Pilot's Response to Stability Augmentation System Failures and Implications for Design

Walter A. Johnson* and David H. Weir†
Systems Technology Inc., Hawthorne, Calif.

The degree to which a pilot can successfully and reliably retain control following a stability augmentation system failure is a key factor in the system design philosophy. Results of a fixed-base simulation of multiple-loop lateral/directional control of a conventional airframe with an abrupt yaw damper failure are presented. They indicate that the postfailure performance is improved if the difference in controlled element dynamics at failure is reduced. One way to achieve this is to utilize a less than optimum prefailure augmentation system so the pilot is required to be in the loop at all times, rather than to merely monitor system activity prior to failure. A model for the pilot's dynamic response is presented which accounts for his behavior during and following failure.

Nomenclature

a	=	inverse time constant
K	=	gain
r	=	yaw rate
s	=	Laplace variable
Y_c	=	transfer function for controlled element
Y_p	=	pilot model describing function
β_g	=	side gust input
δ_a	=	aileron deflection
δ_r	=	rudder deflection
φ	=	bank angle
ζ_d	=	damping ratio of dutch roll mode
ζ_φ	=	damping ratio in bank angle to aileron numerator
ω_φ/ω_d	=	ratio of numerator frequency to denominator frequency for dutch roll mode in bank angle to aileron transfer function
1	=	prefailure
2	=	postfailure
ϵ	=	error
SAS	=	stability augmentation system

Introduction

STUDIES of manual control over the past two decades have been concerned primarily with the pilot's role as an active controller. Recently, however, automation of flight control has expanded to the extent that in many instances the pilot is more of a manager who monitors flight control system activity. With this new role go the problems of a manager, and one of the most significant problems of any manager is what to do when things go wrong, for example, when a flight control system failure results in a transition in the effective controlled element dynamics. Depending on the redundancy of the system, this managerial role may require one of several types of action following a flight control system failure. For example, with a conventional, single channel, electrohydraulic, stability augmenter, about 80% of the failures result in a hardover control input. This requires the pilot to retain control in the presence of the hardover, while adapting to the new dynamics or taking corrective measures.

The failure control problem can be avoided by use of fail-operational systems which merely require the pilot to recognize the first failure and be prepared for a second. An attractive alternative in some cases is to use the pilot's control capability, plus reduced system redundancy, to achieve a design compromise with its attendant savings in weight, cost, etc. To use such a mix of manual and automatic functions, the pilot's capability must be understood and the effective vehicle dynamics must be tailored to make the system degradation at failure as graceful or well behaved as possible.

Past studies have considered controlled element transitions in single-loop tracking tasks,[1-6] and have concentrated on failure identification and the pilot's dynamic response. The current results emphasize multiple-loop lateral/directional control of a conventional airframe with a yaw damper failure which results in an abrupt or step-like change in the dynamics.

Of central concern is the pilot's ability to retain control at failure, and its influence on the design of the augmentation system. If a basically unstable vehicle is augmented to give very good handling, then the transition due to failure will require a large increment in his readaptation and pilot-vehicle performance may be very poor. If, instead, the prefailure augmentation is reduced to place modest demands on the pilot at all times, the increment of readaptation will be smaller and the pilot should perform more capably if and when a failure occurs.

These notions lead to the graceful degradation hypothesis which states: the pilot's transition response and performance are a function of the difference in controlled element dynamics across the transition as well as their absolute forms; the larger the difference the greater the control difficulty and the poorer the performance.

Experiments

The primary experimental task was roll control in a fixed-base simulator with a center stick and a simulated gust input. The failure task involved pilot control with rudders of the dutch roll mode. The roll response was good for all configurations. The dutch roll dynamics were varied from good to poor by changing the level of prefailure augmentation. No manual rudder was required in the good configuration, whereas various amounts were needed in the poorer ones. At failure, the effective airframe dynamics went from good to bad or poor to bad. Both hardover and soft failure modes were simulated. There was no longitudinal control task.

Presented as Paper 68-819 at the AIAA Guidance, Control, and Flight Dynamics Conference, Pasadena, Calif., August 12–14, 1968; submitted August 12, 1968; revision received July 2, 1969. This research was accomplished in part under NASA Contract NAS2-3607.

