International Federation of Airworthiness Conference
Auckland, New Zealand : 20-23 October 1991

Conference Theme
Working together to stop aircraft accidents.

*Paper 9 (Issue 3)**

Breaking through the 10^6 Barrier

R W Howard, CBE, BE, CEng, FRAeS
Chairman, GEC Avionics Limited

Summary

The fatal accident rate in the world commercial jet fleet has been on a near plateau for 15 years since the absorption of the major advances made with the last generation of technology and operational systems.

Notwithstanding continuing technical advances, the increasing size of the world fleet demands more significant improvements if the actual number of accidents is to be reduced.

The paper promotes the view that a further overall increase in the safety of air transport from the high levels currently achieved will not come about by normal evolution. It suggests that what will be required in future is planned, quantifiable improvements compatible with statistical safety assessment methods. Such assessment should encompass the totality of an integrated Air Transport System, as distinct from the current practice of assessing the Aircraft Airworthiness, Air Traffic Management, and Air Navigation systems in isolation.

A relevant history of system statistical safety assessment is reviewed, and certain aspects of cockpit automation and advanced systems which may contribute to future advancement of overall safety are considered.

* Incorporating modifications included in the Conference Presentation.

Breaking Through the 10^6 Barrier

Ron Howard
Chairman, GEC Avionics Limited

1. Introduction

Fatal accident statistics in scheduled airline operations are assembled by a number of aircraft manufacturers and also collected and published by various civil aviation authorities and by ICAO. These derive today from the operation in the West of nearly 10,000 scheduled jet aircraft flying cumulatively over 20 million hours per annum.

Fatal accident statistics are presented in a variety of ratios; numbers per million flying hours, per departure, per 100 million aircraft-km, per landing, per 100,000 landings and others. Each has a particular bias and use. In this paper I shall refer to fatal accidents per million hours, or the corresponding per hour figure, because I feel that this is the ratio which most nearly relates to individual and group concepts of the risk of living.

In the early years of the jet age, the late 1950s, world fatal accident rates for larger aircraft were in a bracket from 2 to 6 per million hours flying, with the large US operators among those achieving the lowest rates and smaller non-scheduled operators tending to be at the other end of the range.

In the 1960s and early 1970s the world rate improved to around 1 to 1.5 per million hours, with little difference between he major countries operating schedules airlines.

Over the past 15 years the world rate has remained on a plateau in the bracket 1 to 1.5 per million hours (1 to 1.5 in 10^6 per hour), but it is encouraging that the operations of some countries have continued to improve and have settled out at levels lower than 1 per million hours.

However the world average for scheduled operations remains fairly steady at a figure somewhat higher than 1 in 10^6 per hour and notwithstanding certain pockets of higher performance this seems to have become a "barrier" waiting for some breakthrough to open the way to significant advances.

Looked at another way, airliners in scheduled operations seem at present to have reached a limit to their safety, and as a consequence the actual number of accidents will continue to increase as more and more aircraft are produced and enter service.

It is the intention of this paper to examine the fatal accident levels as represented by the statistics, speculate on possible future targets for improvement, and explore possible means for achieving them.

2. Aircraft Fatal Accident Rates. A Perspective

All fatal accidents are unacceptable, and it is beyond human power to eliminate them completely. The aim is to keep the rate of occurrence as low as humanly possible.

The question which has been asked many times is "what is a low rate of occurrence"? The starting point has usually been average human mortality rate. Taking as an example, people living in the United Kingdom, the figures from all causes vary from 1 in 10^7 per hour for 20 year olds to 2 in 10^6 per hour for 60 year olds.

So, a first subjective view on the basis of these figures might be that a fatal accident rate in scheduled airline flying of 1 in 10^6 per hour is not at all unreasonable. Thus while any improvement must be acceptable, it might not necessarily be regarded as essential, especially if increases in fares were involved.

However let us look at some other statistics which contribute to the average mortality rate. Firstly, 1 in 10^6 per hour is two to three times worse than the fatal accident risk to car drivers in the United Kingdom since the requirement for seat belts and random breath testing, and sixty times higher than the fatality risk in a passenger train.[1] This reflects not an exceptionally low frequency of train accidents, but rather a greater survivability.

These other comparisons, contrary to the first view, would not seem to justify the current level of risk of air travel.

It might be said in mitigation that the higher speed of air travel over other forms does minimise the "total" risk for most individuals, who do not spend a high proportion of their lives travelling by air. On the other hand, most individuals would probably argue that when they make any significant journey, one form of travel should not put them at a greater risk than another.

There is also another factor to consider. Relative risks are not just a matter for individual acceptance or otherwise. There is also a public dimension.

The public perception is not the same as that of the individual. A multiple fatal accident on a motorway or an accident involving a carriage full of people in a rush hour train, is manifestly more distressing to the public than if the same numbers of people were killed separately in car

accidents. The loss of a wide bodied jet with all its passengers is very much more so.

The number of fatalities involved in each accident is therefore relevant, and it is not unreasonable for the public to expect that accident risk should be lower for transport vehicles which have a high passenger carrying capacity.

We do not have an agreed measure of the level of public distress at which the demand will be for "something to be done", but there are indicators. For example, the total world scheduled airborne passenger fatalities in 1990 were just over 400, whereas road deaths for UK alone were over 4,000. The accidents involved in the former created world headlines, while the latter, although widely reported, had the status of odd columns in inside pages of UK newspapers.

Taking all of these factors into account, if the relative levels of individual and public concern were to be estimated, my own judgement would be that the combined risk and distress factor ratios between car, train and wide-bodied jet aircraft accidents would imply that some improvement in air travel safety is required.

I consider that a fatal accident safety level for air travel at least as low as 0.1 in a million hours, which is one tenth of the present rate, should be pursued.

If this could be obtained across the existing world fleet, it would reduce fatal accidents from one per month to one or two per year, and major disasters would not occur more often that once in 5 to 10 years.

I am aware that to define a specific accident rate target is not a universally acceptable principle. At the moment it is applied only to certain aircraft equipment and systems. There is now a need for wider application combined with the stringent setting of new targets for all elements which contribute to risk, and all their interfaces, so as to arrive at an overall improvement.

In short, if we want to do something about safety, it is far better to work to planned targets than to go on like we are, trusting that the necessary improvements will evolve as a matter of course.

3. The Potential for Significant Improvement

Before examining the potential for future improvements, it would be worthwhile to look at past history and see how we got to where we are, and what we might learn from it.

In the 1950s the world fatal accident rates were not much better than 1 in 10^5 per hour. However there were enormously relevant technological developments in progress, the most important being turbojet propulsion, followed by greatly improved communication systems and air traffic controls and of course the widespread use of silicon transistor and micro-circuit technology. In simple terms these really amounted to large increases in reliability. In the 1950s and 1960s there were also further very significant advances in design for fail-safety and failure survival, and their quantification. All of these together had a large impact on the performance and safety of air transport operations, which became very apparent in the first generation of jet transports and more so in the second generation.

