Onboard Automatic Aid and Advisory for Pilots of Control-Impaired Aircraft

Elaine A. Wagner*
General Dynamics, Fort Worth, Texas 76101

Existing approaches for recovering aircraft after control failures are unsatisfactory. This work reflects consideration of the aircraft control failure problem from a broader viewpoint. Postfailure performance and operating constraints best taken explicitly into account are discussed, and the role of explicit postfailure retrim is presented. Automatic onboard emergency control is seen to be a very significant part of the problem of recovering from a control failure. A rule-based expert system approach was taken to find successful control strategies after jam elevator failures on a C-130 aircraft. The system is described and demonstrated. As a general means of augmenting traditional control reconfiguration, this type of approach seems to be a way of directing use of any remaining control capability, or use of the unusual or counterintuitive ideas sometimes required in recovery. Aspects of a postfailure pilot advisory function, proposed for use in more quiescent periods of flight after first failure accommodation, and also amenable to the expert system approach, are described briefly.

Introduction

CONTROL failures on aircraft are not uncommon. A survey of recent civil aircraft accident reports yielded 25 cases involving failures of controls other than engines.[1] In all but five, most or all of those on board the aircraft perished. In more than half of the catastrophic cases, the flight could have ended safely if the pilot had acted in a correct and timely manner. Reference 2 describes a fascinating complete recovery of a control-impaired aircraft.

Aircraft are increasingly dependent on control for stabilization, maneuvering, and load moderation. However, most of the potential functional redundancy in controls has not yet been exploited. The problem of recovery and control reconfiguration after actuation failures is gaining increasing high-level attention.[3,4]

Generally, when researchers refer to the issue of recovery from control failures, they refer to basic failure robustness of the automatic control—of which most aircraft have some—or to the problem of reconfiguration of that control.[4] ("Reconfiguration" as used herein includes control reconfiguration, restructuring, or real-time onboard control redesign.) Most research on the problem of control failures has been done in the areas of control-loop robustness and reconfiguration. Motivating this have been the considerations that the aircraft must be dynamically stable or stabilized after the failure to have any possibility for recovery, and that the capabilities of the aircraft should be restored to the maximum possible extent.

The importance of changing the automatic control law after a control failure is evident, particularly for higher performance aircraft. But automatic control is limited in some ways. The authority regarding important aspects of controlling the aircraft's flight is given to the pilot. The pilot's perception of the remaining capabilities of the control-handicapped aircraft can be crucial in determining whether the flight will ultimately end safely, as was apparent in accident cases studied.

Depending solely on robust or reconfigured control in the usual sense will generally not be sufficient to allow a control-impaired aircraft to be recovered because of the following:

1) It is probable that neither the pilot nor the traditional automatic control will take into account all of the alternate control capabilities of the vehicle. Automatic control is usually not designed to drive all controls effective in a given control axis. For example, longitudinal control may modulate elevator but not thrust. However, even the use of landing gear, spoilers, leading-edge slats, flaps, and reverse thrust on the ground can all impact recovery in significant ways.

2) Failures often induce significant new constraints on the controllable operation of the aircraft and on the performance that it can achieve. Traditional types of automatic control will not "know" about these constraints and cannot take them effectively into account.

3) A successful recovery control strategy for strenuous failures can be a very complicated multigoal process involving carefully coordinated changes in controls in multiple axes and, in some cases, unusual or even counterintuitive actions.

Given that some of the responsibility for postfailure actions can be expected to continue to rest with the pilot, certain advice and warnings should be provided to the pilot during postfailure flight.

It will be assumed in the following that the failure is a jam failure, although the considerations can apply whether the failure is a jam, bias, or floating surface failure, part of the surface is lost, or whether there has been a failure in the propulsion system. The first section of this paper presents a categorization of the postfailure constraints on operating state and performance that should be taken into account in flying a control-impaired aircraft. In the second section, there is a discussion of the role of explicit determination of appropriate operating points for the impaired aircraft. Next, a rule-based expert system developed to find successful emergency control in the initial postfailure period after jam elevator failures on a C-130 aircraft is presented. As will be discussed, this is an idea that can subsume more usual ways of dealing with control failures but can overcome the limitations of these types of approaches. There is also a description of an integrated onboard recovery aid and advisory system, including an examination of how advisory information could be made available and of pilot-system interface issues. The work reported here was based on discussions with pilots, reported accident cases, and work with a simulation of a STOL C-130 aircraft built in the late 1970s.

Categorization of Postfailure Constraints on Operating State and Performance

In research on the problem of control failures, there has been little discussion of the changes a failure may induce in

Presented as Paper 89-3460 at the AIAA Guidance, Navigation, and Control Conference, Boston, MA, Aug. 14-16, 1989; received Jan. 31, 1990; revision received July 13, 1990; accepted for publication July 18, 1990.

*Engineering Specialist, Control Law Design and Analysis, P.O. Box 748. Senior Member AIAA.

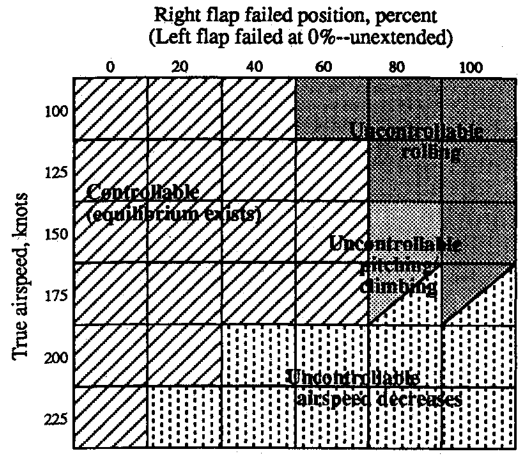

Fig. 1 Effects of asymmetric flap failures on the C-130 aircraft (looking for trim points at 1000 ft altitude).

safe and feasible operating points and on the performance an aircraft can achieve. This information seems to be very important, however, in fully recovering an aircraft after a failure. As will be discussed, this information may be needed in formulating emergency control strategies and advising pilots of control-impaired aircraft.

Postfailure Operating Constraints

The term "controllable airspeed" is already familiar as the minimum indicated airspeed at which a rudder can neutralize yaw induced by an asymmetric engine failure. This terminology can be used in a more general way to refer to airspeed-related limitations of the functioning controls in counterbalancing effects of any type of control failure.

An example of controllability airspeed constraints for a relatively common but, as accident cases showed, still dangerous control failure mode—jam asymmetric flap failure—will support this idea. Figure 1 shows a matrix of discretized airspeeds and flap asymmetries for the C-130 aircraft used in this study. Trim points for steady straight-and-level flight were sought at 1000 ft for zero sideslip. As the figure shows, there may be failure-induced limitations on both minimum and maximum airspeed. The ailerons have limited ability to control the failure-induced rolling when the airspeed is low, and there is limited thrust available to counteract the drag added by the jammed flap at high speeds.

Equilibration by aerodynamic controls achieved at a certain indicated airspeed does not imply that stabilization could be achieved at another airspeed. Jammed extension of a surface may cause effects that vary quite differently with indicated airspeed from the effects of the potentially counterbalancing surfaces. There will be situations in which the side effect from usage of alternate controls will itself constrain the airspeeds for which a failure is controllable.

Stall airspeed can change significantly as a result of a control failure, e.g., when a leading-edge slat fails extended, disrupting flow over the wing, or when a lift-augmenting surface fails retracted. When stall speed changes, the corner velocity—the lowest velocity at which limit load can be obtained—will change. The buffet boundary (airspeeds at which high-speed flow separation occurs) can also change with certain control failures. Gust penetration airspeed is the maximum airspeed at which expected gust loadings cannot result in the aircraft limit load being exceeded. When stall speed changes, so will the gust penetration airspeed.

Control reversal airspeed is the airspeed at which aircraft flight can be sustained with minimum power or thrust required. Failures of aerodynamic surfaces can lead to changes in control reversal airspeed either directly or indirectly through controls used to compensate for the failure. Control reversal airspeed is very important in landing an aircraft. Its value would increase with certain types of control failures, making landing at normal speeds quite dangerous. Even if control reversal airspeed does not change with a failure, the nominal value may be of explicit importance during a recovery. In a certain DC-3 case in which primary aerodynamic lateral and directional control was lost, very careful use of asymmetric thrust would have been needed for control in these axes. Any attempt to control the aircraft with thrust would have decreased total power available for changing airspeed on approach, making operating below the control reversal airspeed more dangerous than usual.

