failure characterization is not overly restrictive because such functional dependence can represent most of the commonly encountered actuator failures, such as position frozen actuators, hardovers, free-floating aerodynamic surfaces, and many more. Hence, the NN-based adaptive secondary actuation systems are designed while disregarding the possible failures of the primary actuators. The online tuned NN of these channels will adapt to the failure driven inputs, interpreted as modeling errors, and compensate for their effect.

Pseudocontrol Hedging

PCH introduces a modification to previous work on NN-based MRAC. It is used to address NN adaptation difficulties arising from various actuation anomalies, including actuator position and/or rate saturation, discrete (magnitude quantized) control, actuator dynamics (including pure time delay), and partial or complete actuator failures.^{10,13,14}

NN training difficulties occur when unmodeled nonlinear actuator characteristics are encountered. For example, unless an adaptive process is protected in some way, saturation that results from failed operation (or from attempting to adapt while not in control) quickly leads to NN windup. The main idea behind the PCH methodology is to limit or hedge the reference model of a MRAC architecture to prevent the adaptive element from attempting to adapt to these characteristics, when they are present, while at the same time adapting to other sources of inversion error for which compensation is possible.

The conceptual idea of the PCH method is to move the reference model backward by an estimate of the amount the controlled system did not move due to selected actuator characteristics (such as position and rate limits, time delays, etc.). In effect, the reference model, which produces the commanded pseudocontrol, is limited or hedged according to the difference between the commanded and actually achieved pseudocontrol. PCH prevents the NN from adapting erroneously to actuator saturation or failure. With PCH, the NN is trained correctly using only achievable pseudocontrol signals. The same concept holds when the pseudocontrol action is due to a different control logic and not the MRAC that incorporates the NN (learning while not in control).

To briefly review the PCH concept, consider the plant dynamics of Eq. (1). The pseudocontrol signal defined in Eq. (3) represents the r th derivative of the commanded signal when no PCH is included, whereas the actuator commands are given by Eq. (4). The dynamic inversion element is designed without consideration of the actuator model. Hence, this actuator command δ_{cmd} will not equal the actuator position δ due to its dynamics, saturation, and/or failure. The pseudocontrol hedge signal v_h is defined as the difference between the commanded pseudocontrol input and an estimate of the actually achieved pseudocontrol. This difference is nonzero only when the commanded actuator position is different from its actual value. To compute this difference, a measurement or an estimate of the actuator position $\hat{\delta}$ is required. This estimate is then used to compute the pseudocontrol hedge as

$$v_h = \hat{h}_r(y, \delta_{\text{cmd}}) - \hat{h}_r(y, \hat{\delta}) = v - \hat{h}_r(y, \hat{\delta}) \quad (15)$$

The PCH signal is next introduced as an additional input into the reference model, forcing it to move back. If the reference model update without PCH is given by Eq. (8), then the reference model update with PCH is set to

$$y_c^{(r)} = h_{\text{rm}}(y_c, \dot{y}_c, \dots, y_c^{(r-1)}, y_c^{\text{pilot}}) - v_h \quad (16)$$

The instantaneous pseudocontrol output of the reference model that is used as an input to the linearized plant model is not changed by the use of PCH and remains

$$v_{\text{rm}} = h_{\text{rm}}(y_c, \dot{y}_c, \dots, y_c^{(r-1)}, y_c^{\text{pilot}}) \quad (17)$$

Hence, the effect of the PCH signal on the pseudocontrol is introduced only through the reference model dynamics.

By the use of the reference model of Eq. (16), Δ in Eq. (10) becomes

$$\Delta = h_r(x, \delta) - \hat{h}_r(y, \hat{\delta}) \quad (18)$$

rather than Δ given by Eq. (5) for the case with no PCH. Any modeling inaccuracies introduced in the PCH path for constructing $\hat{\delta}$ are incorporated into the overall modeling error Δ and, thus, are compensated for by the adaptive control signal v_{ad} .

