

The Crash of AA587: A Guide

Peter B. Ladkin
RVS Group, University of Bielefeld
ladkin at rvs.uni-bielefeld.de
Report RVS-RR-04-03

November 18, 2004

1 The Accident

American Airlines Flight 587 (AA587) was an Airbus A300-600 which encountered some wake turbulence from a preceding aircraft on climb out of JFK New York airport in 2001. During the second wake turbulence encounter, the rudder experienced 5 full deflections in opposite senses in about 7 seconds, and the vertical stabiliser (the fin at the back of the aircraft) broke off at its root due to overload. Control of the aircraft was immediately lost and it crashed into houses in the Belle Harbor area of New York City, in the borough of Queens, killing over 250 people on board and five people on the ground.

The US National Transportation Safety Board (NTSB) public hearing on the crash of AA587 on November 12, 2001, was held on October 26, 2004. Such a public hearing consists more or less of a presentation of a draft of the final report, in particular the conclusions (Findings, determination of probable cause and contributing factors, Recommendations), and comments from attendees. The NTSB's preparatory work and presentations at the hearing were summarised in [DF04, Fio04].

This note considers some of the technical and sociotechnical aspects of the accident uncovered by the three-year investigation.

2 The Interest of the Accident

The accident has been complex to unravel, and the institutional stakeholders still do not agree on the importance of various causal factors, indeed even on whether certain factors are causal at all. Despite initial suspicions, the accident turns out to have no direct digital-computational involvement, except in the increased ability nowadays to calculate the exact aerodynamics of the event retrospectively. Its interest to students of complex heterogeneous systems stems from two themes:

- Questions about control-oscillatory behavior, and

- Sociotechnical (Soctech) systems aspects.

2.1 Control-oscillatory behavior

Feedback control systems tend to oscillate about a goal; they control to reach it, overshoot, reverse control to reach it again, overshoot, and so on. Oscillations can be – normally are – *damped* when the amount of overshoot (the *amplitude*) gets smaller with time and they eventually stabilise at the goal; or *undamped*, in which case the amplitude remains more or less constant and one must undertake other action to exit the oscillations and reach the goal; or *divergent*, in which the amplitude of the oscillations increases – the overshoot gets larger. Divergent oscillation can be dangerous to a structure, for it may be caused eventually to exceed its design strength limits through egregious motion. And divergence can happen slowly, or very rapidly. In the case of AA587, it was seven to eight seconds from the start of the disturbance to when the fin broke off.

In some control systems, such as maintaining heading on autopilot, mild undamped oscillations either side of the theoretically perfect track are normal. One worrisome group of divergent oscillations involves pilot control actions in response to perceived aircraft behavior. Some such events to digital flight control systems have been reported for example in the on-line Risks Forum [Var89] [Var93]. The A300-600 does not have a digital flight control system, but it does contain digital systems such as the autopilot which uses the primary flight control when switched on. It also has certain automatic control-damping systems such as the yaw damper which has digital components.

AA587 exhibited divergent oscillations in yaw, leading within a few seconds to overload of the fin, which broke off. Did it happen all by itself, or was the pilot involved? The question of identifying situations conducive to divergent pilot-aircraft control coupling, in which pilot input in response to aircraft behavior initiates divergent oscillatory flight behavior, was identified as the most pressing problem facing digital flight control systems by a National Research Council committee in 1997 [Nat97] – far more pressing, indeed, than the issues of program and algorithm correctness which have traditionally occupied most computer scientists working on DFCS issues. The coupling of pilot input and aircraft behavior through the control system is traditionally known as pilot-induced oscillation (PIO). The NTSB uses the term aircraft-pilot coupling (APC) in the AA587 documentation. The oscillatory behavior in this accident was the subject of a special report by Professor Ronald Hess, a member of the NRC APC committee [Hes03]. I use the term PIO, pilot-induced oscillation, as does Hess. (Hess warns against the inference that use of this term means that the pilot is somehow to be faulted [Hes03, Executive Summary,first sentence].) I consider this issue in detail below.

2.2 SocTechSys aspects

There are two, maybe three, major concerns that stem from the complexity of what some like to call “system issues”. The “system” meant here is the

inhomogeneous conglomerate formed from manufacturers, regulators and users of transport aircraft, including those that fly them, as well as the kit itself. The phase “system issues” is meant to designate aspects of the whole interaction which are not generally visible when one considers only the constitutory parts as separate entities.

First, astonishing as it may be, it turns out that many pilots, indeed some airlines, did not properly understand the design limits nor the intended use of rudder in large commercial airplanes, nor the basis for certification of a rudder design, and were using it, even being trained to use it, in ways which could result in seriously overloading the vertical stabiliser (“fin”). As indeed happened in this accident.

Such a general miscommunication we can call a three-and-fourpence issue. A joke relates a desperate battle in a war, in which the front-line general sent a message to headquarters to “Send us reinforcements. We’re going to advance”. Relayed verbally from courier to courier, it arrived as “Send us three and fourpence, we’re going to a dance” (three and fourpence was three shillings and four pence in the pre-1972 English currency. The joke is old).

A public NTSB hearing was also held at the end of October 2002. Aviation Week reported then that until the AA587 crash “most transport pilots were unaware that rapid rudder reversals could cause the tail to rip off the aircraft” [FD02]. Indeed, AvWeek’s inquiries had shown this to be the case at the time of the Feb 8, 2002 release of two NTSB recommendations regarding such awareness and training programs [Fio02b]. At the October 2002 hearings, the NTSB “zeroed in on why and how this crucial information, familiar to flight test engineers and manufacturers, is not included in pilot upset training programs” [FD02]. A “rudder reversal” is the rapid movement of rudder from full displacement in one direction to full displacement in the other. It should normally never occur.

The NTSB looked at many potential high-fin-loading events on A300-600 and related aircraft [Dor02c]. In particular, during an incident in 1997, an American Airlines A300-600, Flight 903, had used rudder to recover from a stall [Fio02a] and had undergone a rudder reversal under pilot control. Pilot use of the rudder had overstressed the fin perhaps to beyond Ultimate Load (see below for definition). Use of rudder at low speed, at least in the circumstances of a stall such as this, is foreseen by the certification basis. Rudder use at the airspeed of AA587, some 250 knots or so (1 knot = 1.15 mph), is only foreseen in the case of an engine failure, to keep the airplane straight despite a yawing moment from the asymmetric thrust, and during which the expected loads would be much lower than those generated by a full opposite rudder deflection during sideslip [Dor02b]. Indeed, “both Airbus and NTSB say Flights 903 and 587 are unrelated except for the pilot training issue. Flight 903 was a loss-of-control issue. Rudder input was necessary on 903, unnecessary on 587” [NTS04, Fio04].

One delicate soctech-systems miscommunication aspect concerned how earnest Airbus’s advice on rudder use was after Flight 903 (should they, or the FAA, or someone have been more insistent?), and how American Airlines reacted to that advice. Although American Airlines was reported to have denied that their training program emphasised use of rudder in recovering from upsets [DF04],

the NTSB found that they had done so [U.S04b, Conclusion 8]. Indeed, at the public hearing in October 2002, it was noted that NTSB member John Hammerschmidt had taken American’s AAMP course about the same time as First Officer Molin, the pilot flying AA587, and said that his “AAMP workbook states that rudder becomes primary roll control at high angle of attack” [FD02]. Also, there is the simple question of habit. “A line check airman with a major airline noted that some pilots, when practicing roll upset recovery in the simulator, will “fight initially with aileron, but at the panic point they tend to stomp rudder....” ” [Dor02a].