There are other types of postfailure restrictions on aircraft operating state. Figure 2 comes directly from a reported accident case. It shows the variation with angle of attack and Mach number of the difference between the rolling moment induced by the single failed-extended leading-edge slat and that available from the remaining lateral control resources. The aircraft could not be recovered after roll was allowed to exceed a certain level because Mach number and angle of attack became too large. Potentially, any dependencies (in the

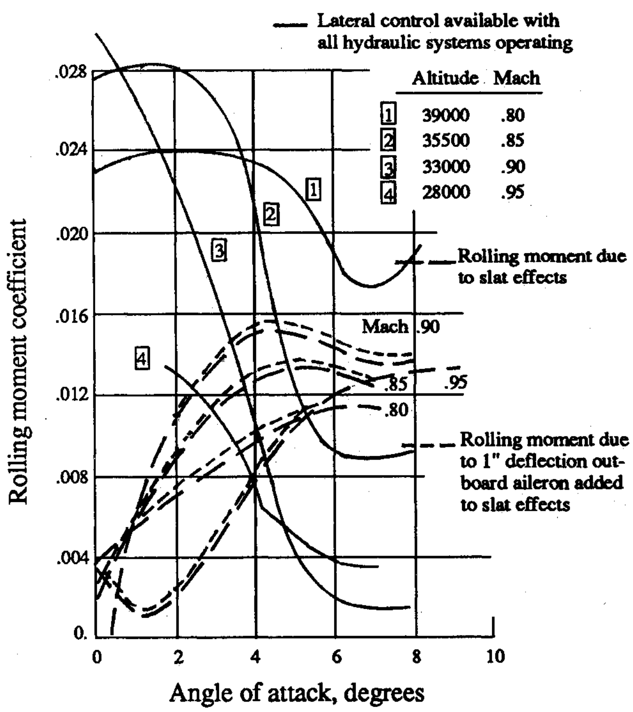

Fig. 2 Rolling moments from extended no. 7 leading-edge slat on B-727 aircraft.

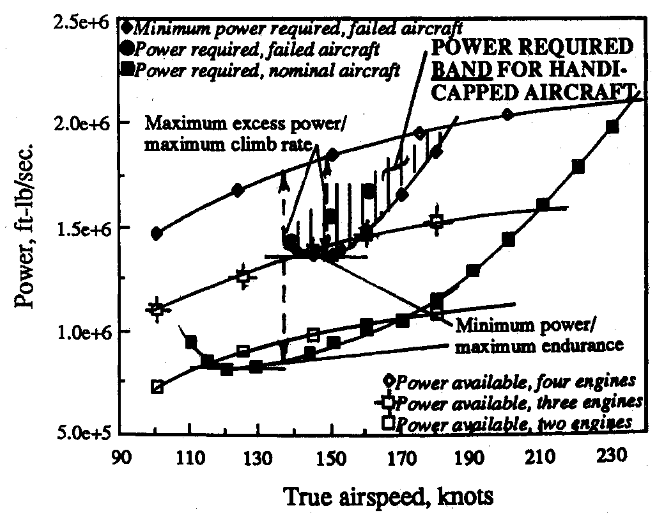

Fig. 3 Power-available/power-required curves, C-130 with elevator jammed at 8.05 deg; steady straight-and-level flight at 10,000 ft.

functional sense) of a dynamic quantity can be restricted by a failure. Note that limitations may be expressed more usefully in terms of certain parameters rather than others (such as bank angle rather than Mach number or angle of attack).

Postfailure Performance Constraints

Any aspect of aircraft performance can suffer greatly when there has been a control failure. The high drag associated with highly deflected jammed surfaces or, secondarily, with the compensating controls, can bring about significant degradation of such basic types of performance as range, endurance, climb angle and rate, and maximum airspeed and altitude. Failure of a wing surface such as a slat or spoiler, by changing the basic aerodynamics of the aircraft, can also lead to significant changes in power required for flight. Engine failures will bring about degradation of all of these aspects of performance. These types of performance can change after a failure and so can the configuration, e.g., values of airspeed and angle of attack, at which maximum performance of a certain type is achieved. Misjudging aircraft performance after failures can lead to disastrous errors in deciding whether a certain destination can be reached, as some accident cases have demonstrated. Maximum performance and how it could be achieved with an impaired aircraft would change over the course of its flight, due to changes in altitude and vehicle configuration and weight. In the most strenuous cases, these variations would have to be accounted for along the way if the aircraft were eventually to land safely.

An accident case in which postfailure performance limitations played a role was one in which the aircraft's leading-edge slats failed to extend on takeoff. The pilot was probably trained to deal with this type of failure, but the aircraft was soon flying at such low speed that available thrust was insufficient to meet the requirements for climb. Two knots of airspeed would have made a difference between the aircraft being able to climb and the gradual increase in drag, loss of height, and the ground impact (and fire) that actually occurred.

Performance constraints induced by a failure can be very significant, even if the failure seems small. Consider a C-130 case involving an elevator jammed symmetrically 5 deg off-nominal pitch down (jammed at 8.05 deg). After the failure, the pitch moment commands intended for the elevator were remapped to the remaining longitudinal controls (more on this reconfiguration later). The impaired aircraft was carefully flown to 10,000 ft and stabilized. Figure 3 shows the marked contrast in power required for straight-and-level flight at altitude with and without the failure. (Points on the new power-required curve were obtained with quadratic programming, using vehicle dynamics linearized at each successive iteration. State and control settings for the cruise equilibria converged fairly slowly, and reducing the state rates from 10^{-3} to 10^{-6} or 10^{-8} could often result in considerable changes in the solved-for state and control settings to give equilibrium flight. A not-fully-converged solution could imply very different postfailure capability than was actually available.)

Keeping the aircraft flying level with this failure required considerable leading-edge flap extension for its pitch-up effects, and the deployment of full pitch-up collective ailerons (the aircraft simulation had been modified so that the ailerons could be deployed symmetrically) and some elevator tab. Deployment of these surfaces, particularly the flaps, resulted in considerable additional drag. The power required ("minimum power, failed aircraft") is considerably changed from the nominal. Figure 3 shows that the airspeed range available for steady, level flight with all four engines operating is approximately halved by the failure. The steady maximum climb rate and climb angle have been reduced considerably, important in deciding whether flight could be continued if climb over an obstacle were required. The maximum endurance speed has changed by about 30 knots true airspeed (KTAS), the maximum range airspeed by about 10 KTAS, and the airspeeds for best climb rate and angle have also changed. Failure of even one of the four engines now could restrict cruise operations considerably, as the three-engine power-available curve of Fig. 3 shows.

Explicit Determination of Postfailure Operating Points

Does explicit retrim have a role in flying a control-impaired aircraft? Suppose that an aircraft is flying with state x_0, control setting u_0, and with state rates $\dot{x} = f(x_0,u_0)$ prior to a failure, and that control i jams $\Delta n°$ off-nominal. Explicit retrim to regain and maintain the prefailure trajectory would involve trying to achieve a new operating point $x_0 + \Delta x$ wtih the unfailed controls at a new setting $u_{f0} + \Delta u_f$ such that

$$f(x_0 + \Delta x, u_{f0} + \Delta u_f, u_i + \Delta n) = f(x_0,u_0)$$

or, in a linear approximation,

$$f(x_0,u_0) + \frac{\partial f(x_0,u_0)}{\partial x} \cdot \Delta x + \frac{\partial f(x_0,u_0)}{\partial u_f} \cdot \Delta u_f + \frac{\partial f(x_0,u_0)}{\partial u_i} \cdot \Delta n = f(x_0,u_0)$$

or

$$A \Delta x + B_f \Delta u_f = - b_i \Delta n$$

A and B are the usual linear model matrices, $\partial f(x_0,u_0)/\partial x$ and $\partial f(x_0,u_0)/\partial u$, respectively. B_f is the matrix of columns of B asociated with the unfailed controls, and b_i the column of B associated with the failed control. Note that controllable airspeed or another postfailure operating constraint could be involved in determining whether a solution to these equations is possible.