* Senior Research Engineer. Member AIAA.
† Principal Research Engineer. Member AIAA.

The basic airframe had bad handling, requiring the pilot to fly with manual rudder plus aileron in the event of an augmenter failure. It had the following transfer functions:

$$\frac{\varphi}{\delta_a} = \frac{-40[s^2 + 2(-0.13)(2.29)s + (2.29)^2]}{(s - 0.024)(s + 5)[s^2 + 2(-0.14)(2)s + (2)^2]}$$

$$\frac{r}{\delta_r} = \frac{-3.2(s + 5)[s^2 + 2(-0.04)(0.4)s + (0.4)^2]}{(s - 0.024)(s + 5)[s^2 + 2(-0.14)(2)s + (2)^2]}$$

The gains were selected for good subjective opinion.

To test the graceful degradation hypothesis, three prefailure, augmented airframes were used; configuration A with good to optimum dutch roll damping, configuration B with only slight dutch roll damping and a value of ω_φ/ω_d equal to unity, and configuration C with slight damping and ω_φ/ω_d equal to 1.3. Closure of the augmenter loop alone did not provide all of the desired configurations, therefore, the stability derivatives N_r' and N_{δ_a}' were changed as the augmenter loop was closed. They were changed back instantaneously when a configuration was failed to the basic airframe, allowing ω_φ/ω_d, ζ_φ, and ζ_d to be varied independently of any other parameters. To keep the augmenter from opposing the pilot, a washout was included with a time constant of 0.25 sec. The resultant bank angle control transfer functions for the augmented airframe were

configuration A (good handling):

$$\frac{\varphi'}{\delta_a} = \frac{-40(s + 0.31)[s^2 + 2(0.70)(1.80)s + (1.80)^2]}{s(s + 0.52)(s + 5)[s^2 + 2(0.70)(1.65)s + (1.65)^2]}$$

configuration B (fair handling, $\omega_\varphi \doteq \omega_d$):

$$\frac{\varphi'}{\delta_a} = \frac{-40(s + 0.26)[s^2 + 2(0.16)(1.96)s + (1.96)^2]}{s(s + 0.29)(s + 5)[s^2 + 2(0.15)(1.95)s + (1.95)^2]}$$

configuration C (poor handling, $\omega_\varphi > \omega_d$):

$$\frac{\varphi'}{\delta_a} = \frac{-40(s + 0.26)[s^2 + 2(0.13)(2.52)s + (2.52)^2]}{s(s + 0.29)(s + 5)[s^2 + 2(0.15)(1.95)s + (1.95)^2]}$$

With configuration A, the pilot could use aileron-only control before the failure, but had to use the rudder after the failure. In B, the pilot needed rudder occasionally before the failure if the dutch roll was excited by the disturbance. With C, some manual rudder control was always required before failure in order to avoid destabilizing the dutch roll. The system block diagram is shown in Fig. 1. Representative values were used for the valve-actuator and rate gyro dynamics.

Two kinds of rudder transients simulated hard and soft failure modes, applied at the instant of augmenter failure. The hard-failure transient consisted of a washed-out rudder step, while the soft failure was simulated by a hold of the rudder signal from the augmenter. As a result of the hold there was no disturbance imposed on the airplane by the failure, and the pilot had no immediate cue in the soft case to indicate that the augmenter had failed. The two kinds of failures were presented with equal frequency in a random sequence.

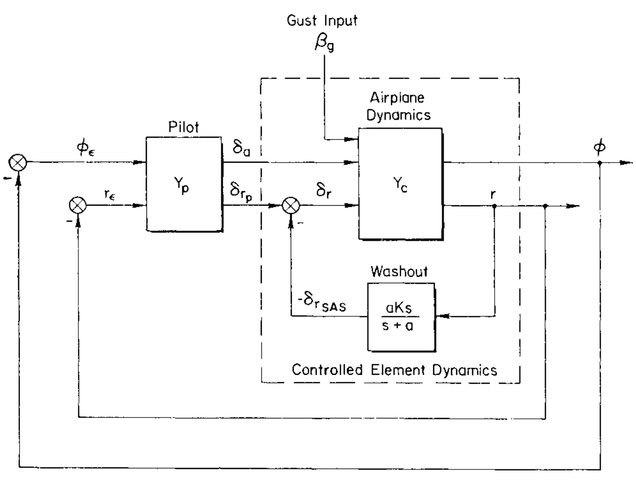

Fig. 1 Block diagram of multiple-loop control task.