Another “systems” issue may be seen by some in the discrepancy between vertical stabiliser strength certification criteria, and loads that can be attained in flight through particular use of the rudder by the pilot. Indeed, a pilot may overstress the fin to breaking through rudder reversals. But one should bear in mind the advice of a “flight controls expert at a major airframe company” (i.e., Boeing or Airbus) that “I think you can break an aircraft in any axis if you work on the controls. Operating aircraft relies on basic airmanship” [Dor02b]. Indeed so. That is why airline pilot training is comparatively lengthy and rigorous.

Yet another “systems” issue can be seen in the following. The “design manoeuvring [sic] speed” of an aircraft, V_a , is defined in a technically complex manner in 14 CFR 25.335(c) (that is, Federal Aviation Regulations 25.335(c)) [U.S]. The definition does not allow easy translation into pilot do’s and don’ts. However, most pilots, including myself, are or were taught that the manoeuvring speed of an aircraft, V_a , is the speed below which one can make full and abrupt control movements without endangering the structural integrity of the aircraft. Not always so, if we are speaking of transport aircraft rudders. V_a for the A300-600 is higher than the 250 kts of AA587. The NTSB says “There is a widespread misunderstanding among pilots about the degree of structural protection that exists when full or abrupt flight control inputs are made at airspeeds below the manoeuvring speed” [U.S04b, Conclusion 16].

3 Some Issues

I enumerate issues raised by observers after the accident and persisting until recently:

1. whether the turbulence encounter was of unusual strength or nature, sufficient to cause extreme aircraft yawing, with the measured side loads of 0.4g;
2. whether the rudder oscillations were caused by the pilot moving the rudder pedals, or whether they were caused by a system flaw (such as discovered during the investigation of a crash of a US Airways B737 in 1994);
3. whether the use of composite materials, rather than metal, for the fin of the aircraft is appropriate;

4. whether the rudder-pedal control design is appropriate (it is different from that on some, but not all, Boeing transports)
5. whether American Airlines trained use of rudder to control yaw during wake turbulence encounters and other disturbance situations
6. whether, and how, the airline had been advised against this practice of using rudder to control yaw during wake turbulence
7. how widely it was understood as industry best practice not to use rudder to control yaw except in certain low-speed flight regimes
8. what the certification standards were for vertical stabiliser strength, and whether these were adequate

The NTSB draft findings [U.S04b] help settle these as follows, with some background added from the reports in the public docket [U.S04a], the reports of the journal *Aviation Week*, as well as some interpretation from me.

3.1 The Wake Turbulence Was Normal

The NTSB noted at the October-November 2004 hearing that “Flight 587’s encounters with wake turbulence were “unremarkable” and aircraft was not in or at risk of an upset event” [Fio04]. During the course of the investigation, the NTSB did engage the services of a number of meteorologists and air turbulence specialists at NASA [Fio02c]. The FDR-recorded side loads of 0.3-0.4g early in the oscillation period were determined to be accurate and to have come through rudder use at those points [Dor01].

3.2 The Pilot Flying Moved the Rudder Pedals

The NTSB determined that the rudder movements were caused by the pilot flying (PF), First Officer Molin, activating the rudder pedals [U.S04b, Conclusion 5]. The flight data recorder (FDR) records control position, but not such things as force applied, and in some ways the cockpit controls are “back-driven”, that is, are moved by the system in response to actuation of the aerodynamic control surfaces of the airplane [Fio01].

The question arises because of the divergent oscillatory nature of the rudder, and the rudder pedal position indicator, in the last seven seconds or so before the vertical stabiliser (“fin”) snapped off. The question whether the PF had done it, or whether it was a system flaw (there are devices such as a “yaw damper” that move a rudder without pilot control input) had still been open at the time of the October 2002 public hearing, almost a year after the accident, although that hearing “proceeded with the implicit assumption that a crewman was moving the rudder” [FD02]. We may presume that the investigation identified no potential system malfunction of the sort that would explain the rudder-reversal behavior.

The cockpit voice recordings (CVR) and the flight data recorder (FDR) indicate that the aircraft was rolling at the time, due to encounter with wake turbulence from a previously departed aircraft. PF was also using control-wheel inputs to counter the roll [U.S04b, Conclusion 12]. The FDR traces were too sparsely sampled (at 2Hz, twice a second) for easy reconstruction of such rapid movements; nevertheless the curves were reconstructed through analysis. Figure 1 shows the NTSB’s reconstruction from [O’C04].

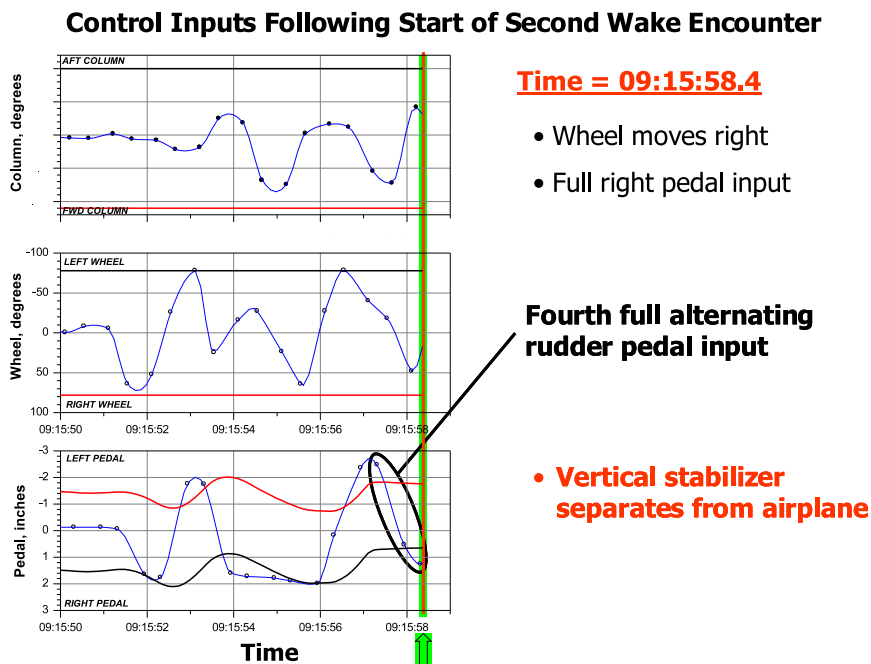


Figure 1: FDR (Reconstructed) Traces of Control Wheel and Rudder Pedal Position [O’C04]

The question arises why PF would have been using rudder to control roll movements; primary roll control is via ailerons and/or spoilers on the wings, activated by the control wheel. There is a control coupling between the yaw axis and roll axis of a traditionally-designed aircraft, so that by putting in yaw inputs, one can affect roll. In small aircraft, use of rudder is important during turning, to control the adverse yaw caused by roll inputs, but in larger aircraft this is negligible, and in any case controlled by devices such as yaw damper. “Most pilots have never used the rudder above 200 kt. and are unfamiliar with its characteristics at higher speeds” [Dor02a]. A roll response to use of rudder in a large transport aircraft occurs with some time delay (hysteresis), which can be of the order of large fractions of a second to some seconds in the case of a

large aircraft such as this.