Re-establishing the prefailure trajectory may not be best. For safety, climbing up or decreasing angle of attack may be preferred. For such cases, retrim might be to a new condition where

$$\dot{x} = f(x_0,u_0) + \Delta r$$

and the state and control setting changes accordingly satisfy

$$A \Delta x + B_f \Delta u_f = - b_i \Delta n + \Delta r$$

After multiple or large failures, there would be pronounced need for explicit information about operating regions in which stabilization or some required value of performance could be achieved. This researcher's experience with recovering aircraft with single jammed controls indicated that explicit retrim alone was usually not a very useful idea in early periods of flight after the failure had manifested itself. "Oppose the (typically large) disturbance" was a much more powerful general idea. Moreover, just having a single retrim point—which a single solution to these linear equations yields—said nothing about its basic reachability, nor, if guaranteed reachable, how to reach it. These could be extremely difficult issues unless the transition was simulated outright. Just to take a simple example, can a certain angle of attack, even if explicit retrim has shown it to be a useful condition, be achieved after failure? With an elevator failure, for example, it may not be achievable, since basic control authority might be lacking. Flying the C-130 aircraft to a given state and given control setting was very difficult. If there is one solution, there will be many[1] perhaps widely spaced solutions to any of these retrim equations, raising questions about choosing among them.

Explicit information about retrim would probably be most useful and important in periods of relatively quiescent flight when the new operating point is clearly reachable. As an illustration of the possible utility of explicit retrim under these circumstances, consider again the C-130 case where the elevator jammed symmetrically 5 deg down and led to such a

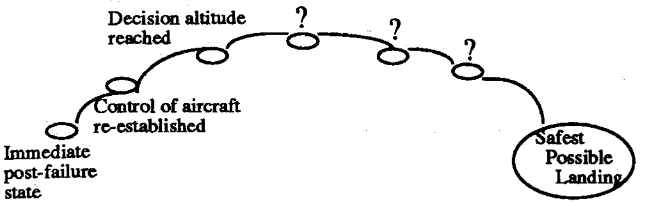

Fig. 4 Near-horizon planning in postcontrol failure flight.

significant decrease in achievable performance. A few additional operating points for steady straight-and-level flight at 10,000 ft were computed for the control-impaired aircraft. These have been superimposed as the higher "power required, failed aircraft" points in Fig. 3. There is a band of power required points at each airspeed, depending on how the failure compensation is distributed among the various viable controls. For example, two separate equilibria were established for the failed aircraft at 150 KTAS, one with 62% leading-edge flaps, angle of attack of -17 deg (a flown-to point in the recovery), and the actual minimum-thrust point, with 51% flaps, angle of attack of -2.25 deg. The difference in required thrust setting was about 8%. Achieving the minimum-power operating point would allow three-engine operation, as well as reducing flap deployment from a level near the structurally limited value. Additional flap deployment could then be available for pitch-up maneuvering, if needed.

The retrim-related idea of control effectiveness has been crucial in some accident cases, although there have been very few where flight after the control failure got past the phase where straightforward "oppose the disturbance" would have been sufficient. However, the following could be considered important later retrimming-type goals: 1) off-load the burden of compensation from controls needed for other purposes, such as maneuvering; 2) move away from dangerous state constraints; and 3) obtain better performance. These may be conflicting goals in some situations. The transitioning between operating points indicated in explicit retrim might be very difficult, because of the use of unconventional controls and the possible considerable degradation of aircraft capabilities after a failure. It may be difficult to decide how difficult or dangerous any given transition between two operating points would be, even where retrim has identified them as being in themselves valid operating points. Deciding what looks like a feasible, safe transition on the basis of simplified reasoning may be possible in some cases. Instructions could then be provided to the pilot to make the transition.

Reaching a new operating point might involve at least temporary losses in desirable operating quantities. Consider a transition to the maximum range operating point in Fig. 3 from another, lower airspeed cruise point. Thrust increases on the C-130 generated a pitch-up moment. Airspeed could be gained only through decreasing pitch by decreasing thrust. Altitude, heuristically a very desirable quantity in emergency piloting, would be lost in the transition and might be dangerous or impossible to regain afterwards.

Many extremely interesting issues related to retrim arose during this study. Reference 1 describes a few accessible mathematical properties of the constant-rate regions of systems of general nonlinear equations, such as those governing aircraft dynamics.

Postfailure Emergency Control

Let us consider a different aspect of the problem of control failures. An important part of the aid envisioned for flying control-impaired aircraft is emergency control in the initial period after the failure has mainfested itself. According to the pilots interviewed, the most important aspect of control failure recovery would be immediate and correct response to failures, particularly in high-speed or ground-proximity operations. Essentially all of the recoverable accident cases studied were vivid demonstrations of the need for immediate emergency control. The time available for required response was generally 4–5 s, but up to 25 s were available for some of the cases studied. Even more familiar types of failures seem not to have been identified by the pilot, even though the aircraft were usually not particularly high-performance aircraft. Pilots interviewed wanted full-authority automatic response after control failures.

Experience with Control Loop Reconfiguration

The most common solution advanced to the problem of emergency postfailure control is to provide failure-robust nominal control or reconfigure the control after the failure. In working with the C-130 to try to recover from jam elevator failures, at first the most conventional reconfigurable control idea was tried: ask that the forces and moments called for by the nominal control be generated by the remaining unfailed controls u_f, or, in terms of the linear model, at each time t,

$$B_f \Delta u_f = B \Delta u$$

where Δu represents the commanded control deflections if all controls were working. Typically, one solves for Δu_f as

$$\Delta u_f = B_f \dagger B \Delta u$$

where $B_f \dagger$ is the least-squares Moore-Penrose pseudoinverse of B_f.

This approach to reconfiguration is not particularly attractive from a control performance standpoint. Stability of the control loop reconfigured in this way cannot even be guaranteed for all initial conditions[1] unless there is full control redundancy—rank B_f = rank B. This reconfiguration is easy since it does guarantee stability and has been proposed by several researchers. One caution: nominal automatic control for aircraft is often of low gain, for safety or other reasons, and thus would have low authority to oppose usually large failure-induced disturbances.

Reconfiguration in this way was unsuccessful in preventing the C-130 from diving to the ground after a 5-deg off-nominal pitch-down failure when the vehicle was flying straight and level at 197 knots indicated airspeed (KIAS) at 1000 ft above the ground. Whether rank B_f = rank B or not, this type of reconfiguration can ask for very large deployments of certain controls, especially if they are particularly or solely effective in effecting changes in certain state rates. This can result in rate and position saturation of the controls (not to mention violat-

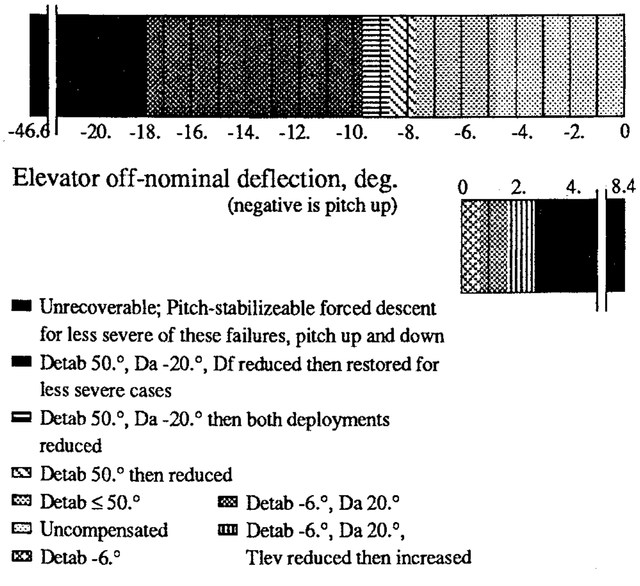

Fig. 5 Recovery spectrum—C-130 elevator jam failures during cruise at 147 KIAS.

ing limitations of the linear model), as was clearly demonstrated with the C-130.

