Progress in the use of automatic flight controls in safety critical applications

R. W. HOWARD, BE, CEng, FRAeS

Director and General Manager,
Marconi Avionics Ltd

1. INTRODUCTION

High authority automatic flight control systems are now in wide use in both transport and combat aircraft. These have required the development of failure survival system architectures, and techniques in safety assessment to ensure that the risk of fatal accidents attributable to the automatic control aspect is acceptably remote.

Most of the current systems are analogue in nature, and employ sufficient redundancy to survive at least one failure during operation. Future systems will employ quadruplex or triplex-monitored architectures, capable of surviving at least two failures, and will be mainly digital in concept. The design and safety assessment techniques developed for analogue systems have been adequate for current systems, but the new digital technology and more extensive use of systems in safety critical applications will demand further development of the various aspects involved.

A major contribution to the design and safety assessment of safety critical systems was made during the past two decades in the programmes for the development of automatic landing for civil aircraft, culminating in the system in the Anglo-French Concorde.

The paper covers some of the details of these early programmes as in these, more than any others, rules were developed for acceptable failure rates and performance levels which need to be met for safe operations.

The same basic techniques developed for these early programmes can be used as guidelines for the design and assessment of the new digital systems but some new problems are presented by the different technology.

2. AUTOMATIC LANDING BACKGROUND

2.1. The requirements

2.1.1. *Early ARB work with UK industry*

In the United Kingdom in the early 1950s, industry, the two main airlines and the Air Registration Board combined their efforts to devise satisfactory requirements for the introduction of automatic landing into two new civil air liners, the De Havilland Trident and the Vickers-Armstrongs VC10.

This resulted in the issue by the ARB in November 1961 of the British Civil Airworthiness Requirements (BCAR) paper 367 'Airworthiness Requirements for Autoflare and Automatic Landing'.

The principal contributions of this BCAR were the concepts of a numerical safety level for automatic landing system certification and the need for an analysis to demonstrate its achievement.

The level of safety required

The ARB defined a figure for the required average fatality risk of 10^{-7} per automatic landing, on the basis that this was a reasonable proportion of the current risk in pilot controlled landings of $0{\cdot}65 \times 10^{-6}$.

In addition to the average risk figure, the ARB also specified later a corresponding maximum specific risk figure of 3×10^{-6} per automatic landing.

It was clear from current data on autopilot system failure rates that risk levels of this order could not be achieved without using equipment redundancy to obtain an automatic failure survival capability. Therefore for the Trident, Smiths designed a triplex automatic landing system, and for the VC10, Elliotts designed a duplicate monitored system, each aiming to give, as a minimum, a single failure survival capability.

From 1961 to 1970 considerable experience was gained in the UK from the certification programmes for these aircraft which advanced significantly the state-of-the-art in designing and analysing high integrity systems.

In the ARB this experience resulted, in June 1970, in Issue 3 of BCAR paper 367. This paper, and the experience subsequently accrued in its application, has been used in collaboration with our European partners as the basis for the development of the relevant Joint Airworthiness Requirements (JAR).

The value of the paper is greatly enhanced by the large quantity of guidance material added as appendices. These cover failure analysis, performance analysis, mathematical modelling, reliability and integrity calculations and suggested methods for incorporating these features into a system safety assessment.

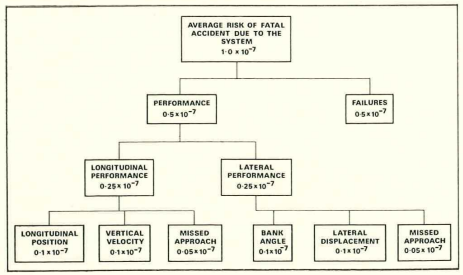

Figure 1. Sample breakdown of average risk.

Paper presented at the Fifth European Pioneers' Day held in Brunswick, Germany on 29th May 1980. Paper No 831.