Simulation Results and Discussion

A nonlinear simulation of a Boeing 747 aircraft¹⁹ is used to investigate the proposed control methodology. The aircraft model consists of six-degree-of-freedom kinematics, linearized aerodynamics, linearized propulsion, and first-order rate and position limited actuators. Although the model is approximate, it captures the relevant behavior of the aircraft reasonably well. In addition, parts of the methodology, excluding the hierarchical structure and the preengagement learning, were successfully evaluated in nonlinear and pilot-in-the-loop simulations as reported in Ref. 8. The primary (nominal) control system utilizes simple linear proportional-integral (PI) controllers, whereas the secondary control systems use adaptive NN-based control systems. If the primary channels represent an existing, certified flight control system, the secondary adaptive control channels introduce a flight safety enhancement retrofit. The controller parameters of the various channels are presented in Table 1. The relative degree of each channel is used to determine the order of reference model and the type of linear compensator to be designed. The approximate control authority of an actuator in a given primary or secondary channel determines the bandwidth of the closed-loop design and of the respective reference model.

Primary Control Channels

The primary control channels are nonadaptive and are representative of typical certified flight control systems currently in use in conventional transport aircraft. Four primary channels are designed: pitch rate control using elevator, forward speed controlled by engines, lateral specific force (sensed lateral acceleration) regulation using the rudder for turn coordination, and roll rate control by ailerons. The pitch, speed, and roll channels are designed with PI linear compensators. The relatively high bandwidth in the pitch and roll channels is specified because of the relatively high bandwidth and control surface effectiveness of the respective actuators. The speed channel is considerably slower, accounting for the lower bandwidth in engine actuation in addition to high bandwidth response in the speed channel not being required. The lateral specific force regulation channel is designed to provide turn coordination and, thus, does not require a reference model.

Pitch and roll rate channels are considered higher-priority tasks, hence, secondary control channels are designed to maintain these functions during failures while utilizing alternative actuators. Thus, the primary channels of speed control and lateral specific force regulation may be abandoned in the event that the engines (and possibly the rudder) are required to maintain these higher-priority control tasks.

Table 1 Controller parameters

Parameter	Primary channels				Secondary channels		
	q/δ_e	u/δ_{th}	a_y/δ_r	p/δ_a	q/δ_{a_s}	q/δ_{th_s}	p/δ_{th}
Relative degree	1	1	0	1	1	2	2
Closed-loop bandwidth, rad/s	3.0	0.3	1.0	3.0	3.0	1.0	1.0
Reference model bandwidth, rad/s	2.0	0.125	—	2.0	2.0	0.5	0.5
Hedging	—	—	—	—	Based on measured $\delta_{a_{\text{meas}}}$	Model-based saturations	Model-based actuator dynamics and saturations
NN gain F	—	—	—	—	750	5000	1250
NN gain σ	—	—	—	—	0.04	0.025	0.05

Secondary Control Channels

The secondary control channels chosen to illustrate the suggested control methodology include pitch rate control using symmetric ailerons, pitch rate control using propulsion, and roll rate control using asymmetric engine actuation. Additional secondary channels are envisioned in an actual implementation, which use such possible actuators as split elevator, spoilers, flaps, and more. Each of these channels is designed using the adaptive NN-based techniques described earlier. Approximate model inversion was performed assuming a linear function between the pseudocontrol signal and the respective actuator commands, hence, neglecting other aerodynamic and kinematic state-dependent effects in the aircraft equations of motions. This simplifies Eq. (4) to $\delta_{cmd} = K v$, where the fixed values for K were chosen independent of the flight condition, using the reciprocal of an average value of each actuator's effectiveness.

Simple first-order lead compensators were designed for each channel. These channels are designed to account for the lower bandwidth and authority of the respective secondary actuation systems, resulting in bandwidths that are about one- to two-thirds of the bandwidths of the primary channels, as shown in Table 1. Output feedback-based NNs are used as the adaptive elements to cancel the model inversion errors in each channel. In each control channel, an RBF NN with five neurons was used. The adaptation gains in the NN training law of Eq. (14) are given in Table 1. Signals used as inputs to the NNs in the two longitudinal channels are the measured angle of attack, pitch attitude, and rate, whereas the lateral channel NN uses roll and yaw measurements only.