Ultimately, the aircraft was successfully recovered when the controls were reconfigured after the failure to take into account rate and position saturation of the remaining functioning controls. More quickly deploying controls served as a stopgap until slower but more effective surfaces could deploy. This successful emergency control was the solution to the following quadratic programming problem, solving for changes in the compensating controls, Δu_f, given nominal elevator command Δu_i:

$$\min[B_f\Delta u_f - b_i\Delta u_i)^T Q(B_f\Delta u_f - b_i\Delta u_i) + \Delta u_f{}^T R\Delta u_f]$$

subject to

$$u_{f\min} < u_f < u_{f\max} \qquad \dot{u}_{f\min} < \dot{u}_f < \dot{u}_{f\max}$$

where b_i is the column of the B matrix corresponding to the elevator. The first attempts at recovering the aircraft via this reconfiguration showed that the elevator's intended effects on velocity, angle of attack, and pitch rate all could not be duplicated with the other controls. By insisting on equal weighting on each of these quantities (i.e., Q = identity matrix), the aircraft was lost. For the successful recovery, Q was a weighting on pitch rate only. This reconfiguration was then used to fly the aircraft with this pitch-down elevator jam and, in another case, after a certain pitch-up failure to stabilization at 10,000 ft and then to pitch-stabilized slow descent that would have allowed the aircraft to be flown safely onto the ground in a no-flare landing.

One choice of Q was as a weighting according to the nominal effect of the elevator on the various state rates, i.e.,

$$Q = \text{diag}\ (b_1^2, b_2^2, \ldots)$$

where b_x is the xth row of the column of the B matrix associated with the elevator's effects. This choice of weighting did not allow recovery from this failure. It simply diverted too many control resources from opposition of the failure-induced nose-down pitch moment.

This pseudoinverse type of reconfiguration was generally modestly successful, but it seemed to be very dependent on the choice of Q and R, raising questions about how these should be chosen to cover wide ranges of failures. Most importantly, however, this reconfiguration was not able to find recoveries for certain other failure cases where other techniques were able to demonstrate that one existed. It may be more useful after the initial disturbance opposition phase of the recovery.

Numerous automatic control reconfiguration ideas were studied. None proved flexible enough to do everything needed to recover an aircraft after a large failure-induced disturbance.

Manual Recoveries

To gain more experience, attempts were made to recover the aircraft after elevator failures using only manual control inputs, i.e., with automatic control on the aircraft disabled. Finding successful emergency control after C-130 elevator failures was usually not immediate and often involved several iterations of piloted simulation. Sometimes the strategy that was finally successful involved rather counterintuitive use of controls, as will be seen, not opposing the effects of the failure but enhancing them temporarily. Much more was involved than simply knowing which controls could be used to oppose a failure-induced disturbance. Multiple controls had to be used in combination and deployed very quickly in many cases if recovery was to be possible at all. However, finding a successful strategy was not too difficult, given the possibility of making a few attempts. The reasoning needed to devise effective emergency control strategies was not deep.

Experience with the C-130 showed clearly that the only way to be certain of an emergency postfailure control strategy was to simulate it with a high-fidelity model. Just 1 deg of control deployment available in some cases made all the difference between recovery and catastrophic loss of the aircraft. In addition, the side effects from usage of alternate controls ("artifact") could be devastating and difficult to foresee. An aircraft is a complicated nonlinear dynamic system and is very sensitive to changes in controls. These characteristics made some recoveries difficult, but made close investigation of all possibilities to get a recovering control strategy very worthwhile.

Recovery Piloting as Expert Behavior

An expert, in artificial intelligence (AI) terminology, has wide breadth and depth of knowledge in a given area, and an expert's problem-solving ability degrades gracefully at its boundaries. An expert can apply knowledge to solve problems effectively and efficiently using shortcuts to eliminate useless or unnecessary calculations.[5] This focusing depends on heuristics built up over long experience working in a given domain. Heuristics are the high-level guidelines that direct what to try in solving a problem. The expert's heuristic approach to solving a given problem is what he or she considers the best approach to try, and it often turns out to be the correct one. The heuristic may turn out to be inappropriate in a given case, and an expert backtracks to try something else. Heuristic information often seems more qualitative than quantitative.

One of the heuristics developed through the work with control reconfiguration and manually flying recoveries after elevator jam failures on the C-130 was that pitch rate was the most important parameter to control. Failure-induced disturbance opposition was almost always a very useful heuristic, although, as will be seen later, there were exceptions. Experience soon made clear that heuristic reasoning and very identifiable human problem-solving ideas (e.g., means-end problem solving of AI theory[6]) were sufficient to find a correct strategy. These techniques fit nicely under the heading of expert-type behavior, and this became the unifying viewpoint for all aspects of piloting control-impaired aircraft. Piloting by experienced pilots is an almost classical expert behavior. The long recovery described in Ref. 2 greatly helped confirm this viewpoint.

There were only five reported accident cases among the 25 studied in which a successful landing of the aircraft was eventually made. In all of these, flight after the failure could be divided into five discrete phases:

1) Regain control of the aircraft.
2) Achieve a safe altitude (on the order of 10,000 ft except where control was regained with the aircraft already higher).
3) Stabilize at altitude, determine landing capabilities, and decide where to land.
4) Approach landing site.
5) Make final approach and landing.

These became subgoals in short time horizon planning, as Fig. 4 shows. Achieving these flight goals reflects the "planning islands" idea, a very powerful human heurism. Driving toward islands greatly simplifies the calculations involved in the overall strategy. The ramifications or the identity of the failure would probably be better known in later phases of postfailure flight, anyway. This planning islands idea is also embodied clearly in the emergency procedures in military flight manuals, which, incidentally, are usually no more detailed than the following excerpt[7]:

> Three basic rules are established which apply to most emergencies occurring while airborne.
>
> 1. MAINTAIN AIRCRAFT CONTROL.
> 2. Analyze the situation and take proper action.
> 3. Land as soon as conditions permit.
>
> Normally, ejection is the best course of action in the event both engines flame out...or positive control of the aircraft cannot be maintained.

If structural damage occurs in flight, the pilot must decide whether to leave the aircraft or attempt a landing. If aircraft is controllable, proceed as follows:

WARNING

- In no case allow airspeed to decrease below 90 KIAS.
- Do not reset wing flaps if significant structural damage is located in the wings.

[1. Communicate intentions to the ground.]
2. Climb to 10,000 feet above terrain (if practical) at a controllable airspeed.
3. Simulate a landing approach and determine airspeed at which aircraft becomes difficult to control (minimum controllable airspeed).

Note

If aircraft becomes difficult to control or approaches a stall, lower the nose and increase power for recovery.

4. If aircraft becomes difficult to control above 105 KIAS (full flap), fly a no flap landing approach. Abandon the aircraft if it becomes difficult to control above 130 KIAS (no flaps).
5. Maintain 20 KIAS above minimum controllable airspeed or 110 KIAS, whichever is higher, during descent and landing approach.
6. Fly a flat power-on, straight-in approach requiring minimum flare and plan to touch down at no less than minimum controllable airsped. Do not begin to reduce final approach speed until the aircraft has crossed the runway threshold and is very close to the runway. Maximum recommended airspeed for touchdown is 105 KIAS (full flaps), 130 KIAS (no flaps).

Expert System for Postfailure Recovery Control

An expert system is a computer program performing within a specific task domain at the level of a human expert within that domain. Expert systems are generally written to have knowledge represented in discrete identifiable chunks rather than being dispersed throughout. This can be very appropriate for modeling knowledge, which occurs naturally in rule-type (IF-THEN) form. (This also allows the program to be easily built up and modified. Modular-type preliminary development of systems in which strong sequencing eventually develops can be done with the sorts of languages used for expert systems.) It became clear by working with the C-130 that a rule-based expert system could be written to automate the process of iteration to find control to recover the aircraft successfully after elevator jam failures.

The elevator on the C-130 aircraft is a large, highly effective surface, and thus failures of 2–3 deg off-nominal out of a 55-deg total deployment range could result in devastating disturbances. The assumption was made that the failure was fully known. To allow a nominal amount of time for failure identification, emergency control was imposed only after a 3-s postfailure delay. [Relaxing this assumption might even enhance the viability of the expert system approach to finding successful recovering control. It may be effective to initiate recovery control on the basis of vehicle response information only, without or prior to the failure being identified. Rational ways of acting in uncertainty can be embodied in the system knowledge. This is fortunate, since all control failures may not be fully identifiable in the time available before response is required. There are interesting issues to be explored in the focused qualitative/quantitative trajectory projection that some pilots[2] can do to determine that response to a situation is needed, and what likely useful responses might be. It should be noted that automatic control can mask the early effects of a failure if the failure detection mechanism (whether an algorithm or the pilot) does not account for it.]