Over the period 1964 to 1968, despite a threefold increase in the number of aircraft in world service, the accident rate per million flying hours was actually halved. The improvements in safety then slowly eased to the present levels, despite continuing technical advances and the large increases in equipment reliability and maintenance efficiency.

What this tells us is that after a period of fundamental improvements, the law of diminishing returns takes over.

Unfortunately we do not see any advances coming over the horizon as significant as the turbojet engine or silicon electronics, certainly none which are likely to improve the 1 in 10^6 level to the extent that the 1950s and 1960s technology advances improved the 1 in 10^5 levels.

So in the near term the prospects for big improvements do not look hopeful.

In my opinion, improvements from the 1 in 10^6 per hour level will be obtained principally by minute attention to the detail of the last area mentioned of advances made in the 1960s, that of fail-safe and fail-operative concepts in aircraft design, and this must also be extended to the air traffic control environment, the navigation environment and the control and management activity of the flight crew.

I also believe that in the pursuit of better fail-safe and fail-operative concepts, formal quantitative safety assessment methods will play a larger part in future than in the past,.

4. Statistical Safety Assessment Methods

An early example of the use of formal statistical safety assessment methods was in the certification of automatic landing by the British CAA Air Registration Board in the early 1960s. This has been widely reported.[2] The method embraces the total systems safety in automatic landing from the viewpoint of both equipment reliability and touchdown performance, laterally and vertically.

Similar methods are widely used for assessing equipment and overall systems reliabilities, and as an aid in synthesising the desired degree of fault-tolerance in design. They are used to a lesser extent for performance assessment, this being applicable mainly to safety critical guidance and controls aspects of aircraft operations.

The method requires the definition of an overall "risk budget" for the "system", in which the survival requirements for major elements are specified. This will involve airborne, ground and flight crew aspects.

The purpose of such assessment was not only to achieve a satisfactory design but to ensure that the certification authority (the CAA/ARB) could be given the necessary evidence that the introduction of any new equipment in the landing and take-off operations would not increase the then-current fatal accident probability. The current figure in the 1960s was 1 in 10^6 per landing, so 1 in 10^7 per landing (in relation to an exposure time varying between 30 seconds and 3 minutes) was chosen as a target.

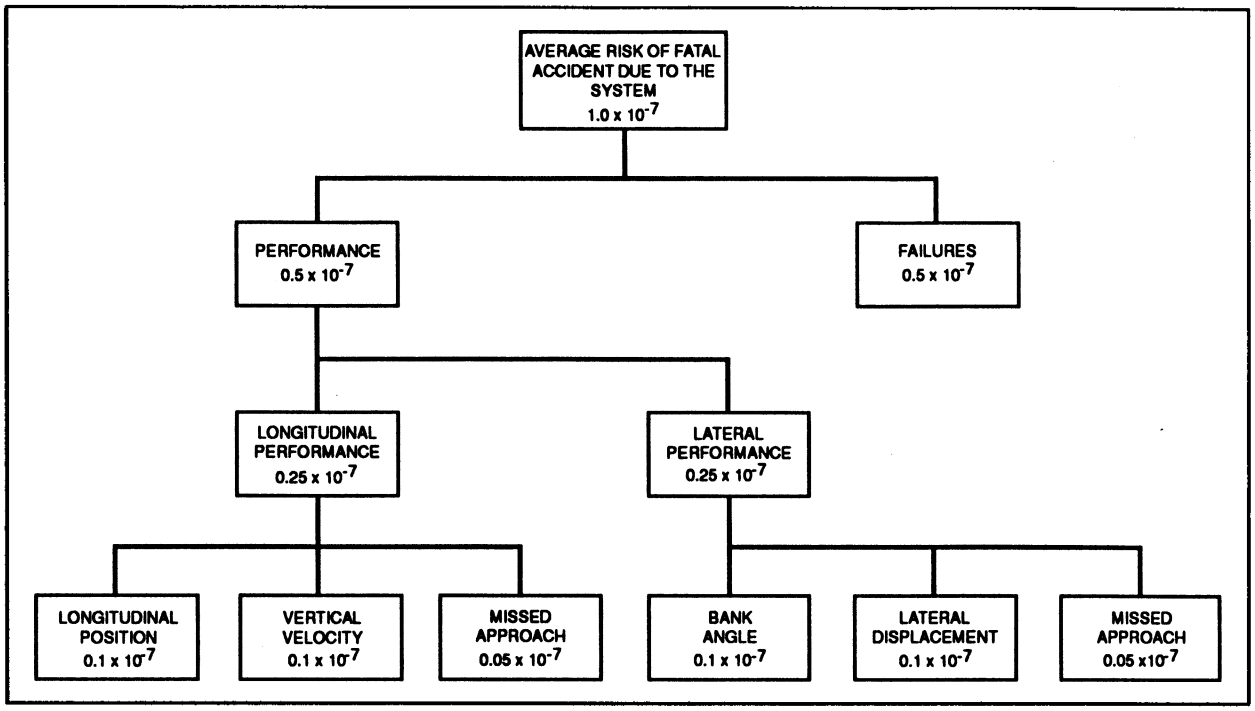

Figure 1. Sample breakdown of average risk

Figure 1 shows an example "risk budget" giving a sample breakdown of average risk for a complete automatic landing system, starting with a 1 x 10^{-7} per landing requirement. This shows also the subdivision of overall average risk into "performance" and "failures" allocations, and in this case a 50/50 split was postulated, 0.5 x 10^{-7} for performance and 0.5 x 10^{-7} for failures.

These two risk elements are applied in different ways.

Looking first at the "failures" figure, this is the maximum allowed as a cumulative failure probability of equipment in the period of exposure, taking into account predicted reliability from cumulative component data (Mean Time Between Failure data) and the effect of fault tolerant architecture employed in the basic design.

Armed with such an overall figure, and the systems architecture, a designer can expand the detail of the risk budget and hence postulate the contribution to survival required by each system element, thus creating the safety specification for each element of the total system.

If any single element cannot alone be proved to meet its risk requirement, then proof of some back-up capability is required, or there must be provided some form of multiple equipment redundancy.

In practical design terms care must be taken at this stage to ensure that if multiple redundancy is used, it can be applied in a manner which avoids any probability of common mode failures. This may demand the use of dissimilar redundancy in some cases, especially in digital systems employing complex suites of software in which all failures may not be quantifiable.

Thus the "risk budget" on equipment reliability becomes an element in an iterative design process which results in a system with a predicted safety which is quantitatively assessable by a certification authority.

The achievement of the performance figure is more difficult to prove.

