JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 20, No. 4, July–August 1997

Design of Deep Stall Protection for the C-17A

Obi I. Iloputaife*
McDonnell Douglas Aerospace, Long Beach, California 90807-4418

The design of the C-17 angle-of-attack limiter system is presented. The C-17, because of its T-tail, has a locked-in deep stall potential in some aft center of gravity configurations. The angle-of-attack system uses alpha measurements from three pairs of vanes for precise angle-of-attack control and to prevent excursion into the deep stall region. The vanes, because of their locations, are sensitive to sideslip angle excursions in varying degrees, therefore, requiring sideslip angle information for proper failure monitoring. A method of synthesizing sideslip and angle of attack using the information from the vanes was developed. With the synthesized angle of attack and sideslip, it was possible to provide a minimum of fail-operational-fail-operational capability with tight failure tolerance. After over 2000 stall tests and assaults to alpha limit, the pilot community has gained the confidence that the system will protect the C-17 from entering the locked-in deep stall region.

Nomenclature

C_m	= pitching moment coefficient
i_H	= horizontal stabilizer incidence angle, deg
P	= aircraft roll rate, deg/s
$\dot{V}_T$	= rate of change of true airspeed, kn/s
$\dot{V}_{T_{\text{disengage}}}$	= attack limiter system (ALS) disengage deceleration limit, kn/s
$\dot{V}_{T_{\text{engage}}}$	= ALS engage deceleration limit, kn/s
$\alpha_{1,\ldots,6}$	= local vane angle of attack, deg
α_F	= fuselage angle of attack, deg
α_{msm}	= angle of attack at minimum static margin of −0.08, deg
α_{stall}	= stall angle of attack, deg
α_{valid}	= computed alpha validity discrete, AOAOK
β	= sideslip angle, deg
β_{valid}	= computed sideslip angle validity discrete, BETAOK
δ_e	= elevator deflection, deg

Introduction

THE C-17A is a high-performance, military cargo airplane. Its basic mission is to deliver outsize cargo routinely to austere airfields around the world. The C-17A can operate routinely out of unpaved runways in forward battle areas, land, and off-load all of its cargo under combat conditions. It can also deliver necessary cargo using its low-altitude parachute extraction system capability.

The flight control system began as a mechanical system with a limited authority stability and command augmentation system. Evaluation of wind-tunnel data revealed a potential inability to comply with a C-17 air vehicle specification (AVS) paragraph that states, "The airplane shall be readily recoverable from all attainable attitudes and motions."

The C-17 is a T-tail airplane and, like many airplanes with a T-tail configuration, has a locked-in deep stall potential. As is seen in Fig. 1, it is extremely difficult to recover from a locked-in deep stall, without special recovery devices like parachutes, using the available pitch control authority. One way to ensure that the aircraft does not attain attitudes and motions where it is not readily recoverable is to prevent maneuvering beyond the minimum static pitching moment angle of attack, shown as static pitching moment margin in Fig. 2. A study investigated various ways of limiting the aircraft angle of attack to prevent excursions beyond that required for maintaining the minimum static margin. The two leading candidates were an angle of attack limiter system (ALS) and a stick pusher.

Presented as Paper 96-3784 at the AIAA Guidance, Navigation, and Control Conference, San Diego, CA, July 29–31, 1996; received Oct. 18, 1996; revision received March 8, 1997; accepted for publication March 9, 1997.

*Senior Principal Engineer, C-17 Avionics/Flight Controls Integrated Product Team. Senior Member AIAA.

The study resulted in the selection of an extremely reliable ALS, instead of a stick pusher, to provide the required protection from deep stall. The stick pusher lacked the required reliability and has a potential for single point failure. A hard-over failure of the stick pusher would require the pilot to carry heavy and untrimmed forces until landing, in violation of the C-17 AVS.

The ALS employs a full authority quadruplex digital fly-by-wire system. This mechanization provides the necessary safety margin, with a minimum requirement of fail-operational-fail-operational for redundancy, and 10^{-9} probability of erroneous angle-of-attack limiting.

This paper discusses the development and validation of the ALS and includes remarks on control law and signal management design.

Angle-of-Attack Measurement

Because vane type sensors provide better accuracy at high angles of attack, six alpha vanes, three on each side of the airplane, are the primary sources of angle-of-attack measurement for the C-17. The vane arrangement involves three groups of left–right pairs: upper, middle, and lower pairs. Vanes in each group are on the same fuselage reference location in the X–Z plane. True angle-of-attack calculation and failure monitoring require one vane pair and one additional vane. The three vane pairs are adequate to meet the accuracy and redundancy requirements. The angle-of-attack vane arrangements and numbering are shown in Fig. 3.

Each vane senses local angle of attack. The ALS then corrects these measurements to true aircraft reference angle of attack before use and dissemination to other functions. However, the biggest challenge was in providing failure monitoring for the alpha sensors to prevent erroneous limiting of angle of attack. Vane measurement of local angle of attack is highly dependent on sideslip angle, and because there are no sideslip sensors in the C-17 flight control system, it made the failure monitoring task particularly difficult.

To prevent nuisance vane failures, an initial failure detection algorithm used a detection threshold of 20 deg. Of course, this implied a propagation of up to 20 deg of alpha error in the ALS. Fortunately, validation of the high-alpha aerodynamics occurred without the ALS during the initial flight test phase. This gave adequate time to observe the in-flight behavior of these vanes and develop a better input signal management (ISM) algorithm.

Analysis of various ISM methods revealed that using sideslip to condition the vane measurements before monitoring provided the only acceptable solution. Sideslip in the monitoring equation allows a reduction of the detection threshold from 20 to 2.5 deg. The lower detection threshold significantly reduces the risk of error propagation without any increase in nuisance disconnects.