With the elevator jammed on this aircraft, the nominal pitch automatic control loop was completely inactivated. The C-130 could stabilize in pitch on its own when the disturbance was small. Its longitudinal motion was damped lightly without the elevator loop, however. Flying the handicapped aircraft to landings was complicated whenever any quick transitions, particularly at low altitudes, were attempted. Some reconfiguration of the pitch damping loop to use other controls would be recommended to decrease the amplitude of the oscillations and reduce the time to stabilize in a new operating condition. Not doing so here, however, does not change the findings of this study.

The development of the expert system began with manual recoveries to gain some initial level of expertise, which was codified in the expert system. Then, with the aircraft initially in each of four different initial stages—two straight-and-level flight conditions, a climb, and a descent—the elevator was failed in 1-deg off-nominal increments throughout its entire range, and the expert system was used to try to find a successful strategy for each case. The recovery strategy was considered successful when the aircraft was flown to stabilization in a climb, heuristically a good goal, although sometimes stabilized descent was the best that could be done.[1]

The expert system was built up incrementally using information from successive recoveries. The type of computations involved invited this type of incremental development. System rules were modified or new rules added as cases were encountered for which the system could not find the successful recovery. Only enough detail was added as was necessary. After some number of recovery cases, no additional rules were needed. Much heuristic reasoning was embedded in the rules, as will be partly described.

To increase the control redundancy of this aircraft, the simulation model was modified so that the elevator tab was available for deployment independent of the jammed elevator (thus assuming that the tab would not bend the jammed elevator, changing the vehicle basic aerodynamic and stability derivatives). The ailerons were rendered so that they could be deployed symmetrically ("collective" aileron). Five alternate longitudinal controls were then available—elevator tab, collective ailerons, symmetric flaps, thrust, and landing gear. The simplification of a strict hierarchy of control usage in the recoveries was difficult to avoid. This was part of the heuristic reasoning involved and was carried over to the expert system. The intuitive approach was to try hardover deflection of successively more controls as needed (added in the order[1] elevator tab, ailerons, symmetric flaps, and thrust) and to examine the response after the addition of each control. The practical effect was that the successful control could be bracketed with few tries between too little and too much.

Another heuristic validated by making recoveries and embodied in the expert system was that, in the initial stages of postfailure emergency control, control changes were "optimally" done simultaneously and as quickly as deployment rates allowed. Certain recoveries required that control deployments be reversed later in the recovery sequence. Even then the aircraft could be recovered successfully if controls were commanded to move as quickly as possible to their new settings.

Using this scheme of applying separate controls in order, the successful recovering control fell along a discrete spectrum of strategies, according to the amount of off-nominal elevator deflection associated with the failure. Figure 5 illustrates the spectrum for failures occurring while the aircraft was flying straight and level at 147 KIAS. The successful recovery strategy varied from no compensation to partial or hardover deployment of all of the available longitudinal controls. In no case was more than one later control redeployment necessary. The control change points were successfully chosen heuristically. The expert system also used straightforward interpolation of control deployments in formulating the recovery strategy, the interpolation strategy being successive midpointing between obvious too much and too little deployment of a control.

The rules of the recovery strategy-finding system were codified in a very commonly used rule-based system language[8] and implemented on a personal computer. The expert system computations were quick. The expert system was written to ask the

```
;
;The advisory is beginning.
;A simulation may be terminated when...
;The aircraft pitch angle starts to exceed 111. deg.
;with a pitch-up failure or
;Ground impact is imminent or
;Stable oscillations in all states are apparent
;and you can determine the approximate
;steady climb rate of stabilized flight.
;Answer all questions below with yes or no.
;Please be careful with your answers: there is
;no explicit checking for inconsistencies.
;
;Try no compensating control to see if the aircraft
;can recover on its own.
;Did pitch rate ever cross zero?
yes
;
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
no
;
;Try the addition of immediate hardover elevator tab
;deflection to 50. deg.
;Did pitch rate ever cross zero?
yes
;
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
yes
;
;Did angle of attack remain below 25. deg. originally
;-or-if it did cross this level did it always later
;decrease through it at least before rising above it
;again?
yes
;
;Did angle of attack remain below 25. deg.-or-
;if angle of attack ever exceeded this level did
;pitch rate reach a minimum before the end of the
;simulation and within 5. seconds after angle of
;attack last decreased through 25. deg.?
no
;
;Add immediate hardover collective aileron deflection
;to -20. deg.
;Did pitch rate ever cross zero?
yes
;
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
yes
;
;Did angle of attack remain below 25. deg. originally
;-or-if it did cross this level did it always later
;decrease through it at least before rising above it
;again?
yes
;
;Did angle of attack remain below 25. deg.-or-
;if angle of attack ever exceeded this level did
;pitch rate reach a minimum before the end of the
;simulation and within 5. seconds after angle of
;attack last decreased through 25. deg.?
yes
;
;Did pitch rate cross back from negative to positive
;if angle-of-attack remained below 25. deg. or after
;angle-of-attack last decreased through 25. deg.?
yes
;
;Did flight path angle remain positive in the initial
;stall recovery-or-if it was negative at some point

;did it later cross from negative to positive?
no
;
;Backing off on the last pitch-down input later
;should be tried. In the last simulation
;if pitch rate achieved a
;minimum after angle of atttack decreased to 25.
deg.
;state the time in seconds that this minimum
;occurred. Otherwise state the time at which angle

;of attack decreased through 25. deg. If angle of
;attack never exceeded 25. deg. state the time at
;which pitch rate was minimized for the entire
run.
28.
;
;Back off collective aileron hardover to 0. deg. at
;28. sec.
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
yes
;
;Did angle of attack remain below 25. deg. originally
;-or-if it did cross this level did it always later
;decrease through it at least before rising above it
;again?
yes
;
;Did flight path angle remain positive in the initial
;stall recovery-or-if it was negative at some point
;did it later cross from negative to positive?
no
;
;Back off elevator tab hardover to 3.65 deg. at
; 28. sec.
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
no
;
;Back off elevator tab hardover to 26.825 deg. at
;28. sec.
;Was the aircraft stabilizing by our definition
;at the end of the sim?
no
;
;Did pitch angle remain below 111. deg.-and-if it did
;so but crossed 30. deg. did it decrease and recross
;this level later?
no
;
;Back off elevator tab hardover to 38.4125 deg. at
;28. sec.
;Was the aircraft stabilizing by our definition
;at the end of the sim?
yes
;
;Was the aircraft settling into an apparent descent?
no
;
;A successful recovery strategy has been found.
;End -- no production true
;(64 productions (560 // 1560 nodes))
;(286 firings (987 RHS actions))
;(10. Mean working memory size (18 maximum))
;(1. mean conflict set size (8 maximum))
;(36. mean token memory size (207 maximum))
;(36. "mean token memory size" (207 "maximum"))
```

Fig. 6 Interaction with expert system, −9-deg off-nominal elevator jam, flying at 147 KIAS (emphasis added).

user about the vehicle response to the recovery strategy it recommended last. Answers to questions chosen and presented by the expert system were used to modify the recommended strategy, and the strategy was applied by the user to the simulation. All of the queries about the response were specifically designed only to require information about objective features of the response, and thus the entire process of finding a successful strategy could have been automated. This should help further the idea of a system that could sweep automatically a wide range of failure cases and make generalizations about successful recovery strategies or about inherent control redundancy of the aircraft. It was possible to accumulate empirical guidelines concerning when the results of using a trial strategy were objectively counterindicative early and terminate the simulation early accordingly. In trying to recover the C-130 after elevator failures, a minute of aircraft response time was more than sufficient to see pitch stabilization in progress if it was to occur at all.

Ultimately, the expert system used an extensive set of rules to decide one of three things: the aircraft stabilized successfully, the failure-induced disturbance was underopposed, or the failure-induced disturbance was overopposed. The following system rule is illustrative:

IF failure was pitch down and pitch rate crosses zero
  but later decreases below its first minimum
THEN the failure was undercompensated

The following rule might then come into play:

IF failure was pitch down
  and the failure was undercompensated
  and the emergency strategy simulated does not include elevator tab usage
THEN try hardover elevator tab deflection to −6 deg
  [its maximum pitch-up setting]

Reference 1 contains a complete listing of the expert system program.