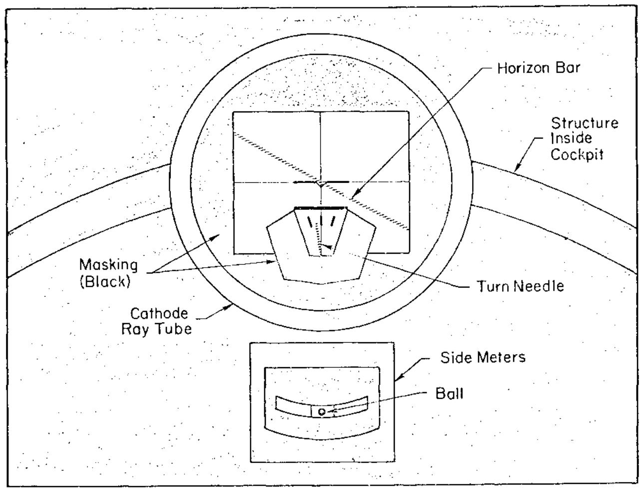

Fig. 2 Cockpit display (indicating coordinated left turn).

The cockpit was fixed-base with a conventional seat, stick, and pedal arrangement. An opaque canopy was closed during the experiments to isolate the pilot from the laboratory. The primary displays of bank angle and turn rate were generated on a two-gun CRT located roughly at eye height and about 12 in. from the pilot. An additional meter was set up to indicate lagged sideslip. The displays were arranged as shown in Fig. 2.

The airplane was disturbed by a random-appearing sideslip gust input consisting of a sum of low-frequency sine waves. The effective displayed bandwidth was determined by the closed-loop airframe dynamics. The rms gust amplitude was 1.4° sideslip.

Two subjects were used in the experiments, one a high-time military pilot and the other a private pilot. They were instructed to "Minimize bank angle and yaw rate as well as you can throughout the run. Try to never lose control of the vehicle after a failure. If at any time you feel tired, we would like you to stop and take a rest so we can get consistent data from an alert pilot." The pilots evolved their own error minimization strategy. They practiced the failure tasks 2 hr a day for several days, until each had accomplished about 200 trials over all configurations. Subsequent to the learning period, their performance and response was substantially the same in a given task. The experimental runs lasted about 2 min each, with the augmenter failure occurring after about 1 min. The runs were grouped in consecutive sets of 5 (with the same prefailure dynamics), followed by a rest period between sets. The data of interest here consist of 90 experimental runs for the military pilot, providing 15 replications for both the hard and soft modes with each of the 3 prefailure configurations.

Basic Pilot Response Model

Prior research[1] has shown that the pilot's response to a controlled element transition in a single-loop task contains four general phases: prefailure steady state stationary quasi-linear control; retention of prefailure pilot adaptation; nonlinear control, large control actions which stabilize the system and reduce the error to some acceptable level, and postfailure steady state, stationary quasi-linear control. Block diagrams of these phases are given in Fig. 3.