PCH was implemented to account for the actuator characteristics. As shown in Table 1, three types of PCH computation methods were used in the three secondary channels. The symmetric aileron to pitch rate channel uses the measured aileron positions δ_{ameas} for hedging, that is, in Eq. (15), $\hat{\delta} = \delta_{ameas}$. In the secondary channels that use engine control, the hedging signal is computed based on a model of the propulsion system, accounting for the case where actuator position cannot be measured directly. The linear compensator design in the pitch rate channel using thrust actuation accounts for a first-order engine model lag, resulting in a relative degree-two system (as opposed to a relative degree-one system with symmetric aileron actuation). Thus, in the propulsion to pitch rate channel, only an engine position and rate saturation model is used for PCH, as is depicted in Fig. 3. The value $\Delta\hat{\delta}_{thc}$ in Fig. 3 is an estimate of the required engine throttle change based on a discrete rate approximation.

A similar design approach in the propulsion-based roll rate channel would result in a relative degree-three case: relative degree two in the asymmetric thrust to roll rate transmission plus engine lag. To avoid designing a dynamic compensator for a linearized plant with three poles at the origin, the engine dynamics were neglected in the

linear design step but accounted for by computing the PCH signal based on a model that incorporated both the engine first-order lag and its position and rate saturations, as shown in Fig. 4. Approximate saturation values and engine dynamic characteristics were used in the PCH computations, whereas the resulting hedging errors are accounted for by the NN, as described in the PCH methodology section.

Simulation Scenario

A simulation scenario is designed to examine the response of the system while under failed conditions, as well as the system response during transients. The primary actuators are failed, and the secondary channels engage with a delay that approximates the lag due to latency in failure identification or a pilot manually engaging the system. The duration of the scenario is 70 s. The pitch and roll rate commands are ramp up and hold, then ramp back down and hold signals. The magnitude of the pitch and roll rate commands is $\frac{1}{2}$ deg/s. For the first 10 s, all primary control channels are engaged and functioning correctly. At $t = 10$ s, the elevator is frozen at its current value for the remainder of the simulation. At $t = 12$ s, the secondary pitch rate control channel using symmetric aileron actuation is engaged. The aircraft's pitch and roll rate motion from this point on is controlled by the ailerons.

At $t = 30$ s, the ailerons are failed by returning their position to zero. For the next 2 s, there is no control of either the pitch or roll modes of flight. At $t = 32$ s, speed control is abandoned, and the engines assume both pitch and roll rate control tasks for the duration of the flight.

This simulation was performed at two flight conditions, sea level at 280 ft/s (0.25M) and 30,000 ft at 600 ft/s (0.6M), corresponding to a hot landing configuration and cruise at altitude for the Boeing 747. This tested the sensitivity of the adaptive design to varying flight conditions. The simulations revealed similar performance at both flight conditions, demonstrating its robustness. Only the sea-level results are presented here.

Sea-Level Simulation Results

The simulation results are presented in Figs. 5–10. A comparison between the commanded pitch rate profile Q_{pilot} and the actual aircraft response Q is depicted in Fig. 5, which also shows the output of the current hedged reference model Q_{com} determined by the active control channel. Relatively low-amplitude inputs are commanded to accommodate the low pitch authority of the symmetric ailerons and engines. There is a small transient in the response following the elevator failure, due to the delay before engaging the symmetric ailerons for pitch rate control. However, after a short transient, this secondary channel provides satisfactory pitch rate performance. Note that in the time period between $t = 12$ and 30 s, the ailerons

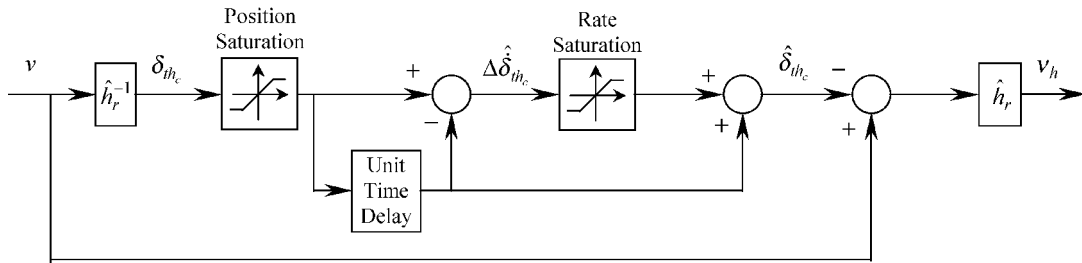


Fig. 3 PCH calculation for the thrust to pitch rate channel.