Example from Use of Expert System

Typical interface between this expert system and the user is shown in Fig. 6, which has come directly from actual system usage. The user-supplied answers to questions asked by the expert system are shown in italics. The failure was a −9-deg off-nominal (pitch-up) elevator jam occurring when the aircraft was flying straight and level at 1000 ft and 147 KIAS. Six expert system-directed attempts (more than usual) were required to find a successful recovery from scratch in this case. The expert system used general but objective features of the aircraft response to check its latest trial recovery strategy and suggest corrections to it. These features included, in this case, state zero-crossings and the aircraft pitch angle remaining below a level associated empirically with looping and, ultimately, loss of the aircraft. Experience with the manual recovery cases had shown what features of the response constituted successful stall recovery, and this had also been embodied in rules. These were used by the expert system to find a successful recovery strategy in this pitch-up failure case. Control dynamics, including actual surface deployment rates and position limits, were included in the simulation model.

In this example, recovery could not occur unless some type of compensation was applied. Figure 7a shows the result of hardover pitch-down elevator tab deployment after the failure. After reaching a high value, angle of attack did decrease through an empirically based safe value of 25 deg. In successful cases, pitch rate began a trend toward pitch-up within 5 s of this. This did not happen here. A better pitch rate response resulted after full pitch-down collective aileron was added, as Fig. 7b shows. Next, to get the flight-path angle to increase to a positive value and thus for the aircraft to climb, aileron was backed off at a later time (completely) as was the elevator tab (to an intermediate value). There is a highlighted exchange in the transcript pointing to the surmisal that some pitch-down compensation must be relieved after pitch rate recovery. Three adjustments were required to obtain acceptable final aileron and elevator tab settings. Figure 7c shows the final successful strategy and the resulting vehicle response.

Only elevator tab and collective aileron deployments were needed to control the disturbance created by this particular elevator jam failure. However, to find a successful strategy in more strenuous cases, the possibilities of also changing flap and throttle settings during the initial part of the recovery, and later reversal of these changes, had to be incorporated into the expert system. In cases where the aircraft stabilized in pitch, but in a descent, the expert system could first recommend initial thrust increase but could backtrack when this prevented stabilization in pitch. Then, using heuristic-based information, it could recommend times and amount of delayed thrust change to get transition to climb. Flap deployments in the simulation were limited structurally by airspeed. Expert system rules were written to decrease thrust or extend landing gear to keep airspeed down so that flaps could deploy fully, if necessary. From a problem-solving viewpoint, keeping airspeed up was a subgoal to the goal of pitch-stabilizing the aircraft. Landing gear, if extended, could be raised to help in gaining altitude during the recovery. By the end of the study, the expert system gave successful emergency control, when it existed, for all ranges of elevator jam failure and for all of the varied initial conditions tested.

An expert system such as this one might be programmed to find recoveries after failures of other types of controls and perhaps after multiple control failures. This type of system might be able to find the minimum altitude loss strategy, which might be crucial in recovery, through incorporation of (probably much more elaborate) bracketing of control strategy.[1] It might be possible to use this type of system to achieve other goals at the end of the recovery, e.g., optimum climb speed. Through certain additions, the recovery stratetgy need not be started from scratch, thus decreasing the number of attempts required. Other types of extensions might involve including rules to refine a workable strategy. Then one could reduce deployments of a certain control, redistributing its part of the failure compensation, in order to hedge against future control needs or future additional failures. Similarly, one might be able to include rules to simplify a strategy with numerous control change points, as might be expected in dealing with higher performance or reduced stability aircraft. Systems of this type might be taught to find recoveries where coupling between control of different axes is involved; e.g., where temporary rolling to let the aircraft nose fall through

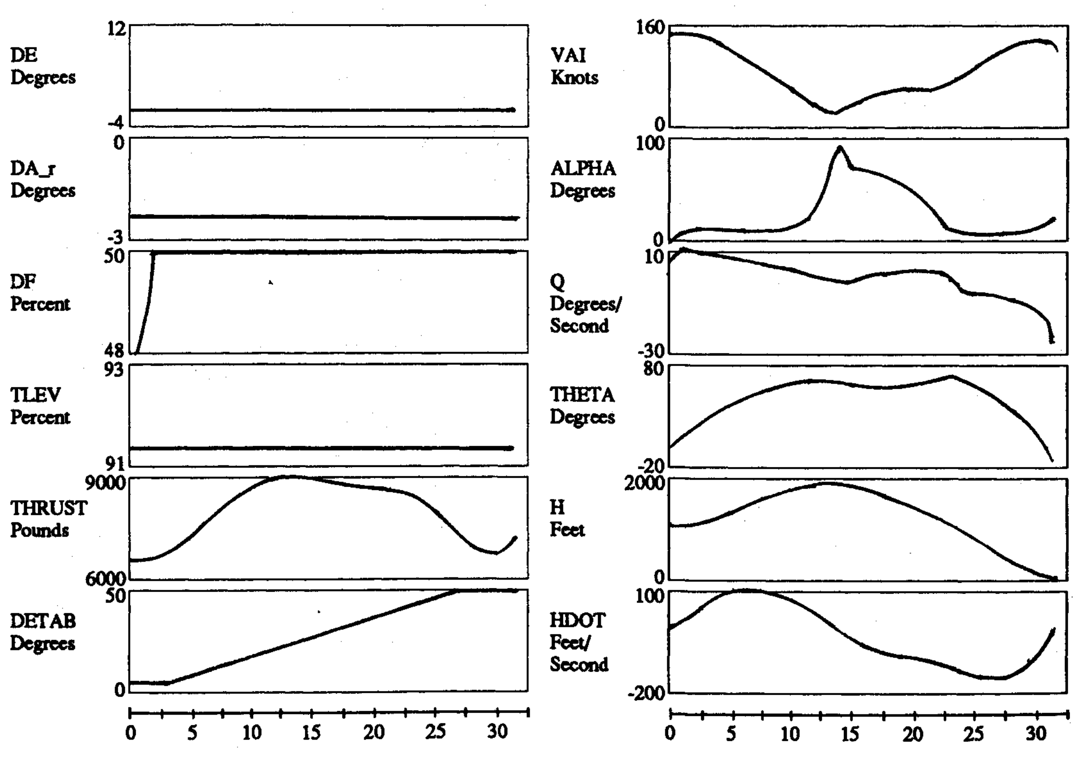

Fig. 7a C-130 aircraft response to −9-deg off-nominal (pitch-up) elevator jam, with expert system-directed compensation; strategy 1 (flaps originally 50%, extension airspeed-constrained; DE—elevator, DA_r—right aileron, DF—flap, TLEV—thrust lever, DETAB—elevator tab, VAI—indicated airsped, ALPHA—angle of attack, Q—body axis pitch rate, THETA—body axis pitch angle, H—height above ground, HDOT—altitude rate).

would be required to recover from large pitch-up longitudinal control failures. If some near-optimum or difficult recovery strategy were being sought, the iterations to find a successful postfailure control strategy would be tedious at best for a human pilot. An automated system could be very sensitive to improvement in the recovery. The system developed was insensitive to specifics about initial condition or delay for failure identification. It changed the strategy using only information about the vehicle response to it.

This type of system could easily be used in conjunction with (possibly reconfigured) automatic control. If an autopilot loop were still engaged, the recovery-finding system would work independently of what control changes were commanded by the autopilot, effectively augmenting or overriding its commands, as necessary. Figure 8 suggests this. Experience with the C-130[1] showed that part of finding a workable recovery strategy could naturally involve determining autopilot usage. Sometimes disengaging the autopilot in one axis was very beneficial in helping the aircraft recover. However, this would be more helpful than necessary if the expert system itself had high authority to command control changes. This type of expert system might be useful in giving some rough idea of the