The October 2002 hearing heard evidence from Captain John Lavelle, who flew a number of times with Molin in 1997-1998. He recalled Molin aggressively using rudder in an oscillatory mode on a B727 to attempt to control roll. This action didn't level the wings, but it did create sideloads due to yaw. They discussed this during the flight. Lavelle recalls that Molin told Lavelle that he was following American Airlines training procedures [Eva02, FD02].

The NTSB concluded that the PF tended to react to wake turbulence using "excessive control inputs" [U.S04b, Conclusion 7], and that his control inputs in this case were "too aggressive" and the degree of his initial rudder pedal activation was "unnecessary to control the aircraft" [U.S04b, Conclusion 12].

3.3 The Fin Met Its Specification

The vertical stabiliser met its design goals and certification standards [U.S04b, Conclusion 6]. Limit Load is defined as the maximum load to be expected in operations; Ultimate Load is 1.5 times Limit Load. The notion of Ultimate Load is based upon the expectation that the structure will fail at higher loads than this. In fact, the aircraft's fin separated from the body of the aircraft at a force of 1.93 times Limit Load, which is some 29% higher than Ultimate Load. Compared against the certification criteria, the composite fin is twice as strong as it "needs" to be and 30% stronger than it is required to be. The question of that word "needs" will be revisited in Section 3.7, where I discuss whether the certification basis is adequate.

Some suggested that use of composite materials rather than traditional metal alloys could have been a factor; and early in the investigation some argued for all A300-600 aircraft thereby to be grounded [Fio02c]. Indeed, the composites broke differently from the way that metal alloys would fail. But substantial investigation showed no unknown difficulties in the use of composite materials. Besides, if a structure breaks at 1.3 times the load at which it is "allowed" to break, one can hardly claim that it was not fulfilling its defined purpose adequately. It is just as likely that one could cause a metal fin to separate under similar circumstances, but for obvious reasons no one has had the temerity to try, and simulating it in static tests would be expensive, difficult, and likely only approximate.

3.4 The A300-600 Rudder Pedal Feedback Design

Opinion continues to diverge among "stakeholders" on the appropriateness of the A300-600 rudder pedal actuation design.

3.4.1 Intended Rudder Use and Certification Criteria

As noted above, the rudder is not a primary flight control. It is used on large commercial aircraft primarily:

- to aid in directional control and stability in turbulence or wind gusts during take-off, and
- to counteract yaw due to asymmetric thrust if an engine fails, which asymmetry can be quite large on a big two-engine airplane such as the A300-600.

These criteria are implicit in the certification standards. Apart from this, it is also used

- in certain landing techniques just before touchdown, at speeds similar to those at take-off;
- in certain circumstances in upset recovery, as for Flight 903 noted above.

The third criterion is an aerodynamic situation similar enough to the first as not to constitute a substantially different requirement.

An encounter with these phenomena would lead a pilot to use rudder in one direction; accordingly the certification criteria are based on a pilot “stomping” on rudder, holding it there, and then releasing it abruptly, representing the most extreme control input possible in response to these phenomena. Rudder reversal is foreseen neither as a pilot response to the above phenomena nor in current certification standards.

3.4.2 Limiting Rudder Travel at High Speed

All large aircraft have rudder-travel limiters, so that the rudder can only move a few degrees when the aircraft is travelling at high speed (otherwise you’d rip off the tail). All modern large transport aircraft use assisted control surface activation, so there arises the question of how you feed the activation back to the rudder pedals: the feedback is necessarily artificial and is referred to as “feel”. It must be explicitly designed.

Some aircraft have rudder pedal movement proportional to the allowed rudder travel. That is, the rudder pedals have constant range of movement over all flight regimes: you move the pedals 75%, and you get 75% of whatever the allowed travel is in that phase of flight: if the travel is 15deg (slow flight) then you get 75% of 15deg; if it is 10deg (fast flight) then you get 75% of 10deg. We can call this system a “ratio changer” system and illustrate the mechanical principles in Figure 3.4.2.

The Airbus A300-600 has pedal movement proportional to absolute travel, so it requires the same movement to move the rudder 10deg at any speed. The rudder pedals are physically prevented from moving further than this, by stops. When the rudder is limited in movement at high speed, then the stops so limit the pedal travel. We can call this the “variable stop” system and illustrate the principles in Figure 3.4.2. The stops are located near the rudder, at the rear of the aircraft, and not near the pedals at the front. This will play a role in later considerations, when the effects of elasticity in the connection from pedals to stop are considered. Rudder travel on the A300-600 is ± 30 deg up to 165 knots indicated airspeed, then it begins to reduce down to ± 3.5 deg at 395 knots

Ratio Changer System

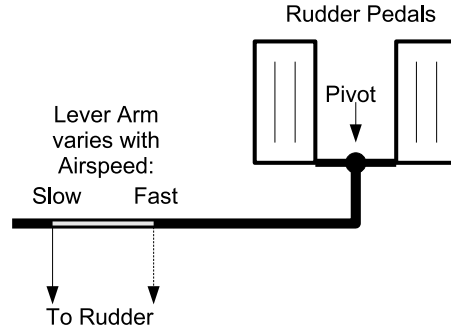


Figure 2: Ratio-Changer Rudder Limiting System, after [Dor02a]

indicated airspeed [Dor01]. At the 250 knots airspeed flown by AA587, the rudder travel limit was 10-11deg [Dor02b].

Variable Stop System

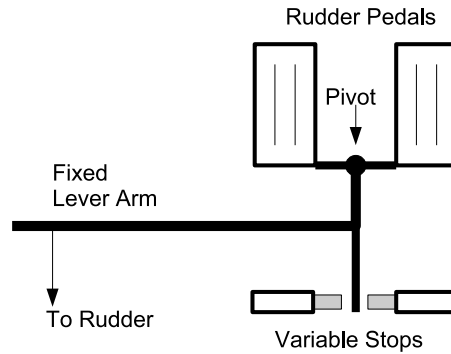


Figure 3: Variable-Stop Rudder Limiting System, after [Dor02a]

[Dor02a] discusses the various rudder travel limiting systems and their “sensitivity” (yaw acceleration per incremental pedal force). There has been considerable discussion over the ergonomics of these feedback methods. But debate has been clouded by partisanship.

At 250 knots (miles per hour times 1.15), the rudder pedal travel on the A300-600 is 1.3 inches, and requires 32 lbs per foot (lbf) of pressure [Eva02]. The B727 flown by Molin with Lavelle in 1997 has a similar 1.3 inches of travel at a similar speed, but requires 40 lbs of pressure [Eva02]. The “breakout” force, namely the force required to start the pedals moving when one wants

to use rudder, is 22 lbf at the airspeed of AA587. So one has to apply 22lbf to get the pedals to move at all, and then a further 10lbf to obtain maximum displacement. Is this too high? Is this just right? One is reminded of the story of Goldilocks and the Three Bears. No one in fact knows. A larger breakout force entails that it is harder for the pilot to activate rudder inadvertently, which could be an advantage in some situations.