2.1.2. *Further development of requirements during the UK automatic landing programme*

Partitioning of risk

As automatic landing developments proceeded, a model for the partitioning of the average risk figure over the various risk elements was developed which is also included in BCAR paper 367 (Fig. 1). There are two major groups, one dealing with performance, and the other with failures.

Some elaboration of this is required in practice and one example is shown in Fig. 2, prepared by Smiths. Some flexibility must be allowed in risk partitioning, but it is important that no single contribution is disproportionately high if there is to be confidence in the overall assessment.

Neither failure nor performance risk levels achieved can ever be proved in any practical series of flight tests, or even in service flying, because the levels are too low, and inordinately large numbers of flights would be required. Statistical analyses are therefore needed.

Failure analysis guidelines

The first part of the overall analysis is the failure analysis, and it is now customary to conduct this in a carefully structured 'top-down' manner through three levels.

First, the overall system architecture or basic failure survival concept is checked. This involves an analysis of the physical design, lane segregation, separation of lane power supplies etc. The second level of analysis aims to identify the nature and probability occurrence of all single failures. This is inevitably a long and tedious aspect of failure analysis. The third level of analysis is the quantitative integrity assessment, which uses component failure rates, knowledge of system architecture, layout, monitoring techniques and safety check periods to calculate the total system failure risk.

Failure analyses are normally conducted by design engineers having a good working knowledge of the system, but up to the present little use has been made of computerised data reduction methods for this task.

Performance analysis guidelines

The second part of the overall analysis is the performance analysis.

The concept of performance variation as a risk factor in automatic landing covers such cases as touching down short or wide of the runway, having excessive vertical, lateral or horizontal speed, or excessive deviation from flight path in an overshoot. The numerous factors which affect performance must be expressed statistically and integrated, not only to calculate the average risk of entering a hazardous situation but also the specific risk arising from any one parameter being on its extreme limit.

Many tens of thousands of simulated approaches using a wide range of aircraft, atmosphere and system

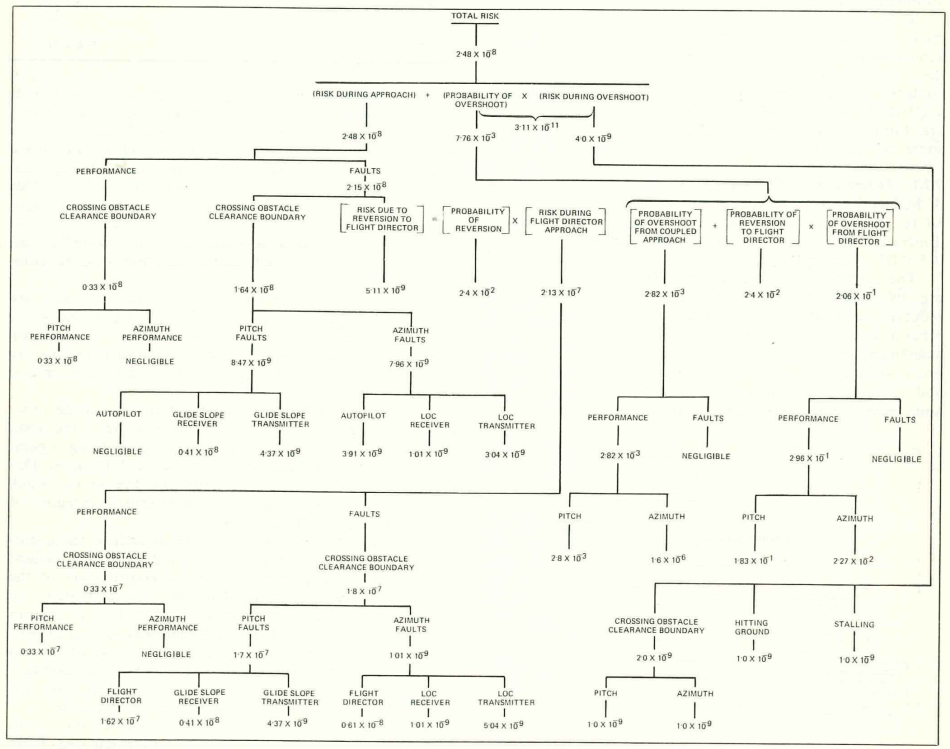

Figure 2. Summary of risks during Category II approach.

parameter variations are necessary to establish performance failure risk levels. Such analysis has been made possible by the existence of suitable hybrid computers and actual flight tests are used only to confirm the accuracy of the simulator model and to test 'worst case' failure conditions.