This is not possible in actual test flying. To predict a probability of 0.5 x 10^{-7} per hour to a reasonable confidence level would require over 5 million hours of flying. The answer is to use a combination of flight testing and simulation. A simulator model can be validated by a minimal amount of flight testing, following which the simulator can be run at a high multiple of real time to assess thousands of flight cases, which, using a Monte Carlo approach, cover a wide range of parameter combinations and variations. Sufficient data can be obtained in this way to predict with adequate confidence, the tails of the distributions on accident probability. Figure 2 is an example carried out for the certification of Concorde automatic touchdown distribution.

Clearly the "risk budgets are much more complex than the block diagram of Figure 1.

Figure 3 is a 1960's example giving a summary of risks during a Category II approach. This includes not only the previously mentioned "performance" and "failures" (faults) elements, but also the risk of overshoots and their associated operational hazards. I include this figure to illustrate not only the level of complexity but also the strength of background, as these assessment methods have been in use for over 30 years.

The method was first promulgated in the United Kingdom in November 1961, when the CAA/ARB issued the British Civil Airworthiness Requirements (BCAR) paper 367 "Airworthiness Requirements for Autoflare and Automatic Landing".[2] From 1961 to 1970 considerable experience was gained in the UK from the certification programmes for automatic landing aircraft, which advanced significantly the state-of-the-art in designing and analysing high integrity systems.