So this stuff is subtle at best, and by no means as black-white as some parties to the debate might suggest. For example, one respected colleague described the A300-600 rudder pedal design as “ergonomically bizarre”. That seems to be an exaggeration. Airbus is known for subjecting its designs to extensive pilot evaluation before service entry, and this aircraft is no exception [U.S02b, pp499-500, Testimony of Captain Armand Jacob]. The science of flying qualities, which is a respected and necessary engineering discipline, remains quite subjective despite all attempts to make it as objective as possible (through rating scales and suchlike. Scales are based largely using the experiences of a given cultural set of pilots as norms, mostly test pilots for obvious reasons).

I don’t know that there is any adequate method of determining what forms of rudder-pedal feedback might be preferable, alternatively inadvisable, and why, amongst those currently in use on commercial transports. One can easily imagine that pilots who use rudder strictly for gust control at low airspeeds and for counteracting asymmetric thrust on engine failure, as foreseen by the current certification basis, would find little reason to fault the current A300-600 rudder control (as Captain Jacob indeed suggested, *op. cit.*). And there appear to be significant cultural differences amongst pilots as to the appropriate use of rudder control, these cultural differences extending wider than the variation in actuation parameters amongst the aircraft themselves.

3.4.3 Oscillatory Phenomena and Associated Characteristics

It has been suggested to the NTSB that the A300-600 appears to be susceptible to PIO in yaw-axis control. Indeed, it has been reported that American Airlines contends that the accident is due to PIO in the yaw axis, and not to its training or to its pilot’s unusual use of rudder control: their submission to the public hearing claims that “the” cause was “the onset of a design-induced, adverse APC event” [DF04].

Hess’s report concludes that the evidence is “consistent with” a PIO phenomenon [Hes03]. The French equivalent of the NTSB, the Bureau d’Enquêtes et d’Analyse pour la Sécurité de l’Aviation Civile (BEA), replied to such suggestions in January 2004 [Bur04].

Hess mentions two other PIO events, to a C-17 and to the Space Shuttle during flight test, by way of comparison. Those events, as well as most reported PIO events such as pitch oscillations to the Gripen and F22 test aircraft and the Gripen accident at the Stockholm Air Show [Var89] [Var93], involve coupling on a primary flight control, as does the majority of the inquiry in [Nat97]. The BEA points out inter alia that rudder is not a primary flight control, and it is certainly not the primary roll control. It suggests that a comparison against

PIO in primary control is inappropriate.

There is a good technical reason to accept the BEA's suggestion. All large transport aircraft have significant control hysteresis if one attempts to induce roll with yaw controls. Such hysteresis is a characteristic feature of controls which conduce to pilot-induced oscillatory phenomena in flight [3,4]. This hysteresis is all but unavoidable if one attempts to control roll with yaw.

Such hysteresis is characteristic of undamped or divergent control-oscillatory phenomena not only in flight control systems. Exceptionally long-period oscillations during strategic human-control attempts have been reported by the psychologist Dietrich Dörner [D03] in a series of fundamental experiments. His participants were to attempt to control the agricultural achievements and the sustainable living situation of a hypothetical African village in a resource-poor environment. The variables, such as water supply, were highly interconnected (using a lot of water leads to current plenty and future drought, for example) and participants were invited to make decisions to control the variables at a sequence of time points, at which they could see the consequences of their last decisions. Most participants exhibited oscillatory control behavior from which they were unable to exit.

This is a general characteristic of human control of feedback systems. Typical PIO behavior happens with a frequency of 2-5 Hz. Dörner's control points were discrete, but represented much longer time frames, of the order of months to years. And there is a middle situation between PIO and Dörner's strategic control. Oscillatory behavior can be caused by inexperience or incompetence, and is usually known as "overcontrol". A typical example occurs when learners or otherwise inexperienced pilots attempt to achieve a precise altitude, and oscillate about it, a phenomenon known as "chasing the altitude". It is common in pilots beginning their instrument-flight training. Such oscillatory phenomena have a typical frequency of a few seconds, and are put down to lack of judgement or skill – indeed, one has to maintain proficiency to keep out of such habits.

Another characteristic of divergent PIO is that the pilot's control inputs are high-gain: that is, hisher control inputs are large in comparison with that control input which would stably achieve the desired state. High-gain behavior can be exacerbated by a rate-limiting design, in which the rate of change of control lags behind (hysteresis) the control input. The rudder does not necessarily move in direct response to the position of the pedals, for example, but moves according to a rate set by the actuators. The position of the pedals commands a rate of displacement, but the rate at which the control surface moves is not directly proportional to the rate at which the pedals are moved if they are "stomped on". A rate-limited rudder design in large transports is, again, all but unavoidable if one does not wish to overstress the aircraft. But high-gain inputs under conditions of rate-limiting are two more characteristics of PIO phenomena.

The rudder of AA587 was rate-limited to 39deg per second. The rudder travel of the A300-600 is limited to 10-11deg at the 250 knot airspeed of AA587, whose rudder would have taken about 0.5 sec to travel from stop to stop [Dor02b]. Since the rudder reversals occurred 1-2 seconds apart, which is 2-4 times as long as the stop-to-stop period, one may conclude that the rudder

rate limit did not play a role in the oscillations, because the rudder had plenty of time to move to the stop before the next reversal was initiated.

One lesson to be drawn from these observations is that there is no hard line between aircraft-pilot coupling phenomena that can be put down primarily to characteristics of the aircraft and such phenomena that can be put down primarily to characteristics of the pilot (inexperience, lack of skill, inappropriate training, and so on). It's a social line: if Joe Blow the test pilot suffers it repeatably in an aircraft, and his colleagues do also, then the aircraft designers had better react. If Joe Blow the learner chases his altitude, then Joe Blow had better practice some more until he doesn't.

Another lesson that can be drawn is that some characteristics associated with PIO are all but unavoidable if one attempts to control roll with the yaw control (rudder) in large transport aircraft at higher airspeeds.

3.4.4 AA587's Oscillating Control Inputs

So what can one say about PIO in AA587? There is no doubt that oscillatory behavior happened. Why did it happen? I summarise the last section. Using rudder to control roll puts the pilot unavoidably in a high-hysteresis control situation. The rudder on large transport aircraft is unavoidably rate-limited (otherwise one could rip off the tail). High-hysteresis and rate-limited control situations can induce high-gain inputs in control, not only in pilots but also in notional "village chiefs" in strategic-psychological experiments. All three phenomena together are components of oscillatory controller-system behavior. However, the rate-limiting of rudder motion was likely not a factor in the AA587 oscillations.

Sensible use of the term suggests acknowledging that PIO occurred, but then asking the separate question whether it can be put down primarily to the PF, or whether, after the initial input, PF was inevitably caught through system characteristics in a control trap which he could not exit.

American Airlines's calling it "an adverse APC event" suggests they propose putting this PIO in the second class: that there is a design flaw and the fix lies with the designer, Airbus. I noted above that it is hard if not impossible for any manufacturer to design a large transport aircraft to avoid high hysteresis under attempts to control roll with rudder at higher airspeeds. So exactly how American Airlines could expect Airbus to have countered these characteristics is unclear.

(Such a contention, even if accepted, does not deal with the question why the PF even stomped on rudder in the first place. We have seen that it is not usual – and certainly not advisable – for pilots to be using rudder on large transport aircraft to control roll at high speed in non-upset situations.)