An important parameter in performance analysis is wind, and considerable effort was expended in the automatic landing programme to devise realistic statistical models.

2.2. Achievement

The application and the development of the new redundancy and assessment techniques in the United Kingdom resulted in some very tangible outcomes in that both the Trident and the VC10 were certificated for passenger carrying automatic landing, and subsequently the Concorde, by UK and France jointly.

By far the most extensive experience has been obtained with the Trident, which has more than 50 000 in service automatic landings to its credit and is approved for operations down to 100 metres RVR with a 12 foot decision height. The 100 metres restraint is imposed for taxying, not for reasons of safety in flight.

The VC10 accrued 3500 automatic landings in service at Cat II equipped airports before use of the system was curtailed in 1974 for cost saving reasons.

The Concorde, to date, has performed nearly 1500 automatic landings in passenger service.

Some representative details of the design and assessment of these systems may bring into proportion the achievements of the automatic landing programmes and the contributions which these have made to later 'active controls systems' developments. The author is best acquainted with the VC10 and Concorde systems, which are similar in concept, and the following sections give some details of these.

2.2.1. *Failure analysis (VC10 and Concorde)*

A budget for the average risk of a catastrophe during a VC10 automatic landing was made up in the manner illustrated in Fig. 3 with system failures allocated $0{\cdot}6 \times 10^{-7}$ per landing.

The system architecture of the VC10 system, comprising two completely separate self-monitored autopilots, reduced the chances of unforeseen critical failure combinations and made the two vital tasks of qualitative assessment and numerical failure analysis relatively simple. Each self-monitored autopilot comprised a main control and monitor lane which were also physically separated, and cross connections between these were restricted to

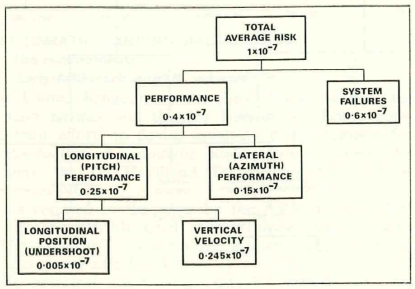

Figure 3. Distribution of average risk: VC10 automatic landing.

those required for signal consolidation (which reduces build-up tolerance discrepancies) and cross-comparison (which is used for failure detection).

The comparators, used for cross-comparison, played a critical role in the system, and they were developed either to be fail-safe or were duplicated, thus reducing the risk of undetected dormant failures.

An indication of the contributions to system failure made by the various risk groups in the VC10 is given in Table I, and the total system average failure probability calculated of $0{\cdot}4 \times 10^{-7}$ was somewhat less than the allocated figure of $0{\cdot}6 \times 10^{-7}$ per landing.

TABLE I
Breakdown of system failures for VC10 automatic landing

Risk groups	Contribution to system failure $\times 10^{-9}$
Elevator servo loop	0·0375
Pitch integration	1
Autoflare logic	9·74
G/Slope computing and flag circuits	2·1
Barometric sensors	0·21
Flare arm relay	0·37
Radio altimeter	15·45
Radio altimeter height switches	1·3
Azimuth computation	0·42
Interlock line by-pass monitors for Cat I	0·72
Special case detected failures	6·60
Special case undetected failures	1·22
TOTAL FAILURE RISK	$0{\cdot}4 \times 10^{-7}$

The Concorde also has a dual monitored autopilot but uses linear integrated microcircuit technology as distinct from the 'discrete component' circuit design of the VC10. It also has digital integrators and an optimised number and positioning of the tolerance reducing consolidation points between control and monitor lanes. This reduces the probability of nuisance disconnects to negligible levels.