Computation of Sideslip Angle

There are various methods of obtaining sideslip, such as installation of sideslip sensors and estimation of sideslip from lateral

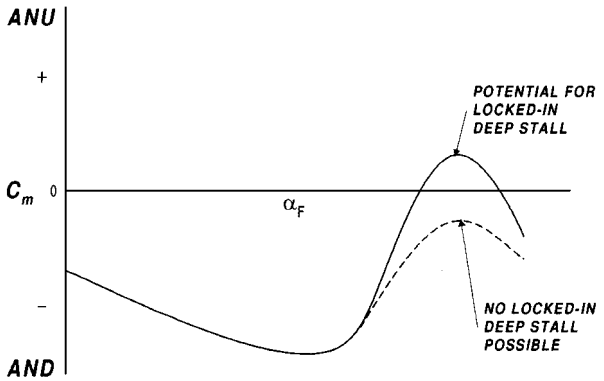

Fig. 1 Pitching moment curve showing deep-stall potential; ANU is aircraft nose up and AND is aircraft nose down. Note: maximum AND authority.

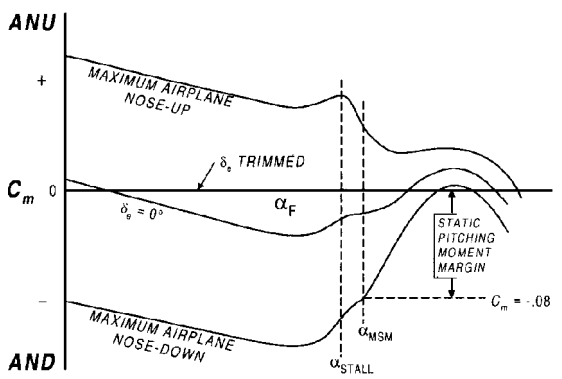

Fig. 2 Available aircraft pitch authority. Note: i_H = constant.

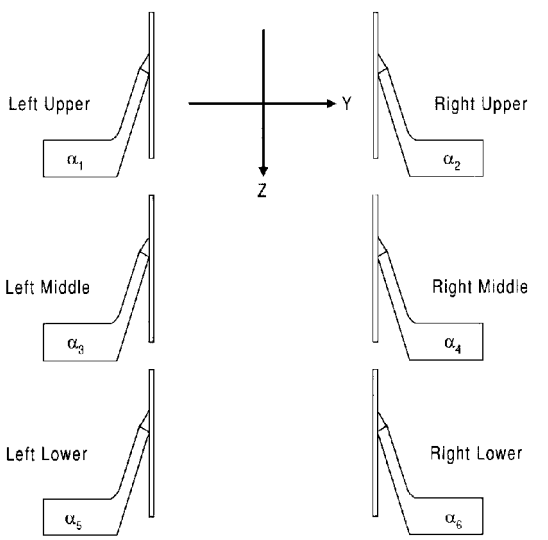

Fig. 3 Angle-of-attack sensor locations.

acceleration. Computation from lateral acceleration is inherently inaccurate in the presence of winds and requires extensive gain scheduling. Also, addition of sideslip sensors on the airplane, at this stage in the program, was not a cost-effective alternative.

After observing how the vanes behaved during the initial high-alpha flight testing, it was apparent that generation of sideslip was possible by exploiting the relationship between the vanes. Analysis of wind-tunnel and flight test data revealed that the angle-of-attack vanes have the following characteristics.

1) Without atmospheric disturbances, sideslip, or failures, vanes in the same group behave alike.

2) In the presence of sideslip (steady), vanes on the same side have a fixed relationship.

3) Vanes in the same pair have a fixed relationship with sideslip and roll rate.

4) Each vane pair has unique regions of angle of attack in which it is insensitive to sideslip.

5) Vanes on the same side behave closer to each other in the presence of atmospheric disturbances than vanes on opposite sides.

It is possible to generate sideslip by exploiting these behavior patterns, using a combination that employs vanes from the same side and the three left–right vane pairs.

Vanes on the same side have a fixed relationship with sideslip resulting in six sideslip estimates after combining as follows:

$$\beta_{13} = (\alpha_1 - \alpha_3)k_1 \tag{1}$$

$$\beta_{15} = (\alpha_1 - \alpha_5)k_2 \tag{2}$$

$$\beta_{35} = (\alpha_3 - \alpha_5)k_3 \tag{3}$$

$$\beta_{24} = (\alpha_2 - \alpha_4)k_4 \tag{4}$$

$$\beta_{26} = (\alpha_2 - \alpha_6)k_5 \tag{5}$$

$$\beta_{46} = (\alpha_4 - \alpha_6)k_6 \tag{6}$$

Also, sideslip can be extracted from vanes on opposite sides (vane pairs) as follows:

$$\beta_{12} = (\alpha_1 - \alpha_2 + K_p P)k_7 \tag{7}$$

$$\beta_{34} = (\alpha_3 - \alpha_4 + K_p P)k_8 \tag{8}$$

$$\beta_{56} = (\alpha_5 - \alpha_6 + K_p P)k_9 \tag{9}$$

where k_1–k_9 are functions of angle of attack, K_p is a function of angle of attack and velocity, and P is aircraft roll rate. Each vane pair has an angle-of-attack region in which it is insensitive to sideslip. These regions span the angle-of-attack envelope and sometimes overlap, thereby resulting in areas where only one of the three estimated sideslip angles in Eqs. (7–9) is available from left–right vane pair combinations. In these regions of angle of attack, the sideslip estimated from the left–right vane combination must be excluded because it is indeterminate. Indeterminacy is treated as a temporary invalidity of that sideslip estimate and is used to exclude it from consideration in the selection process leading to the final sideslip estimate.

Failure of one vane automatically eliminates three sideslip angle estimates. However, without the region of insensitivity, it requires loss of four vanes for complete loss of sideslip. The challenge is in synthesizing the correct sideslip from the nine estimated values in the presence of undetected failures.