The BEA point out that "there is .. a lack of factual data to characterise the sensitivity of A300-600 rudder pedals, and, in general, data on PIO induced by rudder pedal inputs." Indeed so. They also note that, to their knowledge, "no studies have been undertaken on PIO on the yaw axis since the rudder is not a primary flight control" [Bur04]. Also correct. Design-induced PIO is

certainly not proven. Hess concluded only that the phenomena he investigated are “consistent with” PIO.

But Hess observes that the rudder control of the A300-600 is unusually sensitive in one respect (degrees per lb of rudder force) and has an unusually low ratio of maximum force (for full pedal deflection) to breakout force (force required to move the pedal at all). True. The BEA contends that that measure is by no means the accepted measure of sensitivity, or even an adequate one, and note the lack of data.

There are nevertheless reasons to question whether *any* argument based on the precise parameters of rudder control can be sustained. There is no record of exactly what force the PF applied to the rudder pedals during his four “rudder reversals” (stop to stop actuations). Despite the travel being notionally limited to 1.3 inches, the reconstruction shows two of those motions closer to 2 inches travel, and one close to 2.5 inches. See Figure 1. It may well be that the amount of force he put on the pedals would have been sufficient to take them to the stops at maximum rate on any of the transport aircraft with variable-stop rudder design that were considered for comparison. “[I]t may be that a pilot anxious to exit a situation will initially apply full control no matter how the pedal feels” [Dor02a]. If this was so, the particular characteristics of the A300-600 “sensitivity” vis-a-vis any other airplane with a variable-stop rudder pedal system would not have played a role; the same would have happened in any of these other airplanes. If this is so, the Counterfactual Test for the particular “sensitivity” characteristics of the A300-600 rudder pedals would not be fulfilled and thus the “sensitivity” would not be a causal factor. We do not know for certain. But there are estimates of the force that must have been exerted on the pedals to cause the recorded rudder pedal travel.

The rudder travel limiter is mounted near the rudder, at the rear of the aircraft. The rudder motion is constrained by the rudder travel limiter. There is a mechanical linkage between the pedals at the front of the airplane and the limiter at the rear, and this linkage, as any, is subject to elasticity. The pedals may be pushed further than full travel allowed by the limiter; the linkage stretches to accommodate. How much force would be required on the pedals to allow the linkage to stretch to allow the extra 1” or so rudder pedal travel shown on the FDR data points Figure 1? It was estimated by Airbus engineering during the hearings in October 2002 to be 130-140 lbs [U.S02a, page 100].

Maybe rudder pedal travel, only 1.3 inches, was a factor? I have already noted that the B727 has a similar travel at similar airspeed [Eva02]. So the A330-600 does not distinguish itself in this respect.

Whether the rest of the PF’s control inputs would have caused similar “rate saturation” (achieving maximum rate of deflection) and “amplitude saturation” (achieving deflection limit) in other aircraft besides this one is not considered in Hess’s report or elsewhere in the docket. If Airbus’s estimates of applied force based on the FDR data points are correct, the answer is yes: it would have happened in any other variable-stop-limited large transport aircraft. And it is generally accepted that such rate-saturated amplitude-saturated oscillatory rudder control could well have broken the tail off of any of them. If indeed

similar behavior would have been induced in many large transport aircraft by the PF's actual rudder input forces, then American Airlines' argument that the failure was "design-induced" would apply to all those aircraft. A generic design "flaw", then, in all large transport aircraft? Then, in general, one which was known, as exhibited by American's internal memos, mentioned below in Section 3.5, about the dangers of training recovery with rudder, and by Airbus, Boeing and FAA's 1997 advice to American Airlines on such use of rudder [Boe97].

How, then, could one establish that it is a design "flaw" ("flaw" is an evaluative term), rather than a design characteristic about which pilots and their employers are expected to know and thereby to avoid? To put it in terms well-known to computer people: Is it a bug or a feature? Whichever, one expects responsible operators to avoid triggering known bugs.

The "flaw" argument cannot be substantiated on this basis.

3.5 Training in Use of Rudder

The NTSB finds that American Airlines trained use of rudder in recovery methods for wake turbulence encounters [U.S04b, Conclusion 8]. Indeed, a former NTSB member had taken their AAMP training course in which he recalls they did so, as noted above in Section 2.2. A 1997 letter from Captain Tribout, A300 Technical Pilot, to William Wainwright, Airbus chief test pilot, expressed concern about American Airlines teaching use of rudder to control roll in wake turbulence encounters, amongst other situations. He suggests that it is "potentially hazardous" [DF04]. A letter from Paul Railsback, American Airlines managing director of flight operations-technical, to American's vice president of flight operations, at the time of the earlier overload incident in 1997, expressed "grave concerns" about pilots in AAMP being taught to "use rudder as the primary means of roll control in unusual attitude recoveries. This is not only wrong, it is exceptionally dangerous.. American Airlines is at grave risk of a catastrophic upset." [DF04].

And there is Molin's comment to Lavelle, reported by Lavelle, suggesting that he believed use of rudder was trained by American Airlines [6,16].

It is worth quoting at length from a joint Boeing, FAA and Airbus memo to American Airlines on 20 August 1997, concerning their AAMP training [Boe97, Dor02a]:

The excessive emphasis [in AAMP] on the superior effectiveness of the rudder for roll control is a concern. Rudder reversals such as those that might be involved in dynamic maneuvers created by using too much rudder in a recovery attempt can lead to structural loads that exceed the design strength of the fin and other airframe components. The hazard of inappropriate rudder use during wind-shear encounters, wake turbulence, should also be included in the discussion.

Boeing, Airbus and the FAA, then, explicitly notified American Airlines of their concern with using rudder for roll control, that using rudder during wake tur-

bulence encounters could constitute a hazard, and that rudder reversals could overstress the fin. That is exactly what happened with AA587, four years later.

American Airlines's reply to the August 1997 letter said in part:

Let me say this one more time, we do not advocate the introduction of large sideslip angles when flying at high angles of attack. You seem to be predisposed to the belief that we are using rudder first or rudder only. The workbook is not a standalone document and nothing should be inferred without listening carefully to the presentation. In four different sections of the AAMP, emphasis is focused on the fact that when the airplane is not responding to aileron and spoiler control, you should use smooth application of coordinated rudder to obtain the desired roll response. The hazard associated with large of abrupt application of rudder at high angle of attack is clearly exemplified by [NTSB video re-creations of two loss-of-control accidents]

The reply begins in a manner which some might regard as intemperate. It does not address the specific concern expressed by Boeing, Airbus and the FAA about using rudder during wake turbulence encounters, or the specific concern that rudder reversals can overstress the fin. (Besides that, any teacher can attest that if you write something down that is not what you say, what will be remembered is what you wrote, not what you said.)

Despite that contretemps, I think no one would suggest that American Airlines would have taught Molin to stomp on rudder during a wake turbulence encounter.

The question arises how American Airlines could have continued with that form of training even after the company had received explicit notification from Boeing, Airbus and the FAA that use of rudder in certain ways could overstress the airplane and their concern about advocating use of rudder during wake turbulence encounters, and after similar concerns had been expressed by American Airlines's own personnel at a high level. This is a matter for sociotechnical and organisational theorists to explain.