A primary feature also is that it is built around a dual electrical signalling system and a full-flight-regime automatic throttle system.

Control and monitor lanes in Concorde computers are also strictly separated, as illustrated in Fig. 4, so as to ensure that any failures are statistically independent. Clearly identified construction principles such as this are also vital to the qualitative analysis of potential critical failure conditions.

The numerical failure analysis of the Concorde automatic landing system was structured so that the risks during various categories of automatic landing operations could be calculated from the same data base. This was different from the method used on the VC10, which was aimed at certification of the most critical category of operation specified (CAT IIIA).

A 'top-down' analysis was made to deduce the events or combination of events which could lead to a critical or catastrophic failure, and then the contribution of the flight control equipment to that risk was calculated. To reduce the work required in analysing the effect of failures within the computers, the equipment was split into self-contained blocks and the effect of a wide range of failures at the input and outputs of the blocks was assessed. Where the effect of any failures was not clear from the theoretical study, actual hardware tests were made.

Using the results of this form of analysis on all the flight control equipment, together with calculated probability of occurrence, it was possible to ensure that signifi-

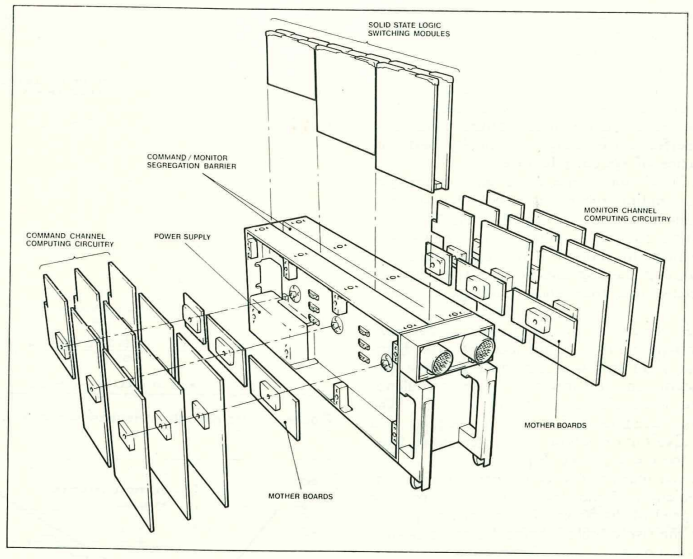

Figure 4. Concorde computer design concept.

cant critical failures were detected and that dormant failures were detected within the established risk periods.

The Concorde failure analysis showed that the targets would be bettered by even more than in the VC10 and this is now supported by the results of 40 000 hours flying and more than 100 000 hours of total equipment operating time.

2.2.2. *Performance analysis*

Turning now to performance analysis, this can only be relied upon if the system model is representative, and a great deal of effort was put in to the establishment of these for both the VC10 and Concorde.

The main elements of the VC10 model are illustrated in Fig. 5.

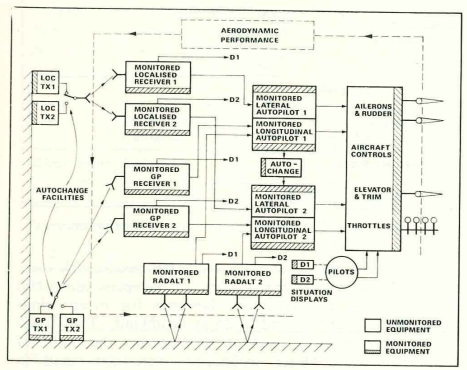

Figure 5. Duplicate monitored automatic landing system: VC10.

In the VC10 case, the model was constructed and matched to the aircraft in a series of progressive tests starting from the basic aircraft response and building up to the total combination of aircraft, autopilot and ILS guidance system. Account was taken of a range of aircraft speeds and weights and the actual hardware dynamic responses including all significant nonlinearities determined from rig (Iron-bird) testing.