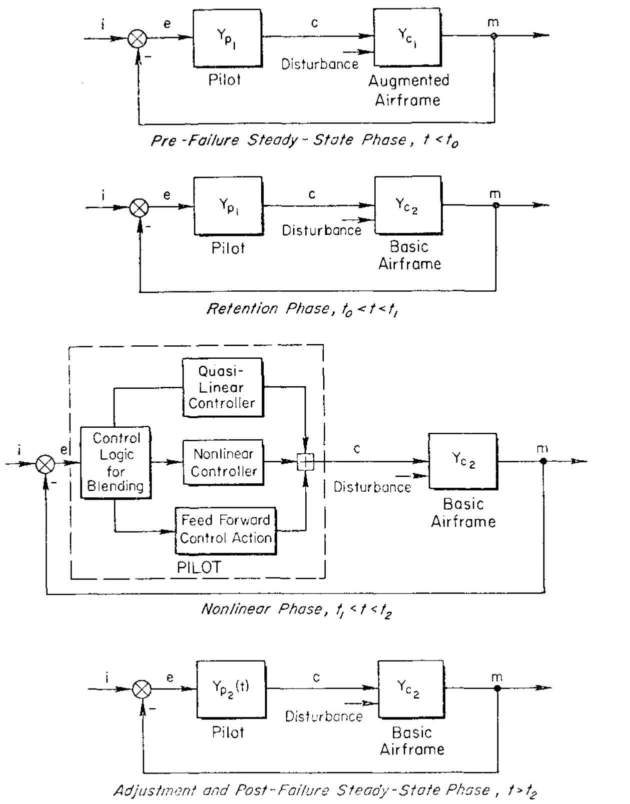

Fig. 3 General pilot model for step change in controlled element dynamics.

The retention phase starts at failure and consists of the prefailure pilot adaptation and the postfailure controlled element operating in a closed-loop fashion on the system error. This frequently results in an unstable condition because the typical failure results in a substantial increase in the controlled element lag. The end of retention is defined as the point in time when an abrupt change in manipulator motion begins, and it is most evident in transitions where a nonlinear control phase can be clearly seen.

The nonlinear control phase starts at the end of retention and continues until the system error has been reduced (approximately) to within the post-failure steady-state envelope. Time optimal control with fixed bang amplitude is an appropriate limiting case.[1] The error envelope may exhibit a further, more gradual decay during the first seconds of the steady-state period, corresponding to an additional adjustment phase[1] wherein the pilot's final adaptation is being achieved by an optimizing process.

The retention and nonlinear control phases are not always present, particularly with highly practiced subjects, fixed base, and relatively mild changes in the controlled element dynamics. With less practice, moving base, fatigue, higher pilot workload, etc., typical of in-flight situations; these phases are dramatically evident. Available moving base[3] and single axis[7] data show catastrophic results when the pilot is unable to progress beyond the retention phase.

The retention and nonlinear control phases comprise the period of transition. Identifying them in the data is sometimes difficult. When present, these phases are evident in the yaw rate and rudder pedal recordings. The transition period of large motions is undesirable and the point of this research is to determine ways of reducing or eliminating it by suitably tailoring the effective airframe dynamics.

Typical Results

Typical examples of failures from configuration A are shown in Figs. 3 and 4. This prefailure configuration was designed to have enough damping that the effective airframe would be subjectively good and no pilot rudder control would be needed. Despite this, the pilot used the rudder intermittently before the failure, as shown in Figs. 3 and 4, presumably because of the high level of gust disturbance (which excited the aircraft) and the fact that the pilot knew a failure was coming. Interestingly enough, the approximate level of pilot-rudder gain which this represents is enough to make the postfailure configuration stable. Hence, the pilot appears to close a low-gain rudder loop with configuration A in anticipation of the failure. With the soft failures involving configuration A, Fig. 3, a retention phase and a nonlinear control phase are both evident in the transition period following failure during which the yaw rate amplitude envelope builds up and then decreases to its steady-state level. The retention phase is the first few seconds following failure where the pilot continues to respond as he did with the prefailure configuration. It has a relatively long duration for this situa-

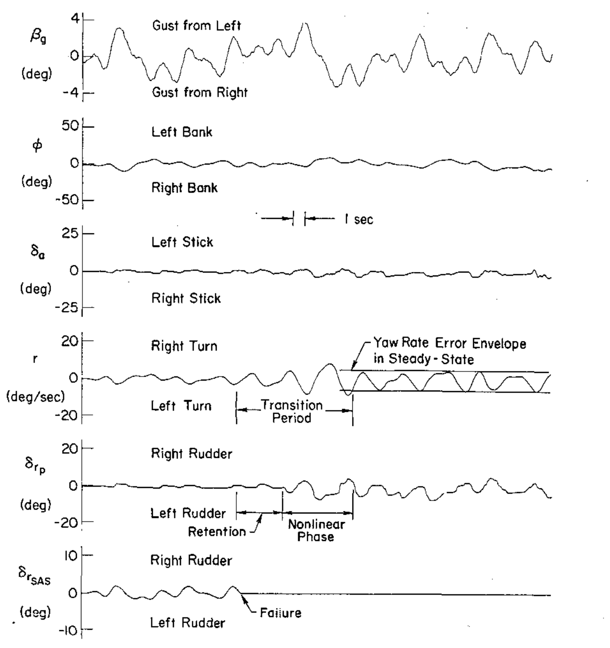

Fig. 4 Example of a soft failure with configuration A.