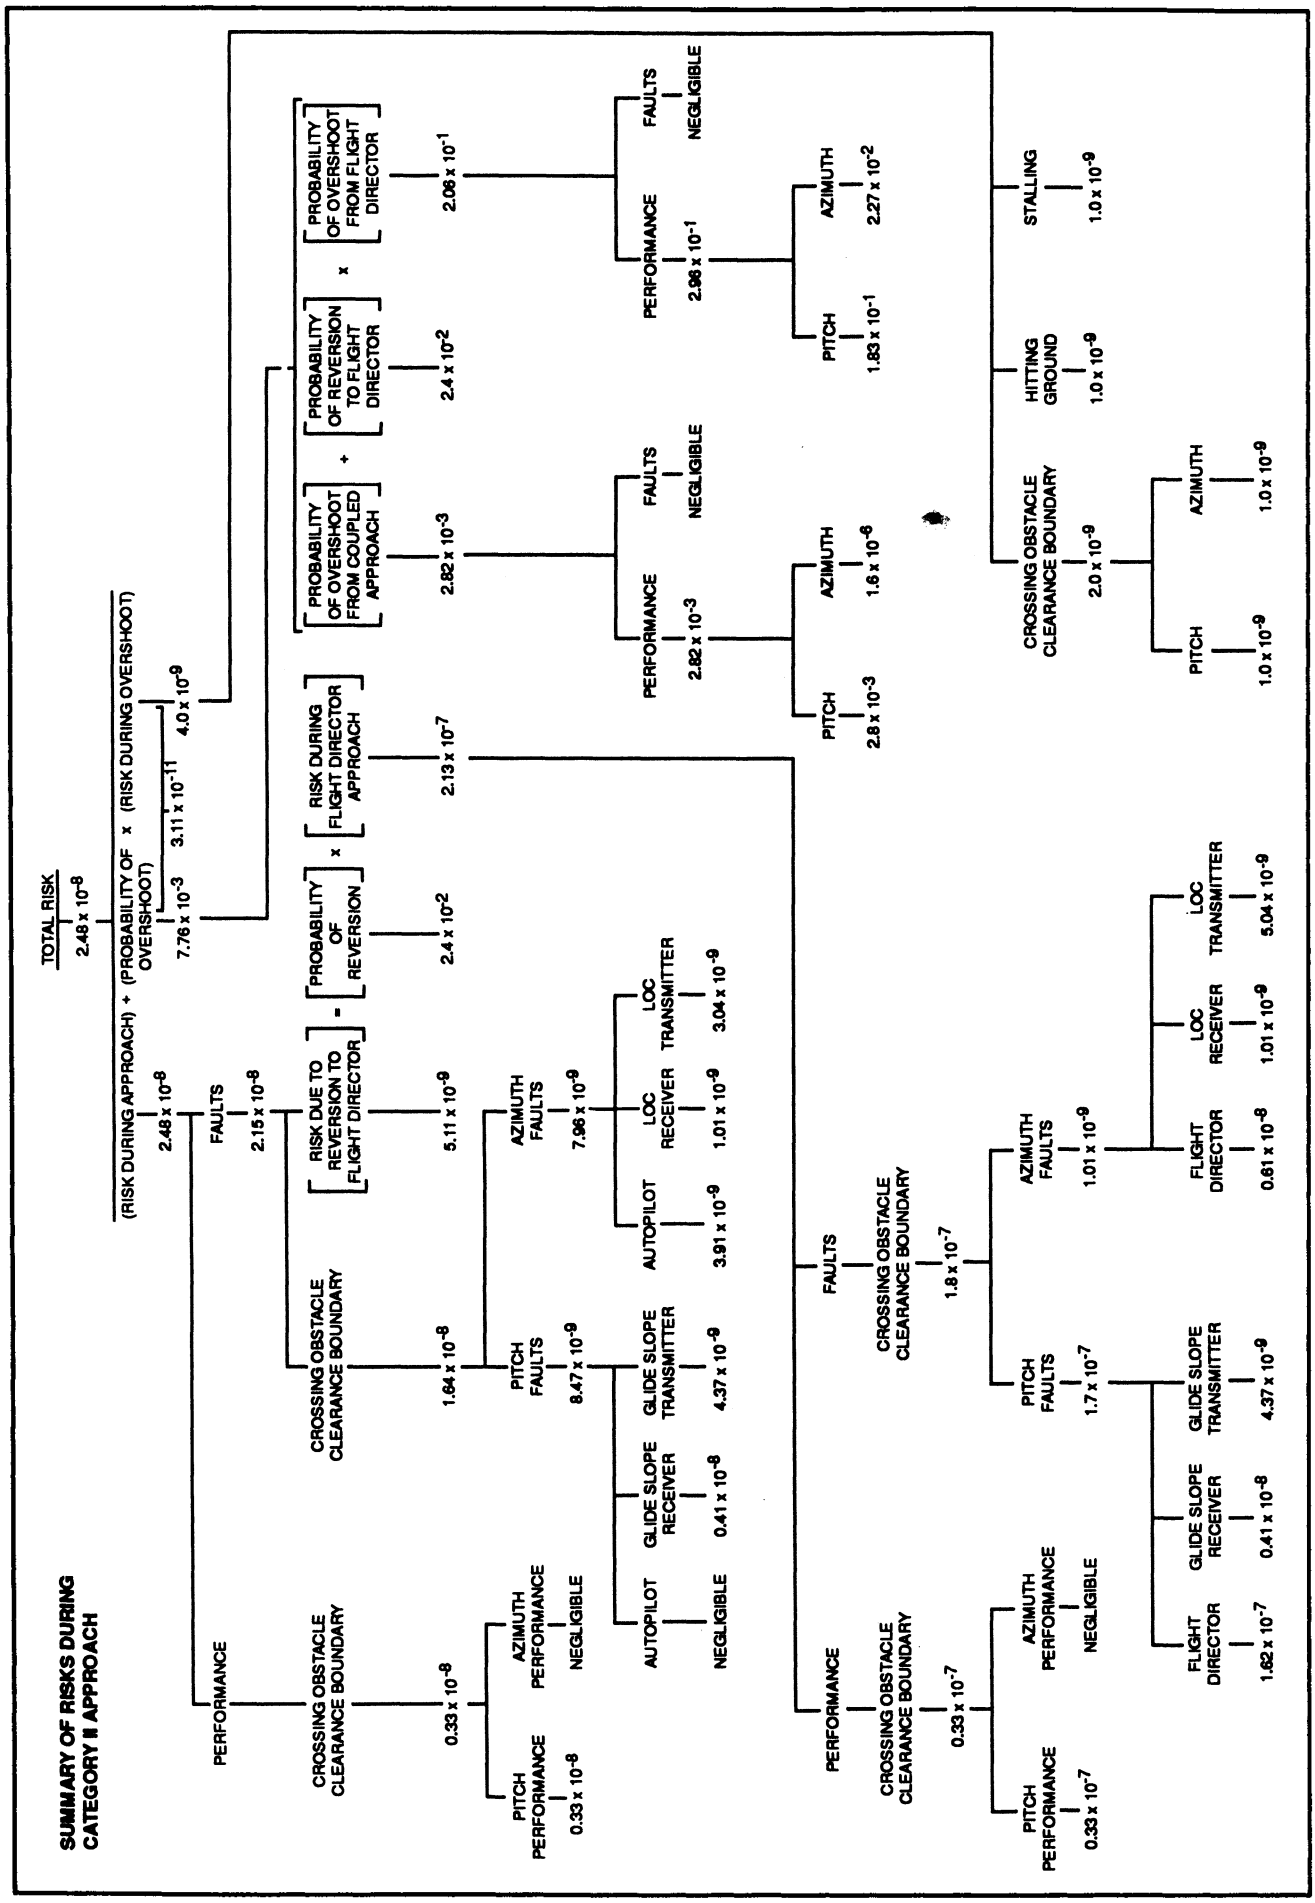

Figure 3. Summary of Risks during Category II approach

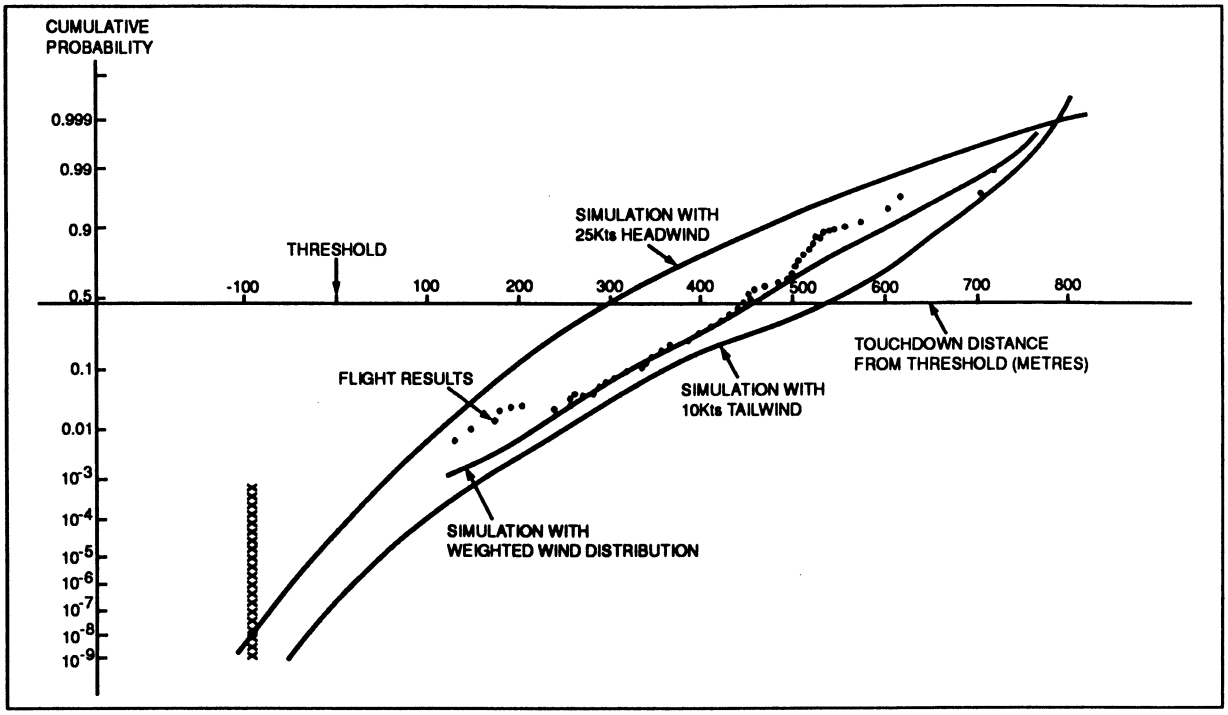

Figure 2. Simulation and flight tests: Concorde

In the CAA/ARB this experience resulted, in June 1970, in Issue 3 of BCAR paper 367 which included restricted visibility operation down to Category 3. This paper, and the subsequent experience gained in its application, resulted in BCAR Paper No. 670 Issue 1 dated September 1976, which was then used in collaboration between the European member-countries as a basis for the relevant Joint Airworthiness Requirement (JAR). I refer of course to JAR-25.1309, well known to all designers and certification authorities. The Advisory Circulars to JAR-25.1309 give specific design guidance for safety assessment against various levels of hazard.

These have been used primarily for airworthiness assessment of individual "safety critical systems". It is the principles outlined which I believe should now become more widely used and extended to the safety assessment of total air transport operations.

At this point I should mention that there is one important aspect related to this, inherent in the automatic landing assessment example, which complicates the task of transferring the technique to all operating sectors of a total transport system.