The final stage was to compare the nominal total simulated system performance in still air with a series of aircraft landings performed in calm conditions.

The results of these automatic landings as compared with the simulation are shown in Fig. 6. It was concluded after this matching process was conducted that the model was an adequate representation of the aircraft performance characteristics.

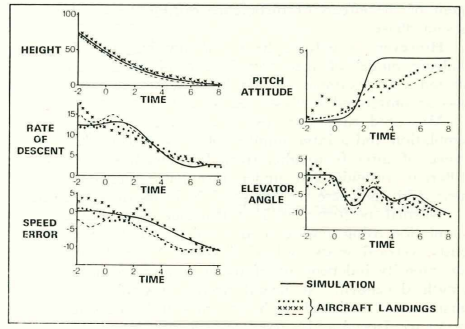

Figure 6. VC10 autoflare model validation.

The second important aspect of the simulation was the wind model, and a standard for automatic landing certification was derived from the best information available in the UK and the USA. Industry, the RAE, and the ARB consulted together in making the choice.

The certification analysis of performance consisted of three parts. First, performance under failure conditions, then a normal performance study, and finally performance in the presence of system tolerances.

One of these, in relation to the VC10, was the analysis to determine the limiting mean windspeed that allowed the safety requirements to be satisfied. For the VC10 the initial analysis indicated, as expected, that the distributions of the critical touchdown parameters were not gaussian and also, that the risk contribution from descent rate greatly exceeded that from short landings or undershoots.

From these early results, an algorithm was deduced, verified and empirically fitted to the descent rate distributions. The final algorithm was a normal gaussian distribution, dependent on performance under turbulence conditions, up to a certain level of descent rate. The portion of results exceeding this level were assigned to a second gaussian distribution whose parameters were fixed as a function of the control laws (Fig. 7).

It was considered that the two component distributions gave a good prediction of the descent rate at impact up to a mean windspeed of 20–30 kt. Touchdown range was less significant in the risk calculation and for this a normal distribution was assumed.

Having developed this algorithm which calculated the risk from the mean and variance of the fundamental component distribution, it became possible to calculate the risk from very small sample sizes. For most cases samples of only 50 landings were used (Fig. 8).

Integration of the conditional probability of fatality with windspeed, coupled with the probability of that windspeed, gives the average risk at any windspeed. If we take an average risk of $0{\cdot}245 \times 10^{-7}$, the portion allocated for heavy landings in Fig. 3, and allow 1 in 10 as the risk of fatality if the designed undercarriage limit of 12 ft/sec is exceeded, the allowable risk is $2{\cdot}45 \times 10^{-7}$ which corresponds in Fig. 8 to a mean windspeed of 18 kt. At this windspeed the specific risk is well inside the requirement and therefore 18 kt was set as the operating limit.

The performance risk analysis for Concorde certification followed similar lines to that for the VC10, however, there were some differences in philosophy.

The design requirement was for certification to 25 kt mean windspeed and this entailed the analysis and assessment of the effects of turbulence during the control law design phase.

However, as before, the initial aim was to achieve a realistic model of the aircraft, control system and ILS station. An example of the high degree to which this was accomplished is illustrated by Fig. 9.

Measured gusts were used as forcing functions to the simulation and a large number of tests were made over a range of aircraft weights from light to heavy, at three different airfields and under turbulent conditions from near calm to severe (in excess of 20 kt mean windspeed).

Normal performance in turbulence was analysed by simulation using a wind model which included horizontal gusts, vertical gusts and wind shear. Vertical gusts had an intensity independent of mean windspeed and a scale length dependent on height above ground. Horizontal gusts were modelled with an intensity proportional to mean windspeed and a fixed scale length. Wind shear was a logarithmic function of height above ground.