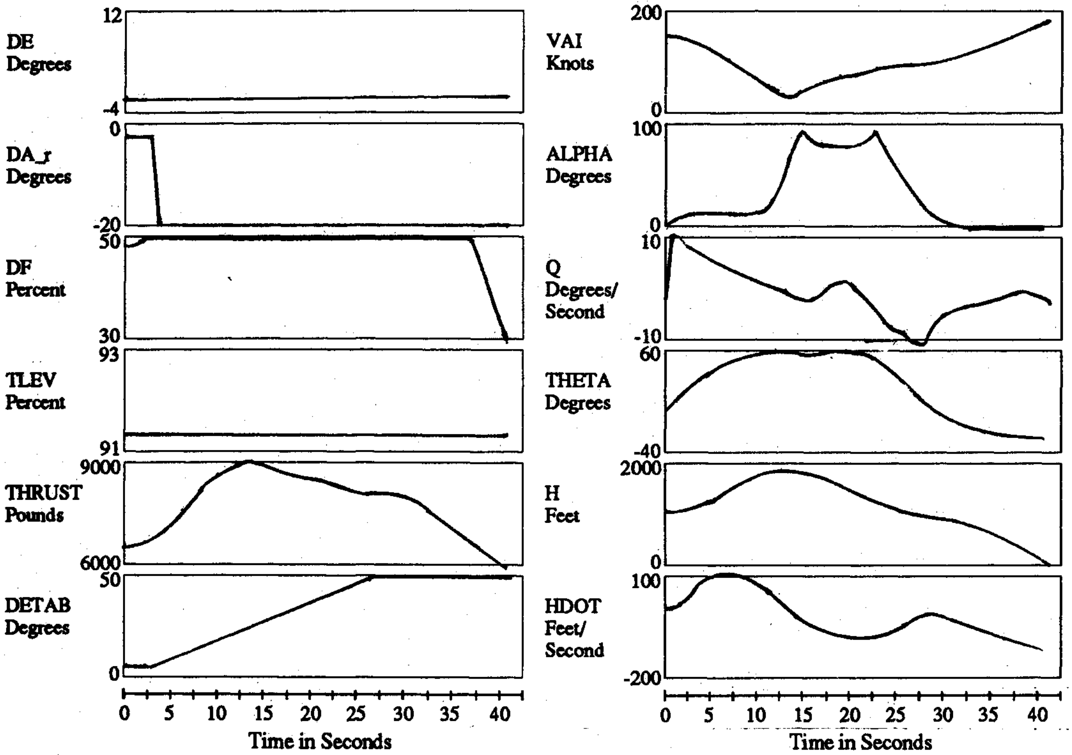

Fig. 7b C-130 aircraft response to −9-deg off-nominal (pitch-up) elevator jam, with expert system-directed compensation; strategy 2.

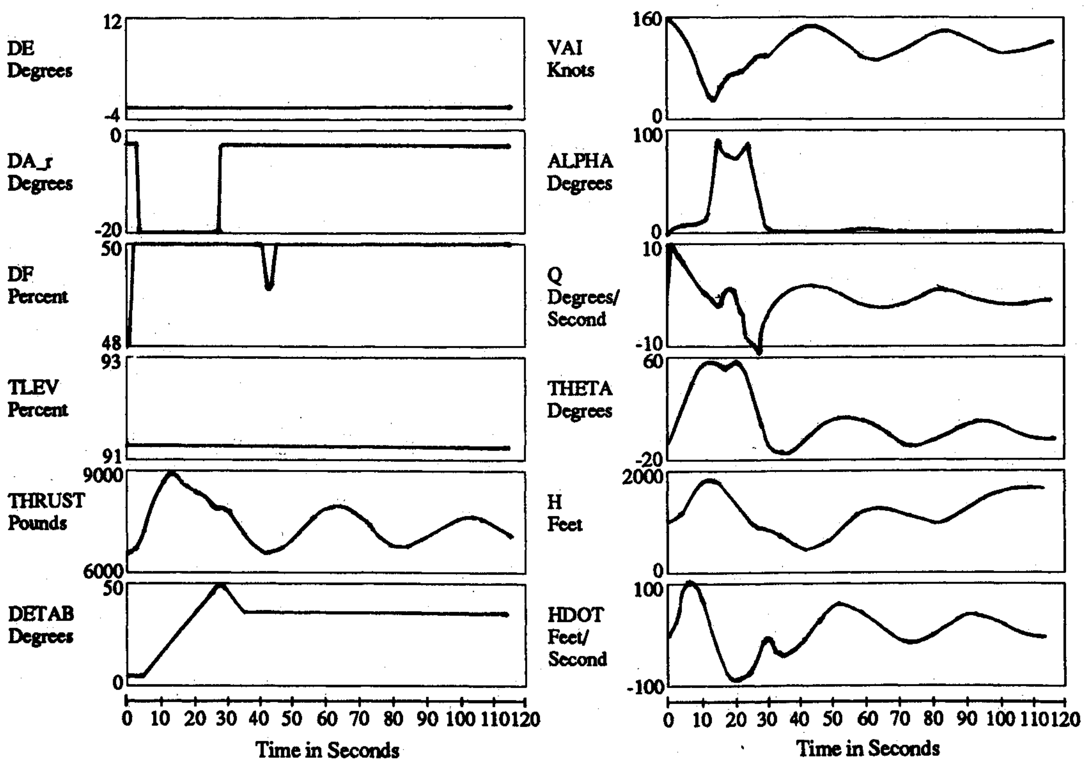

Fig. 7c C-130 aircraft response to −9-deg off-nominal (pitch-up) elevator jam, with expert system-directed compensation; strategy 6; successful recovery from failure.

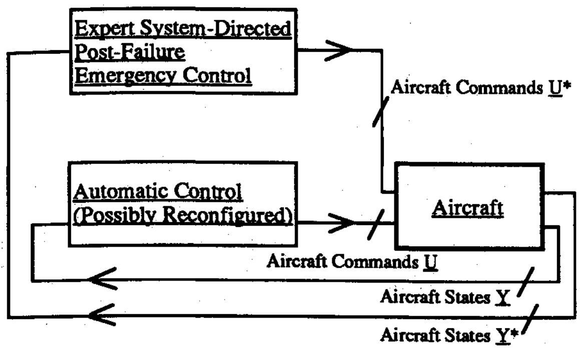

Fig. 8 Expert system to find emergency control used in conjunction with (reconfigurable) automatic control.

basic redundancy of an aircraft configuration alone. If no automatic control were modeled, the expert system might look for control strategies that yielded rate and state zero-crossings but without stabilization being required. It is difficult to establish true control redundancy without some type of outright simulation of the vehicle response to different failures occurring with the vehicle in various initial conditions.

An expert-type system could incorporate many different types of information needed to formulate a recovery strategy. The need to incorporate postfailure constraint information, at least implicitly, should be illustrated. In addition to working with elevator failure cases, many recoveries after C-130 asymmetric jam flap failures were made in this study. In severe flap failure cases, highest priority efforts had to be made to keep the airspeed above the calculated controllability airspeed discussed earlier, so the ailerons would have sufficient authority to oppose the rolling induced by the failure. Increasing thrust alone was not sufficient to keep the airspeed up in these cases. Instead, and counterintuitively according to pilots interviewed, temporary forceful pitch-down of the aircraft using the elevator and thrust increase in the dive to enhance the acceleration had to be part of the recovery. The advantage of the expert system approach is that any counterintuitive recoveries known to be successful could also become part of the knowledge embedded in rules in the system. In a similar way, calculating postfailure performance-type constraints could also become important in compounding a recovery strategy. As one reported accident case demonstrated, part of a successful recovery from a jammed extended surface might involve a strategy that deliberately attempted to exceed structural limitations on the surface, so that it could be broken off. If there were no choice, temporary moderate relaxation of some vehicle constraints might be undertaken. All of these possible aspects of a successful recovery could be embedded in an expert system. Force-moment mapping reconfiguration or other types of postfailure loop reconfiguration could not be expected to produce some of these successful recovery control strategies.

In-flight use of an expert system of the type developed might be very feasible, but it might also be used for finding recoveries from wide ranges of control failures before an aircraft is ever flown. Although it might not be possible to simplify this much, the final result of a broad simulation study might be a failure recovery system with processing consisting of straight-through paths from information about the failure to the successful recovery strategy:

The control failure
(what has failed and how)

Indicated airspeed

Altitude
Other initial states
(as required)

Initial control settings

} → Recovery strategy

The viewpoint that resulted in this type of system for finding emergency postfailure control was quite different from that taken in most failure-accommodating control research, although it is not at all incompatible with it. The advantages are numerous: systems like the one developed here can be written to call for use of unusual controls in unusual ways; they can direct the perhaps counterintuitive strategies that are sometimes required; they can accommodate saturation of controls in a situation in which it can be expected and is very important; and they can include the possibilities of calculating and using information about postfailure operating and performance constraints. This approach is an effective blending of qualitative and quantitative approaches to emergency control.