Figure 4 shows the sideslip computation schematic. Each initial selection algorithm includes all known relevant failure information

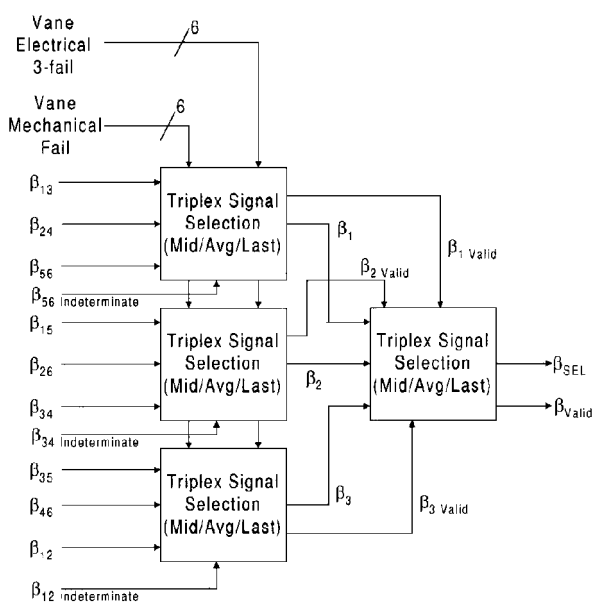

Fig. 4 Sideslip computation schematic.

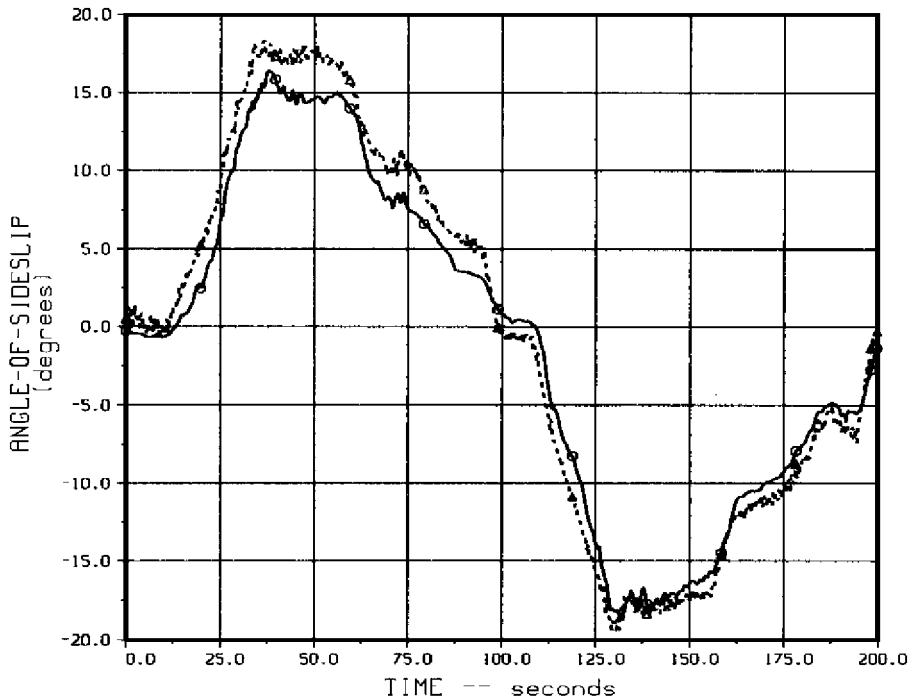

Fig. 5 Comparison of computed and measured sideslip: ∘, computed and △, measured.

about each vane from the last iteration. The failure information is used to narrow the selection process and includes electrical failures (each vane is essentially quadruplex electrical), mechanical failures, and the indeterminate sideslip region. Allocation of estimated sideslip angles to each selection group is such that each vane has only one representation as follows: group 1—β_{13}, β_{24}, and β_{56}; group 2—β_{15}, β_{26}, and β_{34}; and group 3—β_{35}, β_{46}, and β_{12}.

Allocating the nine estimated sideslip angles in this manner trivializes the selection process because each vane failure results in sideslip angles (there are three angles associated with each vane) with different characteristics than other angles in the group. Low sideslip values could be due to stuck vane pairs or from regions where a particular vane pair is insensitive to sideslip. Most failures manifest as higher values of sideslip because a typical vane failure is either stuck or running away, therefore, increasing the value of the difference. Each group uses a triplex signal selection process (middle/average/last). This process involves selecting the middle value when there are no declared failures or indeterminate region on which the particular β_{ij} input depends. Each sideslip angle estimate is valid when its vanes are operating normally and, in addition, for β_{12}, β_{34}, and β_{56}, when they are not in the indeterminate region. When there are only two remaining valid sideslip estimates, the average of the two values would be the output and, finally, the last good value when there are two failures. The output of each group includes the selected value and a validity flag. This validity flag is true so long as there is a valid sideslip estimate in the group. Additional processing of the outputs of the three groups, using the same middle/average/last selection criteria, produces the final sideslip angle. The last selection stage eliminates any errors that might propagate from the initial selection process. The final sideslip angle is complemented with inertial sideslip information from a quadruple redundant inertial reference sensor set to remove disturbance.

Figure 5 is a graph of actual C-17 flight test data showing a comparison of measured (nose boom mounted sensor) and computed (β_{sel}) sideslip angle. The results show a worst-case error of 3 deg during the dynamic portions of the maneuver. This degree of accuracy is adequate for failure monitoring.