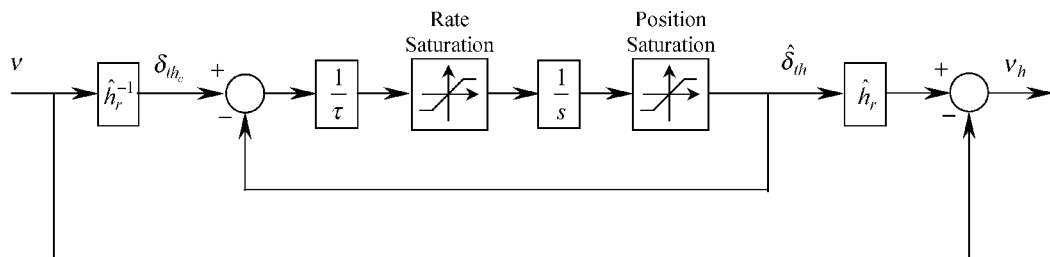


Fig. 4 PCH calculation for the thrust to roll rate channel.




Fig. 5 Pitch rate response.




Fig. 8 Time history of flight speed variation.

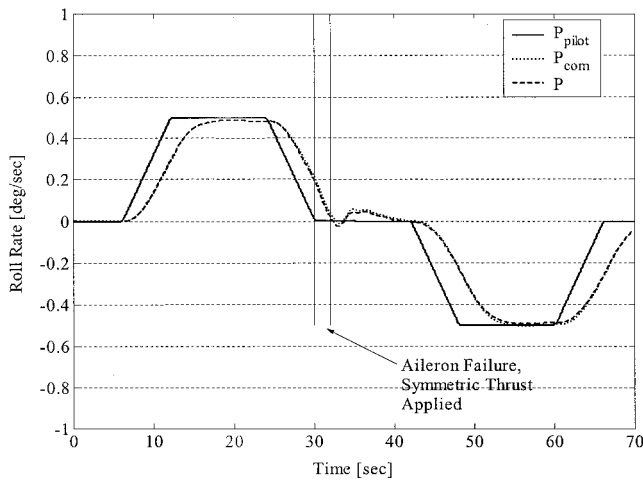


Fig. 6 Roll rate response.




Fig. 9 Inversion error adaptation in the pitch rate channel.




Fig. 7 Actuator time histories.

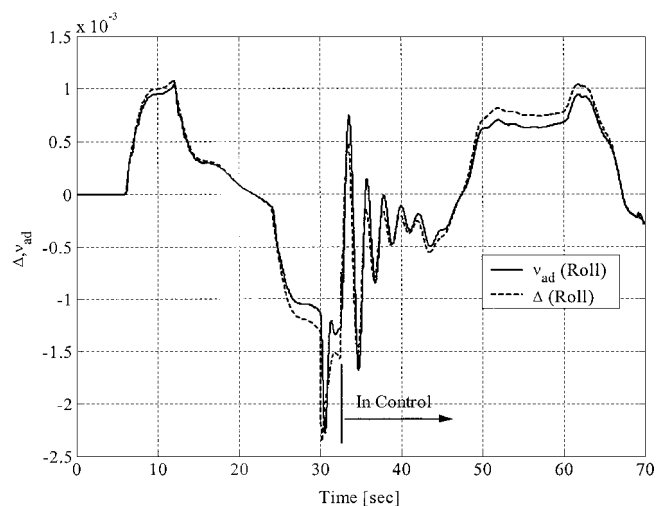


Fig. 10 Inversion error adaptation in the roll rate channel.

simultaneously perform pitch and roll rate control. The quality of the latter can be observed in Fig. 6, showing the aircraft roll rate response. Here too, low-amplitude roll rate commands are issued to account for the low roll rate authority of asymmetric actuation of engines.