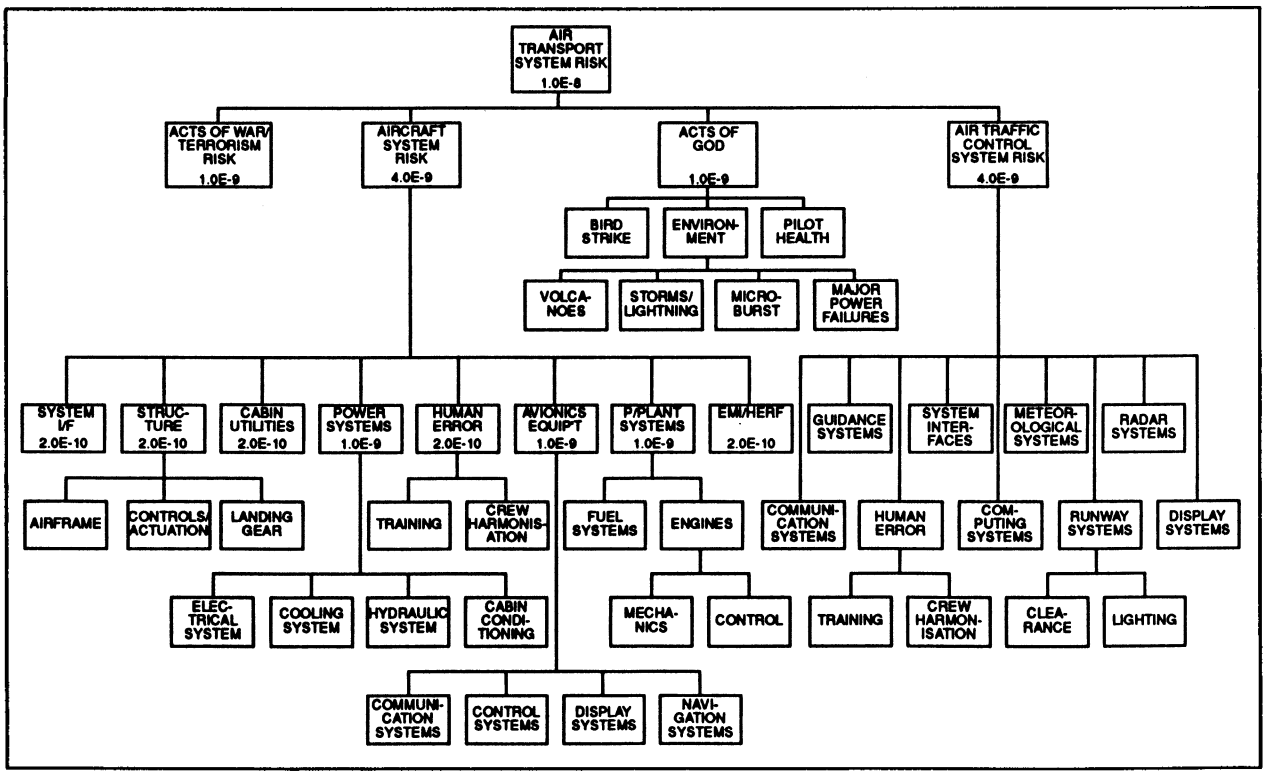

Figure 4. Air Transport System Risk

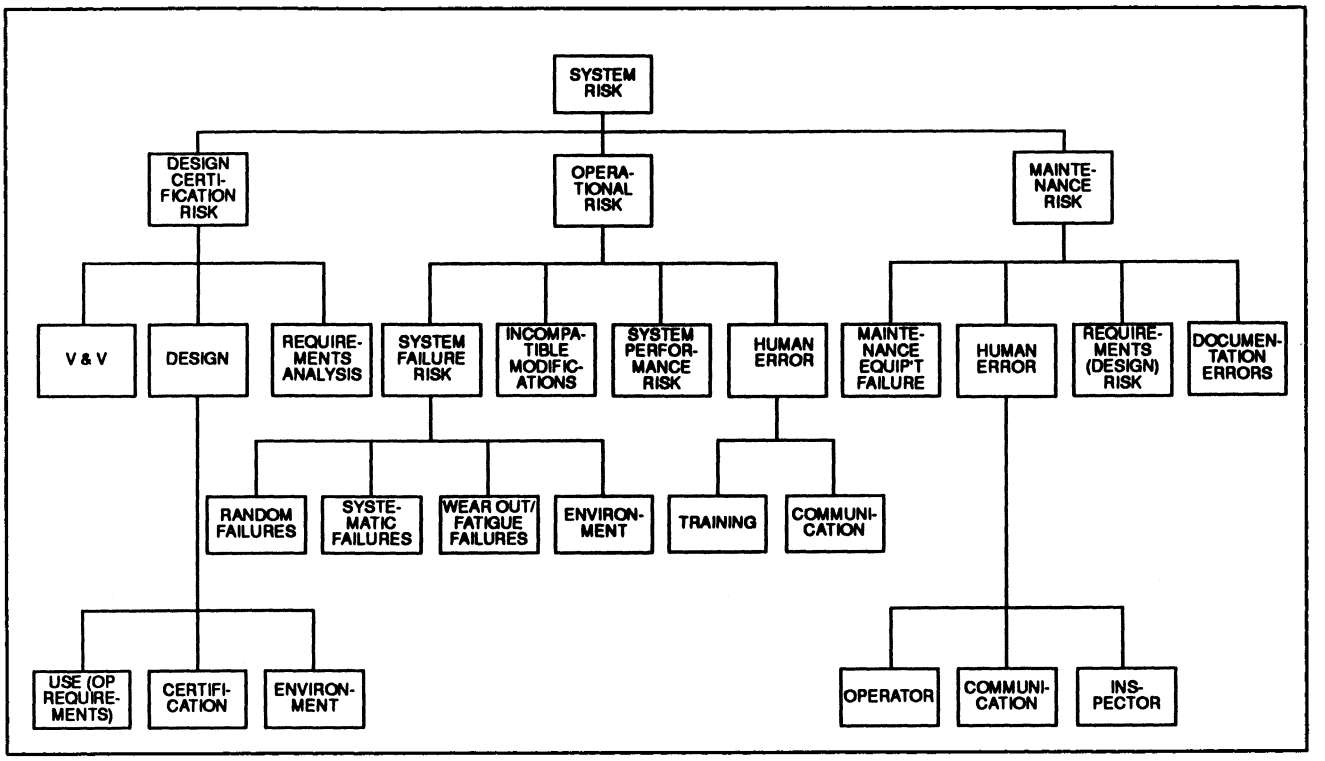

Figure 5. System Risk

This is that automatic landing has no "pilot in the loop" flight crew involvement after setting up is complete and the system is engaged. Hence human error rates are not required to be included in the risk budget.

Systems performance assessments having the human involvement of flight crews and A.T.C. controllers are impossible to speed up on simulators unless the human involvement is simple enough to be itself reasonably simulated. Otherwise all performance assessment must be in real time.

However, despite these difficulties, the application of safety assessment analyses to a total transport system, starting with a 1 in 10^8 per hour target overall risk figure, and including "human-error" elements, might look something like Figure 4.

This total Air Transport system risk may be broken down into a number of contributing risk areas, examples of which are included. Of these, the aircraft system and the air traffic control system are probably the major parts. Each of these top level systems may be further sub-divided into a number of systems each of which has associated with it a number of risk contributions. This decomposition of total risk into risk elements is illustrated in Figure 5 which shows that, in general, each system risk has contributions from the system design, its operation and its maintenance. It will be noted that the operational risk includes both equipment failure effects and performance, as was indicated in the earlier automatic landing example.

Each lower level system contributes to the total system risk during different phases of the aircraft operational cycle.