Failure Detection and Isolation

Signal management of the alpha vanes includes electrical and mechanical failure detection and isolation. Electrical failure detection is performed on each individual vane before use in the mechanical failure detection algorithm. Each vane is electrically

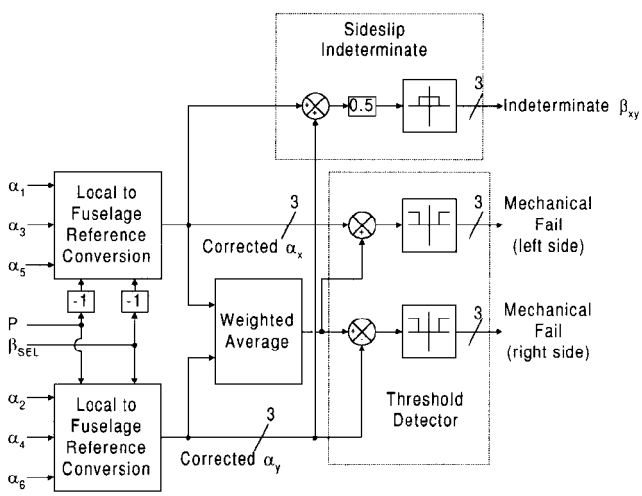

Fig. 6 Alpha vane mechanical failure detection schematic.

quadruple redundant; therefore, initial signal selection of the four electrical inputs detects failures and generates a selected output for each vane. The selected outputs are $\alpha_1-\alpha_6$ of Fig. 3. These outputs are then checked for mechanical failures. Mechanical failure detection requires reducing the output of the vanes to a common reference. Reduction to a common fuselage aircraft reference involves conditioning of the six local angle-of-attack signals with estimated sideslip angle, true aircraft angle of attack, and aircraft roll rate. The chosen reference provides a point on the airplane where all of the vanes read the same angle in the absence of failures. Figure 6 is a schematic of the mechanical failure detection algorithm.

Mechanical failure detection includes rate and position threshold checks. The position threshold check involves comparison of each individual corrected alpha with a weighted average (filtered). When a corrected alpha exceeds the detection threshold, with some persistence, mechanical failure is declared. However, that failure is not latched until the next failure occurs. The inputs to the weighted average algorithm are the six corrected alpha signals, each limited to ± 10 deg of the filtered weighted average. An average of the good vane signals is then computed and heavily filtered for noise.

This signal management mechanization allows a minimum fail-operational-fail-operational redundancy with three vane pairs. On