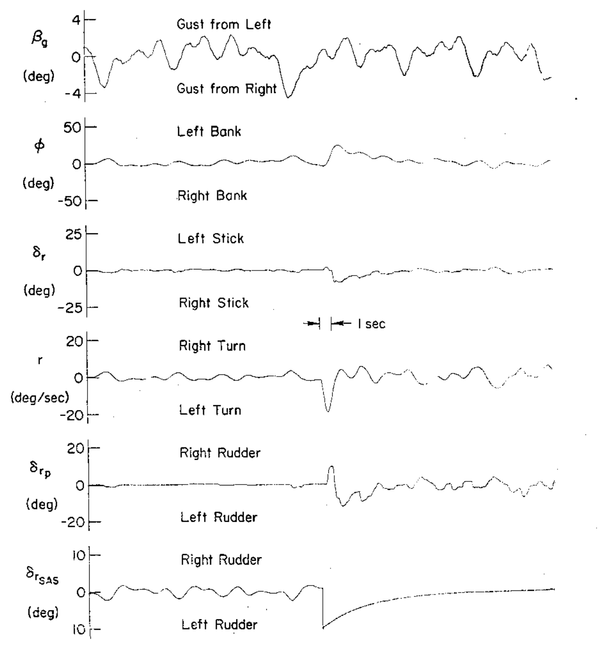

Fig. 5 Example of a hard failure with configuration A.

Table 1 Average results for transition parameters

Transition parameter	Configuration A Soft	A Hard	B Soft	B Hard	C Soft	C Hard
No. of runs with identifiable transition region	7	15	2	15	0	15
Peak yaw rate, deg/sec	5.5	16	5.5	15	5	13.5
Transition duration, sec	4.4	5.5	3.5	9.5	0	7.5
No. of significant yaw rate overshoots during transition	1.3	2.9	1.1	3.5	0	2.9

tion, as shown in Fig. 3. In Fig. 4, the hardover results in an immediate response following a short time delay.

Typical failures from configuration B are shown in Figs. 5 and 6. This prefailure configuration was designed to give a lightly damped dutch roll nuisance mode, but one that was not destabilized by aileron control of bank angle. The figures show that the pilot used larger rudder amplitudes with B than with A. The transition period is not evident in Fig. 5. A rapid recovery following failure can be seen in the yaw rate of Fig. 6.

Configuration C failures are typified by Figs. 7 and 8. This configuration was similar to B except that bank angle control with aileron tended to destabilize the dutch roll mode because of the adverse ω_φ/ω_d ratio. These figures show that the pilot's rudder activity before and after the failure is about the same. On some of the runs involving soft failures with configuration C, it was not possible to detect the failure from the time traces of yaw rate, bank angle, aileron, and rudder because they all look essentially the same (in magnitude and frequency content) before and after failure. An example is given in Fig. 7. This suggests a rapid adaptation for such cases. Rapid recovery following a hardover failure is shown in Fig. 8. The hardover failures for configuration C show a retention time of about 0.5 sec. After retention a large rudder pulse (also lasting about 0.5 sec) was applied in the direction to oppose the input, followed by a crude rudder step of opposite polarity lasting for several seconds. This opposite

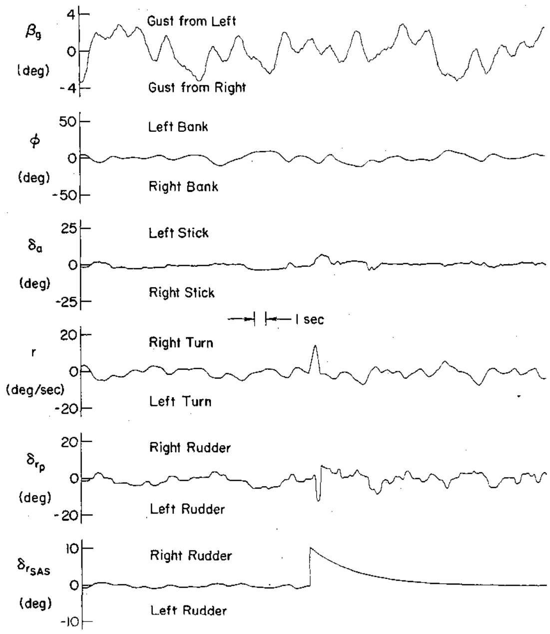

Fig. 7 Example of a hard failure with configuration B.

step was needed to cancel the yawing moment due to the aileron which the pilot was using to remove the bank angle induced by the failure transient, as shown in Fig. 8.