For a significant improvement to be made to the present accident rate it will be necessary to address quantitatively all risks in all operational phases. Some of the phases are shown in Figure 6. It will be appreciated that failures can occur which may not immediately be associated with flight risk, such as during equipment storage and in maintenance, which can lie dormant and manifest themselves at a later more critical time, and these are included.

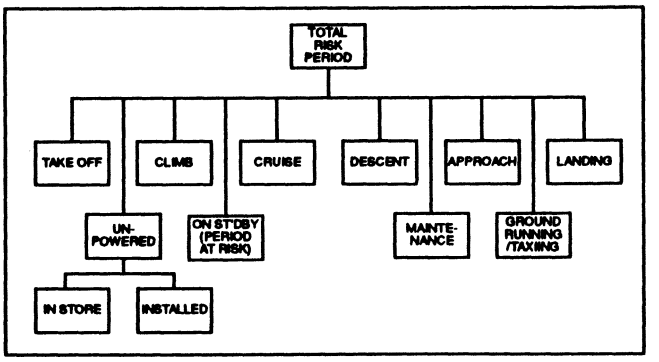

Figure 6. Total Risk Phase

For a particular hazard not all systems necessarily contribute to the risk. For example if we consider the collision risk related to vertical separation of aircraft in mid-Atlantic, then the primary systems involved are the altimeter and its setting and communication between the Flight crew and Air Traffic Control.

Figure 7 shows Figure 4 shaded to indicate the systems contributing to this air collision risk.

It must be stated now that Figure 4 cannot of course be used in its present form to calculate the overall risk even if the lower level contributions were known. The reason is that some lower tier elements in the chart such as aircraft structure, engines and power supplies can have a direct and immediate effect on aircraft safety, and become series elements in a reliability block diagram, their risk contributions being additive.

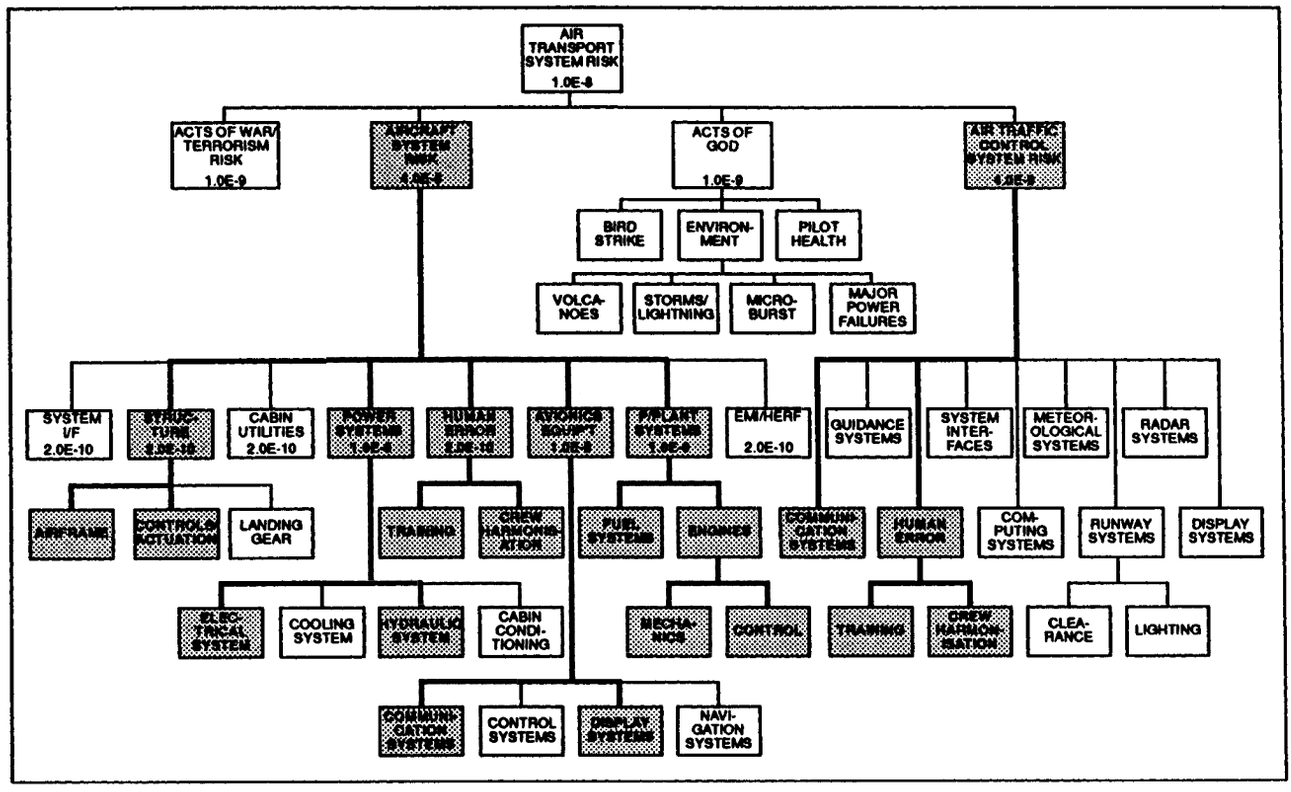

Figure 7. Risk Diagram - Vertical Separation in Cruise

Other elements do not have this direct effect. Mostly they have evolved to provide improved performance, or redundancy to reduce the risk inherent in the simpler more basic functions. Such systems would have to fail in combination to result in catastrophic effects, and as such they are parallel elements in the reliability block diagram, and their probabilities of occurrence are multiplied together.

The diagrams in Figures 4 to 6 are not totally sufficient, but are meant to indicate the complexity of systems and the factors that are used to quantify risk budgets for a total transport system.

To calculate total system risk it is necessary to analyse the system for hazards associated with each of the system functions and to create a risk tree for each hazard using the detailed knowledge of the system elements, their interactions and failure modes.

5. Target Areas for Improvement

Accident statistics and incident reports are constant reminders of the areas of operation in which improvements are required.

These data have indicated clearly for many years that the major area where a significant improvement in safety is required is in the "near airport" category of operations. It is here that flight crew are most active, both in aircraft management and communication with air traffic control.

Statistics show that 50% of all hull accidents occur in final approach and landing and 15% in take off and climb-out. They also show that "flight crew" are given as the primary cause in two-thirds of all fatal accidents.

These figures have also been relatively constant for many years in keeping with the "plateau" in the current overall accident rate.

Unfortunately the reaction to such figures all too often is either that "statistics can be used to prove anything" or "flight crew" or "pilot error" is an all-too-useful way to attribute blame! Yet even a cursory study of fatal accident statistics will show that the "pilot error" element is hardly unreasonable, when viewed in relation to overall human error probability. For example, the two-thirds of the total accidents attributed to "flight crew", when applied to the critical "near airfield" sectors, can imply an average "crew" failure-rate as low as 1.7 in 10^6 over the period involved.. If crews already achieve this as an "average rate" in making errors, then to attempt any fundamental improvement may well be futile. It follows that additional means must be found to achieve the significant improvements which I maintain are now required.