The aileron failure at $t = 30$ s introduces a significant transient in both pitch and roll rate channels, caused by the ailerons' return to their neutral position. Speed control is compromised while engaging the engines for symmetric tertiary pitch rate control together with asymmetric secondary roll rate control. Despite the transient caused

by the aileron failure, the engine's controllers quickly recover. Also, satisfactory tracking is regained, in spite of a slower response with slightly reduced pitch rate tracking accuracy. In particular, note that, during this whole time period, these backup channels have to counteract the effect of the frozen elevator.

Time histories of the actuator positions are shown in Fig. 7. The first 10 s of simulation demonstrate the nominal actuator activities, whereas only the primary control channels are engaged. The differences in the four engines after $t = 32$ s create the asymmetric control




Fig. 11 Longitudinal thrust control comparison: adaptive vs nonadaptive.




Fig. 12 Lateral thrust control comparison: adaptive vs nonadaptive.

actuation needed for roll, whereas the average generates the pitch moment necessary for pitch rate tracking. In the few seconds following engagement, the engines exhibit rate saturation, which results in a significant pitch rate channel PCH, as can be seen in Fig. 5. The engines briefly position saturate at the end of the maneuver. Figure 8, which depicts the flight speed variation, demonstrates that the engines can regulate speed until it is compromised for attitude rate control. Thus, starting $t = 32$ s, the aircraft speed changes according to the pitch rate control channel command and is subject to natural damping only.

Figures 9 and 10 compare the model inversion errors Δ and the adaptive signals v_{ad} designed to cancel them in the secondary propulsion-based control channels. Proper adaptation is demonstrated both when the propulsion channels are active and when engines are used for speed control only.

To demonstrate the importance of adaptation, the simulation scenario was changed to examine propulsion-only control in both the lateral and longitudinal axes. The primary actuators (elevator and ailerons) are failed at $t = 10$ s. One simulation is performed with pitch rate commands only, whereas a second simulation is performed with only roll rate commands. Figure 11 shows the aircraft response with and without NN adaptation in the longitudinal channel, using the same linear lead compensator. Following the failure of the elevator, this secondary channel demonstrates that little control is possible without adaptation. The steady-state error observed with the adaptive controller can be easily alleviated using additional integral action. Figure 12 demonstrates the aircraft lateral response after failure of the ailerons, with and without NN adaptation. Here, the response with the nonadaptive controller produces instability. This clearly demonstrates the advantage of NN-based adaptive control.

Trim Variations

As noted earlier, the approach detailed in this paper is fundamentally nonlinear. A significant aspect not addressed by our numerical results is the variation in trim that can accompany a failure. Although we have not demonstrated this point here, the ability of the adaptive approach to handle trim variations has been extensively evaluated both in piloted simulations and in flight tests in Refs. 8, 9, and 20, which also included a powered approach scenario. Effectiveness in the presence of neglected trim variations has also been addressed within the context of autopilot design for guided munitions in Refs. 21 and 22, wherein the trim tables were completely removed in the design.

Conclusions

This paper presents an NN-based adaptive flight control design methodology and architecture for actuator failure accommodation. Safe flight is maintained by incorporating standby control channels that utilize analytically redundant secondary actuators to continuously maintain satisfactory, or at least stable, operation. The overall dynamic effect of known actuator failures of unknown type and magnitude is compensated by the adaptive NN-based element of the secondary control channels. Each of these secondary control channels is designed for the primary control task, accounting for the dynamic characteristics of the channel, the possible degraded authority of the secondary actuator, and the limited achievable performance. Because the design of the secondary channels does not depend on the architecture of the normal operation controller, the proposed methodology is ideal for safety retrofit of any digital fly-by-wire flight control system.

Acknowledgment

Research sponsored by NASA Grant NAG-1-2243.