3.6 Line-Pilot Understanding of Rudder Characteristics

It does seem that the accident investigation uncovered significant differences of understanding on the use of rudder and its certification basis amongst certain pilot groups [U.S04b, Conclusions 10,16]. One is reminded of the wide variety of advice, some of it inconsistent with other advice, uncovered by the German BFU in the wake of the Überlingen midair collision, on the use of TCAS/ACAS. But TCAS is new, and rudders and their characteristics have been around since the Wright brothers.

3.7 Fin and Rudder Certification Criteria

It seems as if the NTSB has decided that the certification basis for rudder use and vertical stabiliser strength is inadequate. They recommend that modified criteria be developed, and retroactively applied [U.S04b, Recommendations 1 and 2 to the FAA].

As discussed in Section 3.4.1, the certification standards require static tests of rudder loading: you put in full hard rudder and hold it there until the aircraft has stabilised in yaw, then you release it equally abruptly. The forces can be calculated and that sets Limit Load. However, the certification standards do not require the structure to be investigated under application of full opposite rudder during sideslip [Dor02b], nor under rudder reversal, let alone repeated reversals such as occurred to AA587. It has turned out that this was less well-known in the industry than it might have been. Indeed, I knew neither the certification basis nor the details of rudder loading until Michael Dornheim analysed rudder certification in detail in Aviation Week as part of their initial coverage of the accident [Dor02b]. I remember being surprised. Clive Leyman, former chief aerodynamicist on Concorde, pointed out to me that it is very hard to measure and analyse forces on the rudder under such movements [Ley02]. It is obviously possible nowadays, with the highly improved codes that have been made available over the decades since the airplane was certified, and indeed Airbus used such methods accurately to calculate the overload in this accident. Results of their calculations are in the docket [U.S04a].

The current certification basis for the rudder control appears to be consistent with the intended use of rudder. One could imagine that, because one could not well calculate dynamic overloads due to oscillation, the industry (regulators and airplane builders) relied on instilling as Best Piloting Practice that you Just Don't Do That, because you could rip the tail off an airplane. And, indeed, given the intended use of rudder and its certification basis, there is no need ever To Do That.

It is rare that the NSTB determines that the certification basis of such a fundamental property be modified, and even rarer that it recommends compulsory modification of existing airframes to a new criterion.

4 Some Views of Others

4.1 Design Flaw?

A colleague who had read newspaper reports of the public hearing suggested to me that it was inappropriate to cite the PF's behavior as probable cause, because the case was a "design flaw". There are two ways of interpreting such a suggestion:

- there is a specific design flaw with this model aircraft;
- there is a generic design flaw with all large transport aircraft.

First, calling a feature a “flaw” is an evaluative term. I look at what could substantiate such an evaluation.

Consider the first claim. I have pointed out in Section 3.4.3 that characteristics leading to oscillatory behavior are to be expected if attempts are made to control roll with rudder.

Consider the second claim. Fins are not designed to be strong enough to prevent a pilot ripping the tail off under overload during oscillations. Data from this accident would suggest that, to be proof against such phenomena, vertical stabilisers would need to be twice as strong as they currently are. Should they be so built? It is perhaps an unworkable suggestion to try to protect the airframe from the consequences of arbitrary pilot behavior. Dornheim quotes a “flight controls expert at a major airframe company” saying “I think you can break an aircraft in any axis if you work on the controls. Operating aircraft relies on basic airmanship” [Dor02b].

Calling something a “flaw” suggests that it should be fixed. But there is no way at present to fix the phenomena, other than by developing new systems to control the trigger behavior, much the way, say, that modern cars are equipped with anti-skid braking systems. This, I take it, is the gist of the NTSB recommendation.

Let us ask what we might mean by saying something is flawed. If we mean that there is a non-zero risk of catastrophic failure, then one can easily admit that there is indeed a non-zero risk of catastrophic failure, but also point out that this is true of many aircraft structures and phenomena. Indeed, the international standard on functional safety of electrical/electronic/programmable electronic (E/E/PE) systems is explicitly based upon the premiss that there is no such thing as zero risk. So in this meaning the word “flawed” would be a tautology.

So maybe one means by “flawed” that there is an *unacceptable* risk. There are a lot of criteria that could come into play. One is hindsight: of all the aircraft out there, there are apparently only three that are known to have stressed the fin at or beyond Ultimate Load through use of rudder. Two were on A300-600 machines belonging to American Airlines. The third incident occurred to an A310 of Interflug in 1991. Besides that, an Air France A310 exceeded Limit Load but not Ultimate Load in 1999. Three of those four incidents involved recovery from loss of control. There are many ways to break large aircraft if you lose control. But none of those aircraft broke. AA587 broke, but not due to loss of control.

According to Airbus, A300-600 machines have accumulated 16 million flight hours. The A300 fin overstress events are thus very rare, of the order of 10^{-7} per flight hour. A tendency has been noted for a broad measure of acceptable risk across society to be an incident occurrence less frequent than 10^{-6} per hour [Lew90, Page 104], although there is considerable discussion as to what constitutes “acceptable risk” [Lew90, FLS+81]. The “standard” risk of some all-causes catastrophic system failure in a commercial aircraft in flight has been conventionally set at 10^{-7} per flight hour for some thirty years [LT82]. This figure is justified through a rationale based on a broad calculation of 100 independent subsystems, each with a 10^{-9} per hour probability of independent

failure [LT82]. A structural failure due to PIO is not an independent subsystem failure, but a many-causes dependent failure, so the 10^{-7} criterion is appropriate, and the statistics so far fit that. Dornheim widened the sample: “[The] record appears to show that Flight 587 is the first time in civilian jet transport history that a fin has come completely off due to aerodynamic loads. It is on the order of the one-in-a-billion-flight-hours safety rate targeted by the [US Federal Aviation Regulations]” [Dor02b].

So it looks as if an argument that “flawed = unacceptable risk” isn’t going easily to lead to the conclusion that the design was flawed either. It is complicated by the fact that both Above-Ultimate-Load incidents happened to American Airlines crews, so even if the two incidents were broadly to constitute an unacceptable risk, we would be unable to distinguish between that part of the risk due to design issues, and that part due to organisational culture. And this latter causal factor is demonstrated to have been there [U.S04b, Conclusions 8,9].

5 A Superficial WBA

The structure of a causal explanation of the accident is a lot more simple than the above discussion might lead one to think. Many investigations nowadays are hampered by evidential arguments presented by one side or another to emphasise their favored factors and obscure their disfavored ones. But one should not confuse the causal assertions being made or rejected with the difficulty in assessing evidence. Here is how the beginning of a causal analysis, a Why-Because Analysis (WBA), might look. I express it textually and informally. The “BECAUSE” clauses are verified by applying the Counterfactual Test.

The graph of the BECAUSE relation below is easy to draw, and is shown in Figure 5.

5.1 Getting to the Root-Causal Factors

First, I take a shortcut: the accident, as defined in 14 CFR, the Federal Aviation Regulations, stems causally from loss of the fin, and this part does not interest us that much. Interesting is what led to the loss of the fin.