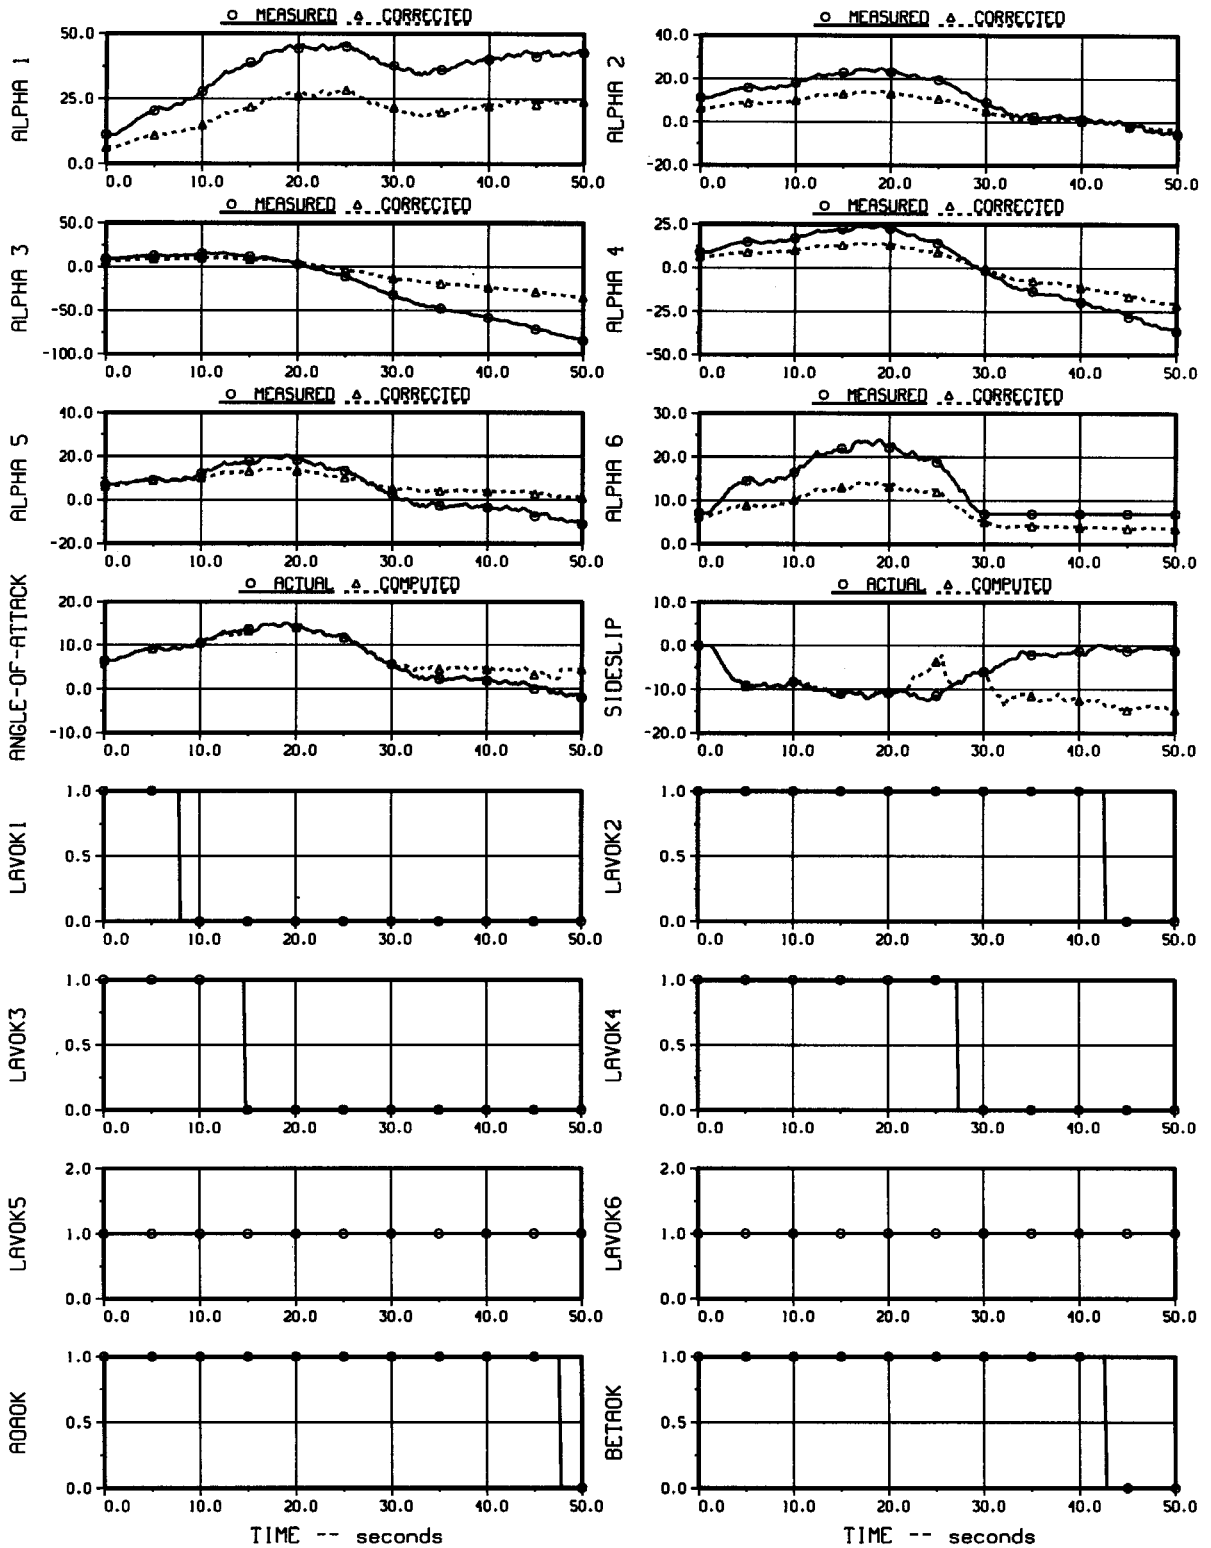

Fig. 7 Simulation results of failure detection logic.

failure of the last vane pair, the pilot is provided enough time to recover the airplane to a more benign attitude prior to ALS deactivation. After detection and isolation of the first two vane pair failures, an ALS fail-op annunciation advises the crew of the condition that one more failure could result in complete ALS shutdown. On detection of the third vane pair failure, the annunciation reverts to alpha limit inop; following this detection, synthesized alpha from inertial and air data measurements provide an additional 5 s of ALS operation, before a complete shutdown of the ALS function.

Failure detection depends primarily on sideslip angle computation. As mechanized, seven of the nine sideslip angle estimates fail with failure of three vanes. However, processing continues until a fourth vane failure occurs. The fourth vane failure will cause loss of the monitoring function because it cannot be discerned which of the two remaining sideslip angle estimates is correct. However, a large variation in the two remaining sideslip angle estimates will cause one of the corrected alpha to trip the detection threshold and correctly shut off angle-of-attack calculation. Although the fourth failure can be detected, it is not possible to isolate to the failed vane. However, this is not fatal because all that is required is to safely shut down the ALS.

A multiple failure scenario that involves detection and isolation of failures on α_1, α_2, and α_3, leaves β_{46} and β_{56} as the only surviving

sideslip estimates. If α_6 fails, then β_{46} and β_{56} would yield erroneous values, and the selection algorithm would simply produce the average of the two incorrect sideslip angles. However, applying this erroneous sideslip angle to the three remaining vanes will produce error in at least one corrected angle of attack significant enough to fail the detection threshold. This is possible because sideslip angle is directly proportional to the angle-of-attack difference between the vanes. A large difference between α_4 and α_6 results in a large sideslip angle (β_{46}), which, when applied in the conversion algorithm of Fig. 6, could result in up to 20 deg difference in angle of attack (note that the purpose of the conversion process is to reduce the vane outputs to a common reference). However, although it is not possible to isolate the failure, early detection makes a safe shutdown of the ALS possible.

Implementation of the signal management algorithm in the C-17 nonreal-time simulation provides the basis for the analysis results presented in Fig. 7. Extensive analysis confirmed the ability of the algorithm to detect and isolate all types and combinations of failures. The simulation results in Fig. 7 are for vane failures in a maneuvering flight consisting of pitch control input, maximum rudder pedal input, and moderate von Karman turbulence.

The simulation involves applying the first, second, third, and fourth failures to the vanes in the following order: left upper vane (alpha 1) slow over (increasing alpha) starting at 1 s, left middle vane (alpha 3) slow over (decreasing alpha) starting at 10 s, right middle vane (alpha 4) slow over (decreasing alpha) starting at 20 s, and right lower vane (alpha 6) freezes at 30 s.

The first failure detection and isolation occurred at approximately 8 s, indicated by alpha vane 1 validity discrete (LAVOK1) going low. The second (LAVOK3) and third (LAVOK4) failure detection occurred at 15 and 27 s, respectively. The freezing of vane 6, the fourth failure in the sequence, caused erroneous isolation of the failure to vane 2 at 42 s. However, the purpose of the final failure detection was served and the ALS properly shut down at 47 s, indicated by computed angle-of-attack validity discrete (AOAOK) dropping to zero.

The results show that computed angle of attack tracks actual aircraft angle of attack until injection of the fourth failure 30 s into the simulation. Inasmuch as vane 6 is a member of the remaining vane pair, injecting the failure into it creates up to 2.5 deg of error in the computed angle of attack.