The data for prefailure configurations B, C, and for the basic (failed) airframe typically show a crossfeed by the pilot of aileron to rudder control. The effect of the crossfeed is to augment N_{δ_a}', thereby reducing or cancelling the adverse effect of ω_φ/ω_d. Although configurations C and the basic airframe are the only ones requiring such a crossfeed, it was found that the pilot also used crossfeeds with A and B (which already exhibited pole/zero cancellation for the dutch roll mode). This may be explained by the pilot learning that a crossfeed was desirable after the failure for better performance, and that a prefailure crossfeed didn't reduce performance. These results support the graceful degradation hypothesis.

The data for all 90 runs were analyzed by measuring typical transition parameters and tabulating these for comparison of

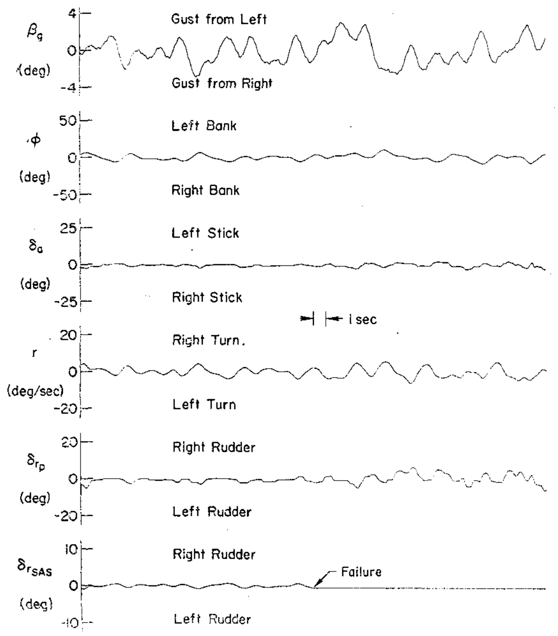

Fig. 6 Example of a soft failure with configuration B.

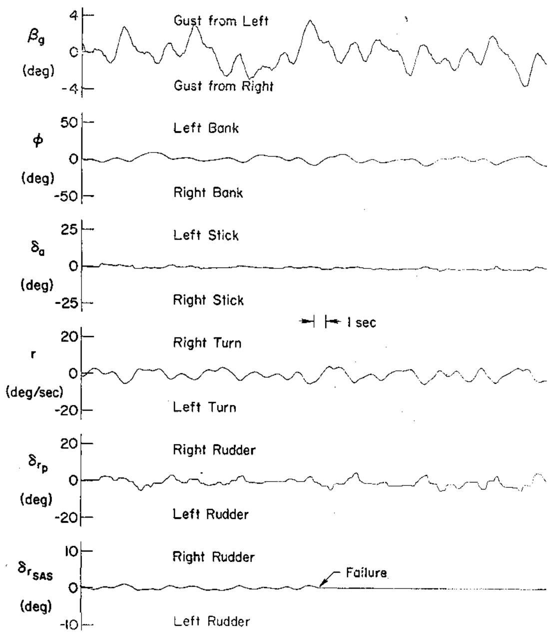

Fig. 8 Example of a soft failure with configuration C.