There are eleven general factors that stand in causal relationships according to the Counterfactual Test:

- [1] The fin broke off
- [2] The fin was overloaded
- [3] PIO behavior
- [4] Wake turbulence
- [5] Pilot control behavior

- [6] Design characteristics
- [7] General pilot culture
- [8] Individual PF characteristics
- [9] AAMP training
- [10] Airbus design choices
- [11] Certification basis for fin strength

They are arranged as follows. I say X *BECAUSE* Y, Z, \dots if Y, Z are necessary causal factors of X (they pass the Counterfactual Test), and Y, Z form a sufficient set of causal factors for the occurrence of X .

- [1] BECAUSE [2]
- [2] BECAUSE [3], [6]
- [3] BECAUSE [4], [5], [6]
- [5] BECAUSE [7], [8], [9]
- [6] BECAUSE [10], [11]

I noted above that there are subtly different ways in which the PF behavior and the AC characteristics can combine to cause the PIO. Maybe the PF stomped on the rudder at first, and then became victim of an oscillatory regime that could not be handled by anyone, and the rudder broke off before he could figure out to just take his feet off the pedals. Or maybe he was actively overcontrolling the airplane. Therefore how the design characteristics (factor [6]) influenced the PIO (factor [3]) is undetermined. However, *that* the design characteristics causally influenced the PIO is a tautology, just as it is a tautology that the pilot behavior (factor [5]) influenced it also.

If we enumerate the causal factors that themselves have no enunciated causal factors, we have [4], [7], [8], [9], [10], [11]. These are the “root-causal” factors of this simple WB-graph. And they correspond well to the general Findings of the NTSB. (An exception is the occurrence of wake turbulence, which is considered a normal event in air travel and is not noted by the NTSB. A WBA identifies it as an environmental causal factor.)

We have seen that considerable investigative effort has been put into investigating the exact causal links between:

- [6] and [3]
- [7] and [5]
- [8] and [5]
- [9] and [5]

- [10] and [6]
- [11] and [6]

This subtlety should not obscure the existence of causal linkage (the “necessary causal factor” relation). The investigative subtlety has mostly gone into determining the relative weight of the causal factors, something which WBA does not do, and which I believe is dependent on many extraneous and not necessarily objective judgements.

5.2 Countermeasures

Countermeasures are actions that can be taken to reduce or negate the influence of certain causal relationships. The basic point is that if any of the BECAUSE links can be broken, the accident cannot happen (because some necessary causal factor is missing). So recommendations attempt to break BECAUSE links for the future, amongst other things. What can one reasonably do?

[4] : Attempting to influence wake turbulence is a non-starter: the inferred turbulence was not more extreme than what many airplanes encounter on a daily basis.

[7] : General pilot culture. Here the NTSB recommended various measures to make pilots more aware of appropriate use of rudder, of its certification basis, and that abrupt control movements below V_a are not necessarily benign.

[8] : This is a non-starter, because the PF is dead. However, one may attempt to influence others that may be like him through awareness and training. Here, countermeasures merge with those related to [7].

[9] : Influencing not just AAMP but all other airline training programs is a must. The NTSB had things to say.

[10] : Although, as we have seen, it is questionable how much Airbus’s design choices could or should be modified to “break the causal link”, the NTSB recommended that this should be revisited.

[11] The NTSB recommended the certification basis be revisited.

So we can see that the NTSB recommendation list covers all the root-causal factors, with the exception of wake turbulence, which is regarded, properly in my opinion, as an environmental property.

6 Some Moral Comments

Along with some of my colleagues, I have been disturbed by what appears to us to be overwhelmingly partisan argumentation during the course of this

investigation. It seems to many as if some major players figured out what was in their future best interests first, and then sought out and promoted exclusively arguments that furthered those interests. I am more inclined to believe that dispassionate inquiry - a willingness to analyse, and to judge according to overt and explicit criteria - is most likely to lead to the goal of accident investigation, which is to improve safety for the future.

Featured in “contention”, as noted by such journals as Aviation Week, have been primarily American Airlines and Airbus (see, for example, the first paragraph of [DF04] and the second paragraph of [Fio04])

Whatever may be their true motivation, Airbus’s contributions are consistent with the simple causal analysis in Section 5. That unfortunately cannot be said of American Airlines’ contributions, which seem to have been concentrated on establishing a hypothesised “design flaw”, and ignoring their own, fairly well established, causal contribution via their training regime, which parts of their own company, as well as Airbus, Boeing and the FAA, queried before this accident. American’s training is beyond reasonable doubt a causal factor in the oscillatory control event that led to the structural failure. A “design flaw” alone cannot have root-caused the accident, according to the WBG. American’s claim as to root causes is simply factually mistaken.

One might query why American Airlines is engaging in public in such demonstrably faulty causal reasoning. I shall leave that to others to explain, as I shall leave to others the explanation of how American could have continued to train use of rudder in upset recovery against both internal company advice and manufacturer and regulator advice.

The best hope for the future of accident investigation seems to me to be to establish an objective standard of causal reasoning to which any public contributions to the investigation can be seen to adhere. I advocate, as usual, WBA.

Acknowledgements

I thank Jim Armstrong, Hal Lewis, Ira Rimson and Harold Thimbleby for helpful comments on an early draft, and Heiko Holtkamp for two of the graphics.

References

- [Boe97] Boeing, FAA, Airbus. Letter to american airlines’ chief pilot, 20 August 1997. available from www.nts.gov/events/2001/AA587/exhibits/239998.pdf.
- [Bur04] Bureau d’Enquêtes et d’Analyse pour la Sécurité de l’Aviation Civile (BEA). Letter to Mr. Robert Benzou, Investigator in Charge, 13 January 2004. available at www.nts.gov/events/2001/AA587/exhibits/290584.pdf.

- [Dö3] Dietrich Dörner. *Die Logik des Misslingens: Strategisches Denken in komplexen Situationen*. Rowohlt Verlag, Reinbeck bei Hamburg, 2003.
- [DF04] Michael A. Dornheim and Frances Fiorino. NTSB wraps up AA587. *Aviation Week and Space Technology*, page 45, 25 October 2004. also available at www.aviationnow.com/avnow/news/channel_awst_story.jsp?id=news/10254air.xml.
- [Dor01] Michael A. Dornheim. Composites experts study flight 587 stabilizer. *Aviation Week and Space Technology*, pages 36–7, 3 December 2001.
- [Dor02a] Michael A. Dornheim. Crash probe triggers rudder design scrutiny. *Aviation Week and Space Technology*, pages 44–6, 1 April 2002.
- [Dor02b] Michael A. Dornheim. Did rudder motions snap off A300 fin? *Aviation Week and Space Technology*, pages 24–6, 21 January 2002.
- [Dor02c] Michael A. Dornheim. Flight 587 probe shows tails have been overloaded before. *Aviation Week and Space Technology*, pages 44–51, 25 November 2002.
- [Eva02] David P. Evans. Upset recovery training and the potential for negative learning. *Aviation Safety Week*, 11 November 2002. available from www.findarticles.com/p/articles/mi_m0UBT/is_43_16/ai_94139244.
- [FD02] Frances Fiorino and Michael A. Dornheim. AA587 Hearing: Rudder 'key' to solving mystery. *Aviation Week and Space Technology*, pages 47–51, 4 November 2002.
- [Fio01] Frances Fiorino. Rudder deflections eyed in AA587 probe. *Aviation Week and Space Technology*, pages 32–3, 26 November 2001.
- [Fio02a] Frances Fiorino. NTSB seeking links between AA587 and earlier incident. *Aviation Week and Space Technology*, pages 52–3, 4 March 2002.
- [Fio02b] Frances Fiorino. Pilots ask: How much rudder is too much? *Aviation Week and Space Technology*, page 44, 18 February 2002.
- [Fio02c] Frances Fiorino. Upset events, wake vortices at center of flight 587 probe. *Aviation Week and Space Technology*, page 39, 22 April 2002.
- [Fio04] Frances Fiorino. Cause and safety effects. *Aviation Week and Space Technology*, pages 50–1, 1 November 2004.
- [FLS⁺81] Baruch Fischhoff, Sarah Lichtenstein, Paul Slovic, Stephen L. Derby, and Ralph L. Keeney. *Acceptable Risk*. Cambridge University Press, Cambridge and New York, 1981.