Computed sideslip angle will track actual aircraft sideslip angle with the first two failures having no effect on accuracy. During the time the third failure remains undetected, estimated sideslip deviates from actual. Tracking resumes once detection and isolation of the third failure occur. Upon the fourth failure, a resulting large deviation in estimated sideslip angle causes an error large enough to trigger detection of the final failure. Estimated sideslip angle becomes invalid when the fourth failure is detected [computed sideslip angle validity discrete (BETAOK) low at 42 s with LAVOK2] because the logic forces it to declare invalidity when there are fewer than three vanes left.

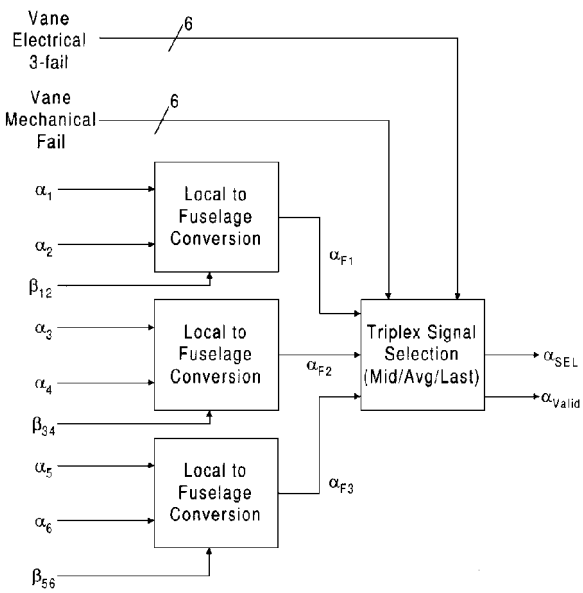

Fig. 8 Angle-of-attack computation schematic.

Selection of True Angle of Attack

The true angle-of-attack computation algorithm involves selection of middle/average/last value from the three processed vane pair data. Although computation of true angle of attack only requires one vane pair (left and right), failure detection considerations dictate a minimum of three vanes (one pair and one) for proper ALS operation. Prior to selection, each local angle of attack is corrected to a common fuselage reference using the sideslip angle associated with its relational pair, i.e., β_{12} for upper pair, β_{34} for middle pair, and β_{56} for lower pair. Aircraft roll rate, pitch rate, and other configuration

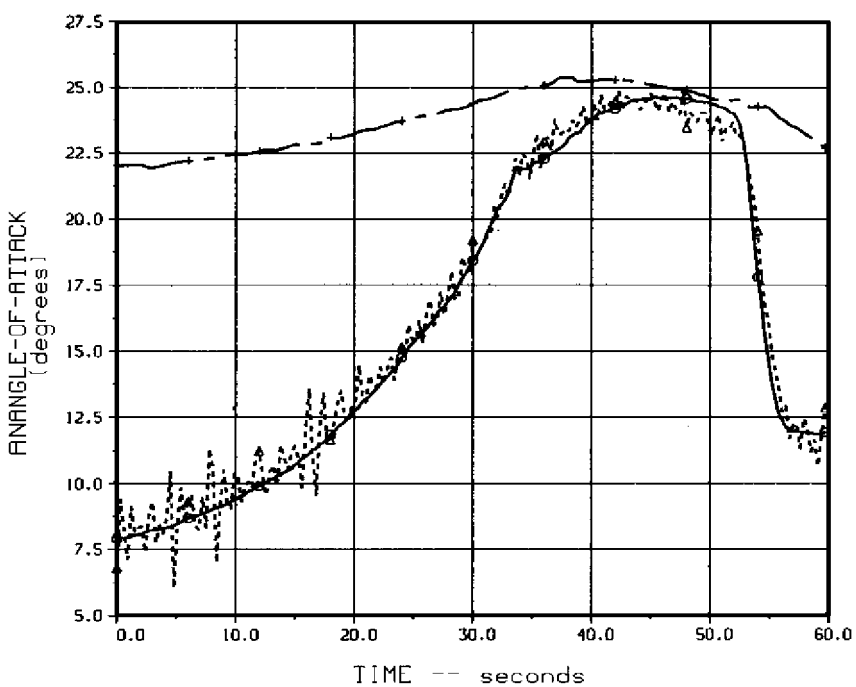

Fig. 9 Flight test comparison of computed and measured AOA: ∘, computed; △ measured; and +, stall.

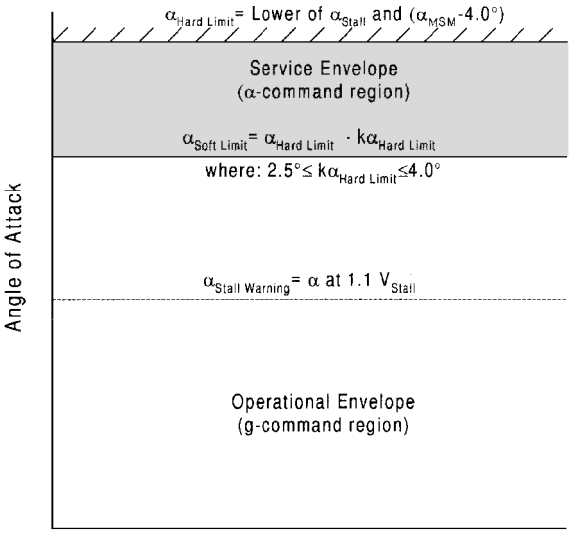

Fig. 10 Illustration of ALS operational envelope.

information are included in the final computation. The simplified schematic is shown in Fig. 8. The selected signal (α_{sel}) is complemented with inertial angle of attack to reduce noise.

Figure 9 shows a comparison of computed and actual angle of attack from C-17 flight test. Stall angle is also superimposed for reference. Computed alpha is α_{sel}, complemented with inertial alpha.