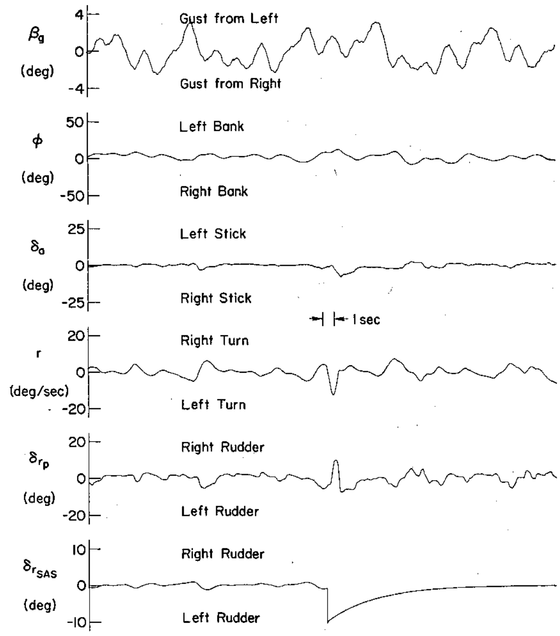

Fig. 9 Example of a hard failure with configuration C.

the configurations. Average results are shown in Table 1. Note that the averages include the runs with no transitions.

The distribution of maximum yaw rate peaks for the soft failures shows that the relative number of runs with yaw rate errors large enough to give an obvious transition region increases as the prefailure dynamics improve. The hard-failure data show that configurations B and C give smaller maximum yaw rates than A. Both of these results support the graceful degradation hypothesis.

The transition was considered to last as long as the yaw rate error peaks remained outside the envelope of the post-failure steady-state error peaks. The distribution of transition durations for soft failures showed that this only occurred with the good prefailure dynamic Since the poor dynamics (configuration C) did not peak above the postfailure asymptotic level, it strongly supports the graceful degradation hypothesis. The results are inconclusive for the hard failures. Perhaps the duration is affected more by the hardover transient acting as a disturbance (the same in all cases) than it is by the change in dynamics.

The distribution of yaw rate error overshoots for the soft failures showed that no significant error overshoots occurred for the poor prefailure dynamics while they occurred frequently with the good prefailure dynamics, again supporting the hypothesis. For the hard-failure data, the distribution of the number of yaw rate error overshoots does not lead to a conclusion regarding graceful degradation.

In summary, most of the soft-failure data supports the graceful degradation hypothesis of better performance being associated with smaller changes in the controlled element dynamics. This was also true for the maximum error peaks of the hard failures. Some of the hard-failure data are inconclusive. There was no evidence to support a hypothesis contrary to that of graceful degradation being associated with smaller changes in the controlled element characteristics.

Revisions to the Pilot Response Model

A composite pilot model for step changes in controlled element dynamics results from the available single- and multiple-loop data, which applies to skilled subjects in practiced situations. It is applicable to the failed loop in a multiple-loop situation if only one loop involves a significant change in the controlled element dynamics.

The data analyses show that the general form of model given in Ref. 1 is still valid with some extension and modification. Numerical estimates are now available for the retention duration, when this phase exists (see Figs. 4–9). The nonlinear control phase may include a learned response (such as a feedforward step) that is merely triggered by the failure. A quasi-linear controller plus feedforward is appropriate for the nonlinear phase, in the hardover cases with a stable retention phase (e.g., configuration C), and in those instances the signal fed forward should resemble the hardover transient. A nonlinear controller plus feedforward is appropriate for the hardover case with unstable retention, but because the response is probably indistinguishable from the nonlinear-controller-alone version. Whichever model is easier to use is appropriate. For soft-failures with some systems, the nonlinear phase is not present. It is present only when required, i.e., when quasi-linear control is inadequate to maintain system errors within some criterion amplitude envelope. Time-optimal control alone during the nonlinear phase is still a valid idealization for the soft failures, as an attainable limiting case, however, some suboptimal control mode is a more accurate match to the data. Postfailure steady-state error characteristics were generally larger than that predicted by an optimized quasi-linear model and this suggests that any adjustment (optimization) phase is relatively long-term (5–10 sec) and does not end immediately following the nonlinear phase. Prefailure conditional adaptation may be present. This means that the pilot may use an adaptation prior to failure which is not a steady-state optimum but which will improve the system performance immediately after failure (such as a modified gain or a crossfeed).

Implications for Flight Control Design

The graceful degradation hypothesis is supported by the experimental results. Therefore, a flight control system design that gives less change in effective controlled element following a failure will give better transition performance.