Integrated Recovery and Advisory System

Let us consider an integrated onboard aid and advisory system. The most desirable system for automated recovery from control failures would have the capability of making an immediate, correct response to failure. The emergency control strategy could be formulated by an expert system such as that presented. The pilot would be informed of all automatic control actions, but the system should relinquish control once the aircraft was stabilized. (An explanation must be available on demand for advice given by an expert system, according to traditional AI thought.) There would likely follow phases of relatively quiescent flight, and the transitions between flight phases would generally be done slowly and carefully. This was evident in the few successful reported failure cases and is the result of the pilot's natural inclination to reduce performance demands on a handicapped aircraft. Based on the considerations already established, a system should be provided to support the pilot with advisory information for the rest of the flight. The ideal pilot advisory system interface should have the most important information displayed continuously. Other information about each general flight phase should be kept updated and available to the pilot on demand.

A postfailure advisory system must be able to decide which types of information might have changed significantly after the failure. The goal is selective substitution in the pilot's knowledge. Piloting is an expert behavior and has evolved naturally to being very efficient from an information standpoint. It makes sense to evaluate and present information that pilots would naturally choose, and to try to use their heuristic assessments of what information needs to be recalculated. When faced with a significant failure, a good pilot would know when to evaluate for more precise information, but on the basis of a more qualitative assessment of the situation. The knowledge involved in supporting the advisory would probably be more broad than deep, and probably not extensive. The overall aid/advisory system would probably be a hybrid blending of qualitative rule-based processing directing quantitative computation. This again suggests use of an expert system approach to implement the advisory system. Most languages used for expert systems have a mechanism for calling functions in a standard programming language as part of a rule's THEN actions.

Pilots suggested the following questions about residual control capability as those they would ideally want answered:

1) How much roll capability remains?

2) How much control do I have over vertical acceleration, angle of attack, airspeed, and sink rate?

3) How much sideslip can I use?

4) Am I "committed to land?" Can I put the flaps down for a landing?

Answers to these questions can be impossible to determine or even convey. Any strong notion of "control capability" must be a function of the aircraft operating state. Even at a given state, it is a function of the available range and allowable rate of deployment of the viable controls, the structural limitations on their deployment, and their "artifact." Furthermore, use of a control may be limited by operating constraints perhaps induced by the failure. These types of considerations apply in nominal aircraft operation, but then available control capability need not be assessed precisely, because what is feasible is generally "known" by the pilot from training and experience.

Fortunately, when pressed, pilots wanted less to be apprised of the new values of the quantitative maneuvering figures of merit as to know whether or not the capabilities are significantly degraded and the configurations at which they are optimized. Some general sorts of reminders, such as "avoid (adverse) sideslip" when there has been an engine failure, were also wanted. Pilots want to be able to use standard or standard emergency procedures (which have been practiced and are more automatic) as much as possible. An advisory system should be able to list and reference standard flight manual information and checklists, although, as seen in the first section of this paper, this sort of advice probably should be augmented with quantitative advice assembled from more fundamental information.

Pilots are accustomed to watching indicated airspeed closely. Most flying is done according to airspeed guidelines, and changes in these that are judged large according to some measure by the advisory system must be put forward very clearly. Certain background calculations anticipating possible additional failures should also be performed. Reference 9 suggests that some of the most important pilot advisory information of the type being recommended here be superimposed on a head-up display when possible and gives suggested display formats.

Reference 1 lists suggested flight phase advisory information for the postfailure flight phases identified earlier, after control of the aircraft has been regained through automatic assistance. It includes information on how the need to assemble some advisory information might be established and programmed. One of the most important considerations is that the failure of a given control results not only in the loss of at least some of its own controlling functions but, possibly very significantly, in the degradation of those of the failure-compensating controls or changes in state (such as decreased angle of attack for pitch down effect). Calculations to support postfailure advising need to be keyed to both failed and compensating controls, particularly compensating primary controls, and compensating changes in state. The advisory system should contain enough rules to be able to judge when to make the inference that the extent of off-nominal control or state is such that a given aspect of aircraft performance, for example, must be evaluated explicitly.

Returning to the issue of explicit retrim, the trajectory an aircraft is flying might be achievable with greater efficiency or with more potential for maneuvering reserved. This could arise in any flight phase and seems particularly important in flying a control-impaired aircraft. The pilot might not often be able to ask the system for better trim points. The onboard aid and advisory system should probably infer needs on the basis of current apparent steady-state flight. In some cases, upcoming needs might be inferred and retrim calculations proceed on that basis. The possibility of doing some rational type of continual automatic background retrim should be considered.

Being able to implement the proposed recovery aid and advisory system would probably not require unrealistic computational and display capability. Its goals and features would mesh extremely well with the aircraft installations being developed in current programs to design cockpit aids for fighter and transport aircraft.[10]

Conclusions

Analysis of aircraft accident cases, discussions with pilots, and experience flying simulated recoveries have shown that existing research on the problem of control failures has not yielded entirely adequate answers. Familiar failure-robust or reconfigurable control ideas have discrete limitations in finding recoveries even when they exist. Failure-induced changes in operating and performance constraints cannot be taken explicitly into account. These approaches cannot utilize all remaining control resources in the complicated, unusual, or even counterintuitive ways that were shown here to be sometimes required. A different, wider perspective was developed. This perspective was embodied usefully in the expert system idea, which became the backbone for a proposed automatic aid and advisory system. Piloting is an almost classical expert behavior.

Types of failure-induced changes in constraints have been discussed. These constraints, although rarely considered in other research, can become crucial in formulating emergency control strategies and advising pilots flying control-impaired aircraft. Elements of the proposed automatic emergency control system have been discussed. A rule-based system to find successful control after elevator jam failures on the C-130 aircraft was developed and has been considered extensively. Systems of this type seem to be very effective, and extensions to this approach have been presented. They can conceivably be installed onboard an aircraft, and, as is detailed, could be important resources on the ground. This type of system can and should be used in conjunction with more familiar failure-robust or reconfigurable automatic control.

The advisory function of the onboard aid and advisory system is seen to follow usage of full-authority postfailure automatic control. There has been a treatment of what types of information about continued flight must be calculated. The rule-based approach can provide a way of determining useful qualitative advice about flying a control-impaired aircraft and deciding which types of more quantitative information must be calculated. There was also a discussion of postfailure retrim and some demonstrations of the impact of this which are believed to be novel. Numerous interesting and substantiating concrete examples for all of these points have been included. The automatic aid and advisory system proposed here, including rule-based systems for find emergency control and directing advisory calculations, seems to promise a good, viable solution to the problem of aircraft control failures.

Acknowledgments

A National Science Foundation Graduate Fellowship is gratefully acknowledged, as are the facilities and financial support provided by the Charles Stark Draper Laboratory and support from the Office of Naval Research (Contract N00014-82-K-0582) and NASA Ames and Langley Research Centers (NASA Grant NAG-2-297). Lena Valavani and Wallace Vander Velde were research supervisors.

References

[1]Wagner, E. A., "On-board Automatic Aid and Advisory for Pilots of Control-Impaired Aircraft," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA, Dec. 1987.

[2]McMahan, J., "Flight 1080," *Air Line Pilot,* Vol. 47, July 1978, pp. 6–10.

[3]Montoya, R. J. (ed.), "Restructurable Controls," NASA CP-2277, Aug. 1983.

[4]Chandler, P. R., "Self-Repairing Flight Control System Reliability and Maintainability Program Executive Overview," *Proceedings of the National Aerospace and Electronics Conference,* Vol. 1, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 1984, pp. 586–590.

[5]Waterman, D. A., *A Guide to Expert Systems,* 1st ed., Addison-Wesley, Reading, MA, 1986.

[6]Cohen, P. R., and Feigenbaum, E. A. (eds.), *Handbook of Artificial Intelligence,* 1st ed., William Kaufmann, Los Altos, CA, 1982.

[7]*Flight Manual, USAF Series T-37B Aircraft,* T.O. IT-37B-1, Secretary of the Air Force, Washington, DC, July 1982.

[8]Brownston, L., Farrell, R., Kant, E., and Martin, N., *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming,* 1st ed., Addison-Wesley, Reading, MA, 1985.

[9]Weinstein, W., "Control Reconfigurable Combat Aircraft Development, Phase 1—R&D Design Evaluation," Flight Dynamics Laboratory, Wright-Patterson Air Force Base, OH, Rept. AFWAL-TR-87-3011, May 1987.

[10]Kinnucan, P., "Building a Better Cockpit," *High Technology,* Vol. 6, Jan. 1986, pp. 49–54.