Angle-of-Attack Limiter Control Laws

Normally, the C-17 pitch axis control is in the *g*-command configuration. However, to provide maximum deep stall protection, the control system reconfigures to an alpha-command configuration when the angle-of-attack limiter system activates. Activation occurs as angle of attack approaches the α-command region defined in Fig. 10, or during sustained periods of high-airspeed bleed-off rate. Figure 10 shows the splitting of the C-17 angle-of-attack envelope between the *g*-command region (normal operational envelope) and the α-command region (service envelope). This division minimizes ALS interference with the handling qualities of the aircraft inside the operational envelope. However, at extreme attitudes with rapid and abnormal deceleration, ALS activation may occur within the operational envelope to preclude overshoot into deep stall.

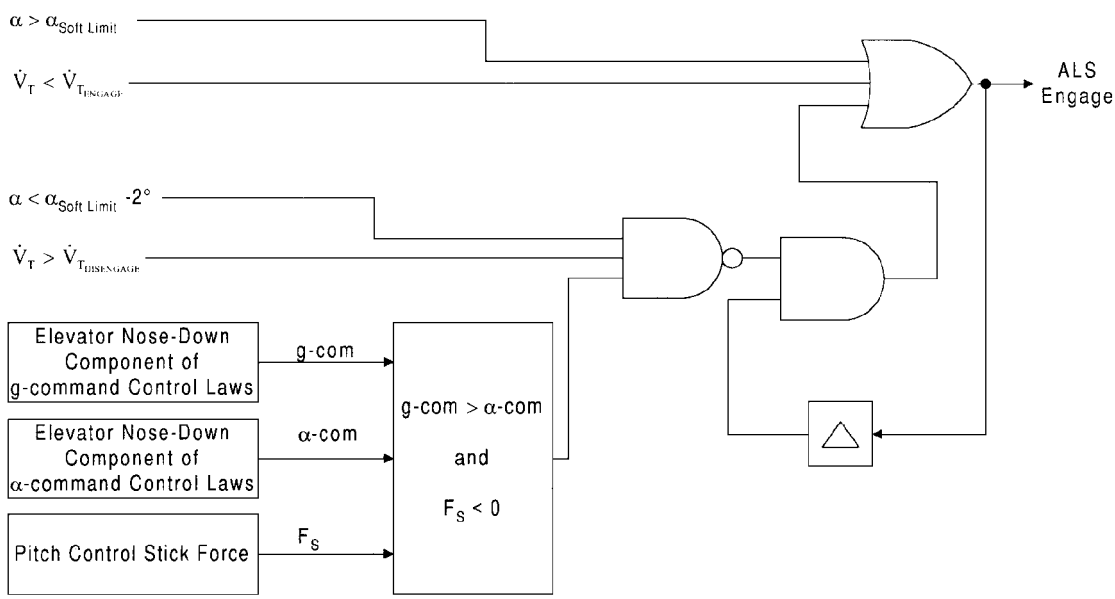

Fig. 11 ALS switching logic.

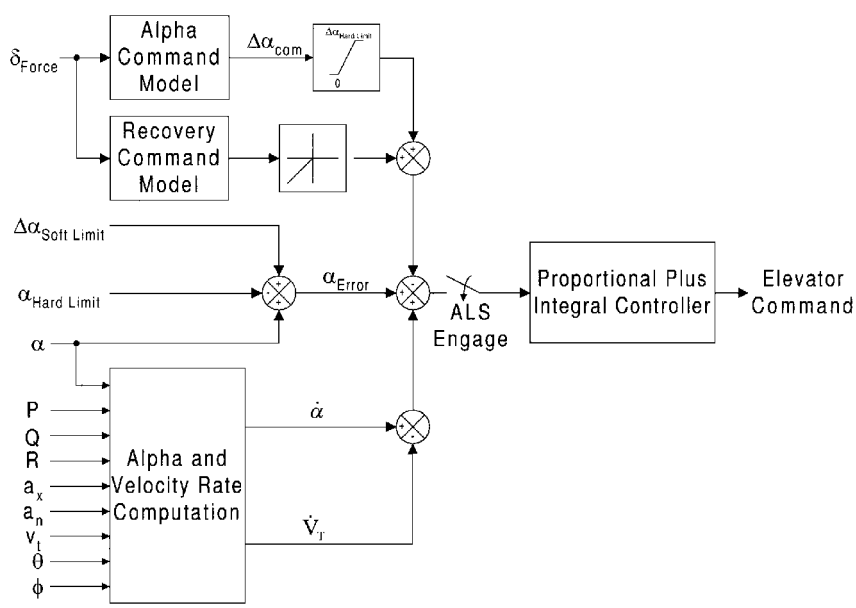

Fig. 12 ALS control law schematic.