The transition performance with hardover failure was relatively invariant with the amount of change in the dynamics at failure. The performance with soft failures tended to vary as a function of the changes in dynamics, but it was generally better than that for hardovers. The relative invariance of the hardover failures is a consequence of the retention times being about the same due to alerting. It suggests that if alerting can be designed into (or accompany) soft failures, performance will be improved.

For soft failures, the crux of the entire graceful degradation situation lies in the closed-loop system stability during the retention phase; that is, the stability of the prefailure pilot model, Y_{p_1}, controlling the postfailure vehicle, Y_{c_2}. If $Y_{p_1}Y_{c_2}$ is stable, closed loop, then a critical situation generally does not develop after a failure. This can be achieved by suitable tailoring of Y_{c_1} and Y_{c_2}.

A flight control (e.g., stability augmentation) system should be designed to keep the pilot in the loop prior to a failure rather than to allow (or require) him to merely monitor the controller-vehicle system activity. This might be accomplished by giving the pilot the task of controlling a simultaneous model of the actual system. One payoff for this is obvious—the pilot will have a shorter retention period. For example, a pilot who has his feet on the pedals at the time of failure because he is in the loop will have a shorter time delay prior to initial rudder input than a pilot who is merely an observer and has his feet on the floor. For some kinds of failure, a long retention period or delay can be catastrophic.

A pilot tends to use conditional adaptation, which means that his prefailure control activity is a function of the probability of failure and the severity of the transition. If $Y_{p_1}Y_{c_2}$ closed loop (or Y_{c_2} alone) presents difficult control characteristics, then the alert, practiced pilot will anticipate a failure by modifying his prefailure adaptation, Y_{p_1}, to improve

$Y_{p1}Y_{c2}$. This will be something other than the stationary adaptation appropriate to Y_{c1}. Consequently, a control system should not only make prefailure flight control available, but it should be designed so that considerably improved $Y_{p1}Y_{c2}$ stability is possible with simple modifications in prefailure control activity (as was the case in the multiloop experiments).

The data have implications for pilot training. During training of the pilot subjects in the experimental tasks, performance following a failure improved rapidly as the number of practice runs increased. Even hardover failures gave little performance degradation when encountered regularly. Any failure can be difficult if it represents a significant change from recent experience (as was found during early failure practice runs following only steady-state control of each configuration), which implies that pilots should have frequent refresher training to keep them current with regard to flight control system failure situations.

Conclusions

The experimental results indicate that the graceful degradation hypothesis is valid. Thus, an improvement in tracking performance can be expected during a transition in controlled element dynamics if the difference between the prefailure and postfailure dynamics is diminished. In practice, this means that the augmented vehicle should be less than ideal, and the basic airframe should have some minimal level of controllability, if a critical failure can occur.

A previous model for human pilot dynamic response during controlled element transitions[1] was verified and extended. The verification shows that the following transition response phases do exist in general: prefailure steady-state quasi-linear control, retention of prefailure pilot adaptation, nonlinear control, and postfailure steady-state quasi-linear control. The extensions relate to the existence of hardover (deterministic) signals in the system and occur in the nonlinear control and postfailure steady-state phases.

References

[1] Weir, D. H. and Phatak, A. V., "Model of Human Operator Response to Step Transitions in Controlled Element Dynamics," CR-671, Jan. 1967, NASA.

[2] Sheridan, T. B., "Time-Variable Dynamics of Human Operator Systems," TN-60-169, March 1960, AFCRC.

[3] Sadoff, M., "A Study of a Pilot's Ability to Control During Simulated Stability Augmentation System Failures," TN D-1552, Nov. 1962, NASA.

[4] Young, L. R. et al., "The Adaptive Dynamic Response Characteristics of the Human Operator in Simple Manual Control," TN D-2255, April 1964, NASA.

[5] Elkind, J. I. and Miller, D. C., "Adaptive Characteristics of the Human Controller of Dynamic Systems," TR-66-60, July, 1966, AFFDL.

[6] Elkind, J. I. and Miller, D. C., "On the Process of Adaptation by the Human Controller," *Third Congress of the International Federation of Automatic Control*, London, June 1966.

[7] Weir, D. H. and Johnson, W. A., "Pilot Dynamic Response to Sudden Flight Control System Failures and Implications for Design," CR-1087, June 1968, NASA.