References

1. "Aircraft Accident Report, United Airlines Flight 232," PB90-910406, NTSB/AAR-90/06, National Transportation Safety Board, Washington, DC, Nov. 1990.
2. Ostrof, A. J., "Techniques for Accommodating Control Effector Failures on a Mildly Statically Unstable Airplane," *American Control Conference*, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 1985, pp. 903–906.
3. Gao, Z., and Antsaklis, P. J., "Stability of the Pseudo-Inverse Method for Reconfigurable Control Systems," *International Journal of Control*, Vol. 53, No. 3, 1991, pp. 717–729.
4. Ochi, Y., and Kanai, K., "Application of Restructurable Flight Control System to Large Transport Aircraft," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 2, 1995, pp. 365–370.
5. Pachter, M., Chandler, P. R., and Mears, M., "Reconfigurable Tracking Control with Saturation," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 5, 1995, pp. 1016–1022.
6. Burken, J. J., Lu, P., Wu, Z., and Bahm, C., "Two Reconfigurable Flight Control Design Methods: Robust Servomechanism and Control Allocation," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 3, 2001, pp. 482–493.
7. Siwakosit, W., and Hess, R. A., "Multi-Input/Multi-Output Reconfigurable Flight Control Design," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 6, 2001, pp. 1079–1088.
8. Rysdyk, R., Leonhardt, B., and Calise, A. J., "Development of an Intelligent Flight Propulsion Control System: Nonlinear Adaptive Control," AIAA Paper 2000-3943, Aug. 2000.
9. Calise, A. J., Lee, S., and Sharma, M., "Development of a Reconfigurable Flight Control Law for the X-36 Tailless Fighter Aircraft," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 5, 2001, pp. 896–902.
10. Johnson, E., Calise, A. J., and El-Shirbiny, H., "Feedback Linearization with Neural Network Augmentation Applied to X-33 Attitude Control," AIAA Paper 2000-4157, Aug. 2000.
11. Bull, J., Mah, R., Hardy, G., Sullivan, B., Jones, J., Williams, D., Soukup, P., and Winters, J., "Piloted Simulation Tests of Propulsion Control as Backup to Loss of Primary Flight Controls for a B747-400 Jet Transport," NASA TM 112191, April 1997.
12. Burken, J. J., Burcham, F. W., Jr., Maine, T. A., Feather, J., Goldthorpe, S., and Kahler, J. A., "Flight Test of a Propulsion Based Emergency Control System on the MD-11 Airplane with Emphasis on the Lateral Axis," NASA TM 4746, July 1996.
13. Johnson, E., and Calise, A. J., "Neural Network Adaptive Control of Systems with Input Saturation," *American Controls Conference*, Inst. of

Electrical and Electronics Engineers, Piscataway, NJ, 2001, pp. 3527–3532.

¹⁴Johnson, E., “Limited Authority Adaptive Flight Control,” Ph.D. Dissertation, School of Aerospace Engineering, Georgia Inst. of Technology, Atlanta, Nov. 2000.

¹⁵Calise, A. J., Hovakimyan, N., and Idan, M., “Adaptive Output Feedback Control of Nonlinear Systems Using Neural Networks,” *Automatica*, Vol. 37, No. 8, 2001, pp. 1201–1211.

¹⁶Isodori, A., *Nonlinear Control Systems*, Springer-Verlag, Berlin, 1989, pp. 147–161.

¹⁷Funahashi, K., “On the Approximate Realization of Continuous Mappings by Neural Networks,” *Neural Networks*, Vol. 2, No. 3, 1989, pp. 183–192.

¹⁸Hornik, K., Stinchcombe, M., and White, H., “Multilayer Feedforward

Networks are Universal Approximators,” *Neural Networks*, Vol. 2, No. 5, 1989, pp. 359–366.

¹⁹Nelson, R., *Flight Stability and Automatic Control*, McGraw-Hill, New York, 1989, p. 260.

²⁰Brinker, J. S., and Wise, K. A., “Flight Testing of Reconfigurable Control Law on the X-36 Tailless Aircraft,” *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 5, 2001, pp. 903–909.

²¹Calise, A. J., Sharma, M., and Corban, J. E., “Adaptive Autopilot Design for Guided Munitions,” *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 5, 2000, pp. 837–843.

²²Sharma, M., Calise, A. J., and Corban, J. E., “Application of an Adaptive Autopilot Design to a Family of Guided Munitions,” AIAA Paper 2000-3969, Aug. 2000.