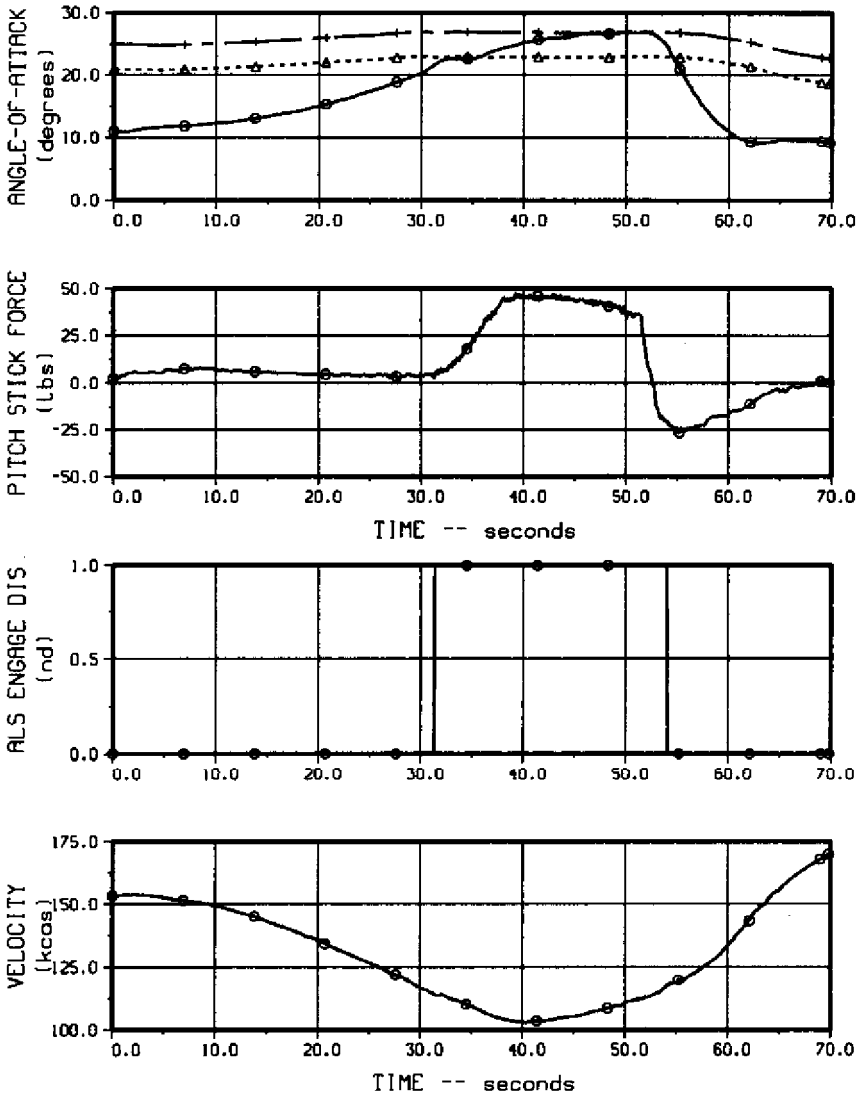

Fig. 13 Flight test result from a 1-*g* stall: ∘, alpha; △, soft limit; and +, hard limit.

Deep stall protection starts with the onset of stall warning characterized by an aural "stall stall" and a high-frequency stick shaker at $\alpha_{\text{stall warning}}$. The stall warning envelope delineates the limit of the operational envelope and stall warning onset signals the pilot to initiate recovery of the airplane. However, if the aircraft angle of attack continues to increase (either by pilot input or by decrease in speed), ALS activates when alpha reaches $\alpha_{\text{soft limit}}$. At this point, ALS flashes on the heads-up display and alpha stops increasing. At the soft limit, the pilot must apply back pressure on the stick to command additional angle of attack, up to the hard limit ($\alpha_{\text{hard limit}}$). Hard limit is the lower of α_{stall} and the angle of attack that provides a minimum 4-deg alpha margin from the alpha at which $C_m = -0.08$ (Fig. 2).

Disengagement of the ALS after recovery in the operational envelope became a major issue during development. Satisfactory disengagement requires that any transient be benign and not lead to a secondary stall. The resulting logic is shown in Fig. 11.

Once ALS is engaged, either by alpha exceeding soft limit ($\alpha > \alpha_{\text{soft limit}}$) or deceleration rate greater than some preprogrammed value ($\dot{V}_T < \dot{V}_{T_{\text{engage}}}$), it latches so long as the pilot is applying back pressure on the control stick (Fig. 11). The deceleration rate at which ALS engages is a function of calibrated airspeed and varies from 6 kn/s at 100 kn to 15 kn/s at 350 kn. When the pilot relaxes back pressure or pushes forward on the stick, the system compares the elevator commands from the g-command and the α-command control laws to determine which is providing the most recovery moment (nose down elevator). When the g-command feedback recovery component becomes greater than the α-command recovery component, the ALS is unlatched and ready for disengagement. Actual disengagement occurs if alpha is 2 deg less than the soft limit and deceleration rate is less than a preprogrammed value, which is also a function of speed.

Using the recovery component to determine an appropriate disengagement point prevents possible occurrence of secondary stalls by ensuring that, when switching out of ALS occurs, any transient would be in the recovery sense.

Figure 12 shows the ALS control law schematic. It is a simple alpha-command system with rate of change of angle of attack ($\dot{\alpha}$) for short-term damping and longitudinal acceleration ($\dot{V}_T$) for phugoid damping. The alpha-command model allows commands up to hard limit with a very low alpha per stick force ratio. The recovery command model has approximately 10 times the alpha command model gain to facilitate recovery from stall. The ALS is integrated with the normal pitch control system and is operational at all times.

Figure 13 shows a typical flight test result for a 1-g stall maneuver. The time traces show the phenomenon inherent in the ALS: a temporary stop at alpha soft limit (indicated by ALS engage discrete going high at 31 s) and the additional alpha response to pilot control input up to the hard limit thereafter. Application of additional force on the stick after attaining the hard limit would not result in additional alpha response.

Validation of the C-17 ALS in both flight test and pilot-in-the-loop simulation involved combinations of control inputs representative and, sometimes, more severe than those specified in MIL-S-83691 (Ref. 1). These tests involved more than 2000 stall points, without adverse incidents.

Conclusion

This paper has presented the very complex systems and methods required to provide deep stall protection for the C-17 military transport airplane. It shows that a minimum of fail-operational-fail-operational performance can be provided with three vane pairs without employing sideslip sensors, even though vane performance monitoring requires accurate sideslip angle information.

Sideslip angle is synthesized from the observed behavior of the vanes with respect to each other. The accuracy obtained from this synthesized sideslip angle is adequate to provide full failure detection and isolation up to the third failure and allows for a safe recovery and shut down after the fourth failure.

Acknowledgments

Special thanks to pilots John Burns and Donald Brown for their assistance and patience throughout the development of the alpha limiter system. I also extend my sincere appreciation to Timothy Dalhstrom for his advice and inputs in shaping the control system philosophy, Gregory Hoffman of Lockheed Martin Corporation for his work in the signal management algorithm, and all other engineers who helped make the system possible.

Reference

[1] "Stall/Post-Stall/Spin Flight Test Demonstration Requirements for Airplanes," MIL-STD-83691, March 1971.