- [Hes03] Ronald A. Hess. An inquiry into whether a pilot-induced oscillation was a causal factor in the crash of American Airlines Flight 587, December 23, 2003. Published in the NTSB public docket as Human Performance Group Chairman's Factual Report, Addendum 2, December 2003. available at www.nts.gov/events/2001/AA587/exhibits/288388.pdf.
- [Lew90] H. W. Lewis. *Technological Risk*. W. W. Norton & Company, New York and London, 1990.
- [Ley02] Clive Leyman. Personal communication, 2002.
- [LT82] E. Lloyd and W. Tye. *Systematic Safety*. UK CAA Publications, London, GB, 1982.
- [Nat97] National Research Council. *Aviation Safety and Pilot Control: Understanding and Preventing Unfavorable Pilot-Vehicle Interaction*. National Academy Press, 1997.
- [NTS04] AA903 factors affecting AA587. Board Meeting of October 26, 2004, 26 October 2004. Available at www.nts.gov/events/2001/AA587/presentations/AA903_factors.pdf.
- [O'C04] John O'Calloghan. Aircraft performance presentation. Board Meeting of October 26, 2004, 26 October 2004. Available at www.nts.gov/events/2001/AA587/presentations/02_airplane_perf.pdf.
- [U.S] U.S. Federal Government. Federal Aviation Regulations Part 25.335, Design airspeeds. available at www.risingup.com/fars/info/part25-335-FAR.shtml.
- [U.S02a] U.S. National Transportation Safety Board. Transcript of public hearing, american airlines,flight 587, belle harbor, new york, october 29, 2002, 29 October 2002. Available from www.nts.gov/events/2001/AA587/trans_021029.pdf.
- [U.S02b] U.S. National Transportation Safety Board. Transcript of public hearing, american airlines,flight 587, belle harbor, new york, october 30, 2002, 30 October 2002. Available from www.nts.gov/events/2001/AA587/trans_021030.pdf.
- [U.S04a] U.S. National Transportation Safety Board. Exhibits from the Public Docket of Aviation Accident, In-Flight Separation of Vertical Stabilizer, American Airlines Flight 587, Airbus Industrie A300-605R, N14053, Belle Harbor, New York, November 12, 2001, NTSB/AAR-04/04, October 2004. available at www.nts.gov/events/2001/AA587/exhibits/default.htm.

- [U.S04b] U.S. National Transportation Safety Board. Synopsis of Report of Aviation Accident, In-Flight Separation of Vertical Stabilizer, American Airlines Flight 587, Airbus Industrie A300-605R, N14053, Belle Harbor, New York, November 12, 2001, NTSB/AAR-04/04, October 2004. available at www.nts.gov/publicctn/2004/AAR0404.htm.
- [Var89] Various Authors. Gripen pitch oscillations. In Risks 8.32 and 8.49, 1989. catless.ncl.ac.uk/Risks/8.32.html and catless.ncl.ac.uk/Risks/8.49.html.
- [Var93] Various Authors. Gripen airshow pio. In Risks 14.81, 14.82, 14.85, 14.89, 1993. catless.ncl.ac.uk/Risks/14.81.html,
catless.ncl.ac.uk/Risks/14.82.html,
catless.ncl.ac.uk/Risks/14.85.html,
catless.ncl.ac.uk/Risks/14.89.html.

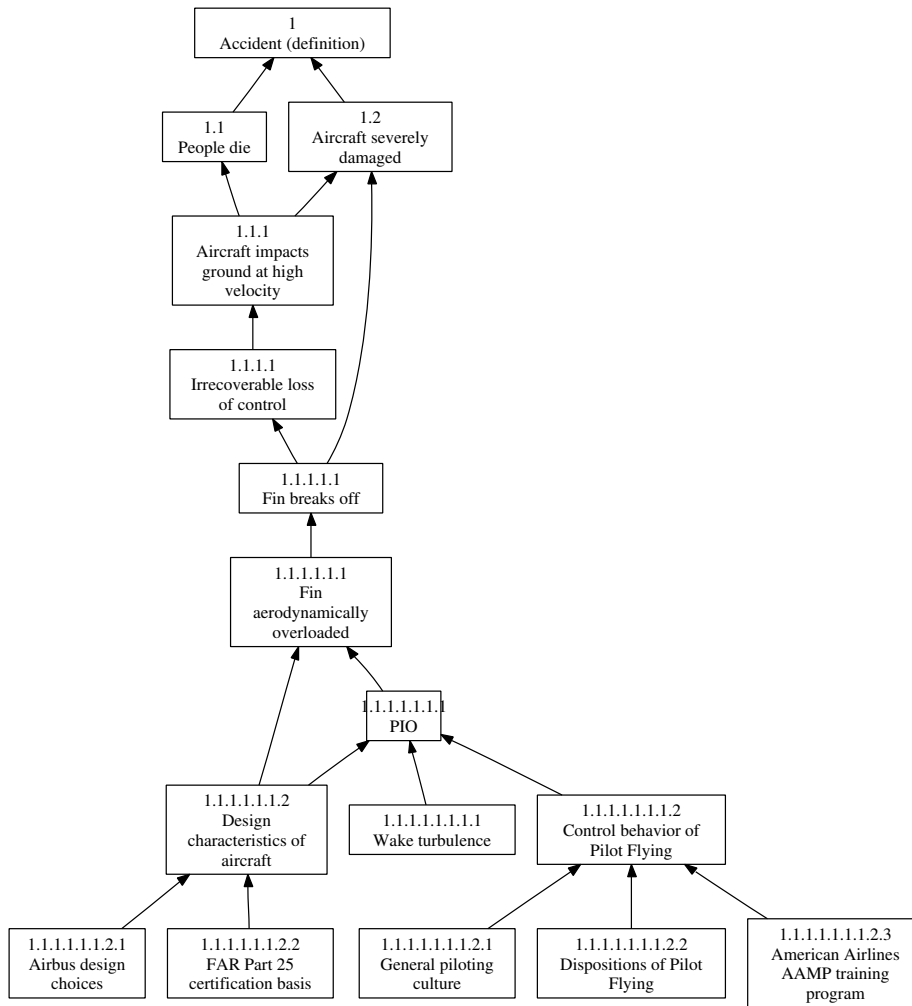


Figure 4: The High-Level Why-Because Graph of AA587