So while it is essential to analyse and take action on all accident causes and incident reports to prevent recurrences, it is unlikely that the overall accident rate can be significantly improved by this. This is because at the 1 in 10^6 per hour level the chances are that one corrected rare event will merely be replaced by another rare event, as yet not experienced. In an ASRS (Aviation Safety Reporting System - NASA) Narrative, one pilot said, "I am sure I will do something else just as dumb in the future - but it wont be the same thing"[3]. It should not be a surprising comment. A 1 in 10^6 per hour event is one which a pilot has less than a 2% chance of experiencing in the whole of his flying career!

So, at the low levels of risk now achieved, it becomes increasingly necessary to have the means to survive failures deriving from "human-error", as well as to eradicate the causes of those which are discovered.

So an alternative attitude to the "pilot error" element in accident statistics is to take a "pro-active" view of the

figures. As discussed previously, we should employ the accident and incident reports to construct incident-survival operational models and fail-operative total system architectures. These, if promoted in a top-down concept for air transport, including the flight crew and ATC operators, should have a substantial effect in improving actual safety. By this means it can be ensured that unreasonable demands are not made on the flight crew or ground controllers, and that their contribution to overall risk is acceptable.

There is no doubt in my mind that the relative flatness of the world airline accident rate is to some extent a measure that the ultimate limit has been reached of error-free human involvement, in a very complex task environment. Hence, further to the previous observations, if substantial improvement in safety in the future is to be achieved, massive automation and information assistance to the flight crew is required, coupled with the implementation on the ground of parallel concepts in air traffic control means.

Charles Billings has referred to such a concept in the aircraft as Human-centred Automation.[3] I would expect the concept of Human-centered Automation to extend outside the aircraft into air traffic control and hence in principle to the total world Air Transport system.

In the following sections a range of automatic and cockpit systems is considered which, in keeping with the views expressed, could make significant contributions to the future safety of air transport operations.

6. Assistance to the Flight Crews from Existing Technology:

6.1. The Use of Automatic Landing

As 50% of accidents occur in the final approach and landing, it would be pertinent to assess what the effect might be if automatic landing were adopted by all aircraft for all landings. At present the influence of automatic landing on the accident rate is negligible.

Automatic landing systems are installed on less than 30% of current scheduled commercial aircraft and on these it is used only in 5 to 10% of landings. Flight crews tend to use it mainly in low visibility conditions, but otherwise, understandably, they prefer to do manual landings in order to maintain the "currency" of their skills. Hence automatic landing is probably used in less than 2% of all landings. The level of use is also affected by the relatively small number of runways in the world equipped to the necessary standards.

Automatic landing was developed in the early 1960s to increase the safety of landing in poor visibility. The risk was expressed in terms of fatal accidents per landing. If the automatic landing facility was not available, the safe alternative was to divert to another airfield. The period at risk was therefore small, typically the three minutes to touchdown after confirming that the facility was available and committing to a landing.

This low period at risk allowed the target safety of less than one fatal accident in 10^7 landings to be achieved, even with the relatively low reliability of the equipment available at that time.

If automatic landing was to be relied upon to increase the safety of all landings, regardless of weather conditions, the system must remain fully available from take-off to touchdown, with the period at risk thereby increased to the mean flight time.

This is not an impossible hurdle. Automatic landing systems installed in the new generation of transport aircraft have two advantages over their earlier counterparts. Firstly, they are digital rather than analogue, and have more accurate airborne sensors and ground guidance installations. This gives a higher performance and hence a lower risk. Secondly, the equipment is also considerably more reliable. Taking these into account, it is not unreasonable to expect an improved performance by a factor of 5 and improved channel reliability of 20.

The effect of including these two improvements in the risk budget (refer to Figure 1) suggests that the overall system design accident probability in final approaches and landings could be reduced from 1 in 10^7 to approximately 1 in 10^8.

If such systems were installed and used for all landings, it would imply an accident rate about 30% lower than is achieved in current operations in that sector of a flight, but this would amount to less than 10% on the existing automatic landing equipped fleet.

Hence the universal adoption and use of "automatic landing" would appear to offer a significant but not large improvement in safety. This could however be a conservative assessment. It can be said that automatic landing and its method of use, does give protection against the majority of the approach and landing hazards, in that once set up and engaged, which could in many cases be as early as during part of the let-down, it would dominate the risk budget to the exclusion of other elements. It also has the advantage that its safety of operation, and any future improvement, is quantifiable.

6.2. Fly-by-Wire

As the use of Fly-by-wire increases in future, many contributions to safety should accrue. It is likely to give a net improvement in the incidence of fatal accidents due to jammed or severed controls, and it can incorporate a wide range of manoeuvre or incidence limiting control laws as protection against mishandling in emergency conditions. These are being increasingly used in new aircraft. Fly-by-wire also offers the possibility for achieving common or nearly common handling characteristics between different aircraft types, with consequent pilot training advantages, including a reduction in the need to maintain type-currency.

While quantification of the improvements offered in future by existing Fly-by-wire concepts is not entirely possible, it should make a considerable contribution to the enhancement of safety in the long term.

There are however further prospects.

Fly-by-wire also offers the means in the future of using different combinations of primary and secondary flying controls and throttles to assure continued safe performance after a major controls failure.

With the pre-provision of suitable computing control laws, reconfiguration of controls could be done automatically to counteract the effect of flying control or engine control losses due to rare failures, structural failure or even sabotage. Proof of concept flight testing of such a system has already been carried out by NASA on an F-15 aircraft. Such a system could be entirely automatic in operation, such that the loss of any capability in controls or engines, or extreme failures could immediately be compensated. Suitable displays would be necessary to alert the crew to the actions taken and give an assessment of any further changes necessary to achieve safe completion of the flight.

It should well be possible, from a study of controls-related fatal accidents, to quantify the potential improvement in safety offered by such a capability.

The need for reconfiguration as described above could normally be determined from information about response or otherwise of any control (effector) to normal commands. Further information might be obtained from the processing of images from external cameras now being installed on many aircraft.

6.3 Air Navigation Systems

Considerable effort on a world-wide basis is being devoted to the requirements for upgrading air navigation and control systems.

At present these are directed mainly towards improvements in the coverage and accuracy of Communication, Navigation and Surveillance (CNS) systems and the fusion of these into satellite-based Air Traffic Management (ATM).

The potential improvements in safety which the availability of the new system offers are at the moment however receiving only secondary consideration as the main concepts are progressed.

As the desire for reducing separations on oceanic and continental routes increases, and congestion in terminal areas grows, there will be an increasing need to set safety targets compatible with a total air transport system, as outlined generally in section 4. When a safety budget is constructed, this will almost certainly demand "fail-operative" characteristics with a possible multiplication of information sources, data transmission and interrogation and confirmation communications. It appears likely that the systems now under consideration will provide such capability if fused into an appropriate overall systems architecture.

There are many examples which illustrate this.

In long haul oceanic or continental airspace operations where Automatic Dependent Surveillance (A.D.S.) systems are relevant, the probability is that large scheduled operations aircraft will have both Inertial Navigation System (INS) and Global Navigation Satellite System (GNSS) installations and possibly other navigation source data, which will in total give both accuracy and fail-operative survivability to the A.D.S. position information. Traffic Alert and Collision Avoidance Systems (ACAS/TCAS) will support this further and be greatly enhanced if other ADS position information on aircraft in the same vicinity were transmitted back to each aircraft, including its own reported position. If all such aircraft position data, and the sources, were subject to statistical risk assessment on both performance and failure aspects as described in section 4, then improved separations could be determined in relation to quantified safety targets.

In terminal areas, approach systems used in radio auto-coupling and automatic landing, at present Instrument Landing Systems (ILS), are fail-operative and have accuracy monitoring. It is possible that INS/GNSS accuracy down to terminal area handover could provide further cross checks on individual position before coupling to the landing aids.

There is a further safety back up due to the presence of ATM Collision Avoidance and TCAS, and the currently vital Ground Proximity Warning System (GPWS). The larger scheduled aircraft will therefore progressively be provided with a dissimilar multiple fail-operative capability in the sector of operations where the accident risk is highest. This should have a significant effect on the reduction of the incidence of Controlled Flight into Terrain (CFT) in near terminal "red sector" areas, with sufficient redundancy to eliminate nuisance problems.

On communications there is now the possibility, in terminal areas having SSR Mode S Data Link, to duplicate voice clearances by up-link transmission to displays or printers, using unambiguous data formats.

Cockpit map displays will in future also play an increasing part in the validation of data coming from multiple sources, in presentation of ADS returns, and in ensuring unambiguous interpretations. This could make a major contribution in preventing those instances of wrongly-directed avoidance manoeuvres, confused ATC let-downs, turning the wrong way and disorientation in complex terminal areas using a multiplicity of diverse aids.

7. Assistance to the Flight Crews : Long Term Research

The subjects considered under section 6 involved capabilities or systems already within reasonable range of future application. There is a range of research programmes proceeding in the United States and Europe aimed at generating considerable assistance in the cockpit in the longer term. These cover independent "electronic" duplication of flight crew activity, for monitoring purposes, as well as research into "expert system" assistance for more complex operational or emergency situations.

These systems are under development primarily for future military applications, and civil applications are not currently defined. There are however potential safety improvements inherent in the increased information generation and its effective transfer to the flight crew.

7.1. Active Situation Appraisal System

The concept of a robust error-tolerant pilot cockpit interface is being explored to investigate the ways a cockpit might;

- contain the means to recognise an incipient incident,
- present readily assimilated information which might use a selection of media, including indicators,

displays, aural and visual alarms and synthesized speech.

This will require the computerisation of typical operational information and aircraft performance parameters, so that the system can;

- identify an emergent problem,
- diagnose the cause and validate it from independent sources,
- evaluate crew response,
- advise on and/or prompt the remedial action.

The term "robust" indicates that the "Interface Manager" can handle a wide range of situations reliably, with its performance degrading only slightly under the most unusual circumstances.

Initially such concepts would apply only in limited areas. An example might be to give advice on exception handling for certain safety critical situations such as potential fuel exhaustion. Advice could also be given on incipient hazards such as low energy situations, with indications to the crew of the various alternatives and their consequences.

The system might include, as these become available, independent imaging source information such as is provided by Millimetric Wave, or Infra Red sensors, or Laser Radars.

A successful system will require extensive investigation into typical human errors arising due to misinterpretations or stress or excessive workload, for which the NASA/ASRS and UK "Chirp" Reporting Systems and other databases will provide valuable information on the building of models of the human error process.

7.2 Incidents Data Storage and Access

A vast amount of information is derived continuously from incident reports. Over many years general information and 'type-specific' failure information has led both to modifications and flight crew alerts.

The technology is becoming available to allow the automatic coding and prioritisation of such information and this is an area where research into the application of Artificial Intelligence techniques could be beneficially applied to assess the true cause from apparently conflicting information. In future it its conceivable that such information could support a real-time on-board advisory system for emergency situations.

The development of ASRS and "Chirp", or corresponding information, into a format suitable for use in such a system has some appeal.

7.3. "Pilots Associate" Expert System

Work on the US Pilot's Associate Programme and the UK Mission Management Aid (MMA), both military research programmes, has laid a strong foundation for a so-called "electronic crew member" to advise and support human crew during times of exceptional work load.

Such a "crew member" would;

- Fuse or collate information from many aircraft systems and sensors, associating related items of information,
- Assess this information against the flight plan and current data and give an assessment of the situation,
- Provide an advisory function as required in the formulation of alternative plans and present these to the flight crew.

The Pilot's Associate is currently a large scale multiple computer workstation-based simulator using expert systems as decision makers and advisers.

The MMA is also a multiple workstation-based simulator using conventional software and some "rule-based" systems techniques.

As yet the computing is not available at a suitable size, weight and cost, but at the current state of progress with Reduced Instruction Set Computer (RISC) machines or equivalents it should be available within 4 to 8 years. The remaining hurdle will then be the certification of the relevant safety-related "rule-based" software.

Such systems might also find use in the longer term future in Air Traffic Management.

8. Summary and Conclusions

- The safety of air transport operations needs to be improved in the future and agreement must be reached on a new target and areas to be pursued in order to achieve it. A world-wide fatal accident rate of at least 1×10^{-7} per hour must be achieved, and perhaps 1×10^{-8} per hour or lower should be used to determine lower-tier system element design targets.
- Particular attention should be given to the improvement of "near-airport" operational safety by significant increases in both flight automation and cockpit assistance for the flight crew.
- Airworthiness, Air Navigation and Air Traffic Management safety requirements should be fused into an integrated systems approach, with all safety aspects of design quantified by the universal use of standard statistical safety assessment methods.

Acknowledgements

The author acknowledges the assistance given in the preparation of this paper and the assembly of background data by Messrs G Belcher, R I Bishop, T G Hamill, M F Moulton, F G Oates, I S D Stitt, J R Taylor and M J Tooze. His thanks are also due to the UK Civil Aviation Authority and the Boeing Company.

The views expressed in the paper are entirely those of the author and not necessarily those of GEC Avionics Limited or others.

References

1. Safety in the Air. Ronald Ashford. RAeS/IMechE Joint Lecture, Belfast 17th February 1988

2. Progress in the use of automatic flight controls in safety critical applications. R W Howard. Fifth European Pioneers Day Lecture, Braunschweig 29th May 1980. (Published RAeS Journal October 1980)

3. The view through the looking glass - at ourselves : The Importance of incident reporting for aviation safety. Charles E Billings M.D. RAeS Wright Memorial Lecture. 13th December 1990