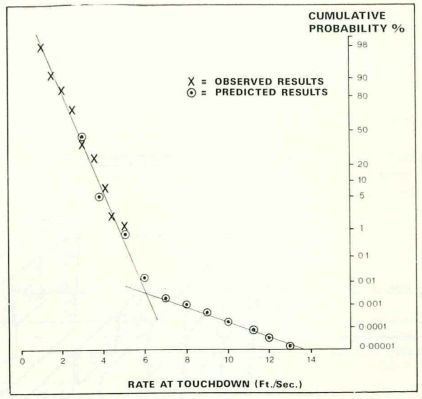

Figure 7. Height rate at touchdown: cumulative probability.

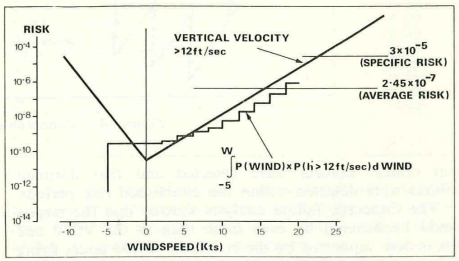

Figure 8. Risk of $\geqslant$ 12 ft/sec vertical velocity at touchdown.

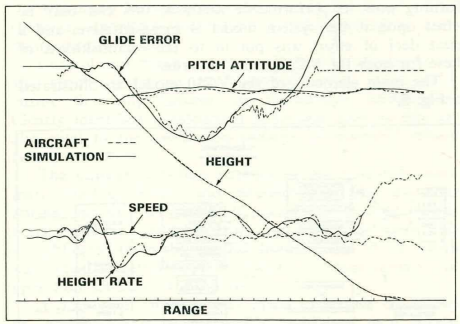

Figure 9. Concorde dynamic model validation.

Initial estimates of average risk and specific risk were calculated by performing groups of samples of 10 000 and 1000 simulated automatic landings for mean windspeeds from 10 kt tailwind to 25 kt headwind. The 10 000 landing samples were used to help predict the 'tails' of the distributions. Other techniques were also used to increase confidence in the distribution tails by cross correlation of associated parameters.

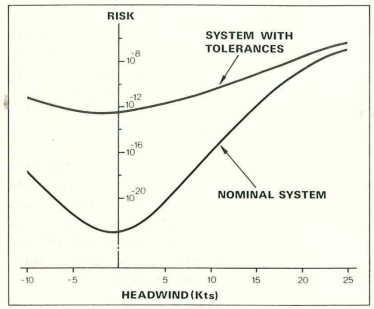

Figure 10. Conditional probability of fatality versus windspeed.

To take account of parameters variability and tolerance a Monte Carlo approach was adopted. This consisted of assigning distributions to all aircraft, control function and ILS ground based parameters.

For these Monte Carlo systems the analysis was repeated over the range of mean windspeeds and a new risk curve obtained. These results and the combined result (touchdown distance plus vertical speed) for the nominal system are plotted in Fig. 10.

This allowed a true average risk to be computed taking account of typical variation of all other parameters. As can be seen, at low mean windspeeds the variation of all parameters adds considerably to the risk. However, at high windspeeds the effects of turbulence became increasingly significant and the two results became closer. At 25 kt the conditional probability of fatality is less than an order of magnitude greater than for the Monte Carlo analysis.

Comparing the Concorde TSS 1-2 requirements for Category III certification against simulation predicted performance (Table II) shows that the requirements are well satisfied.

All parameters are either an order of magnitude better than the TSS 1-2 requirements for risk parameters, or within 60% of the TSS 1-2 requirements for performance parameters.

Comparing aircraft results with simulation results also gives good agreement (Table III).

The aircraft parameters were measured from more than 150 automatic approach and landings made during the certification phase. The reported mean windspeeds varied from 23 kt headwind to 10 kt tailwind (in fact most of the approaches were in high winds) and the approaches were evenly split between three airfields with their respective ILS systems.

The fact that no missed approaches occurred due to excessive beam deviations is not considered statistically significant for the sample size.

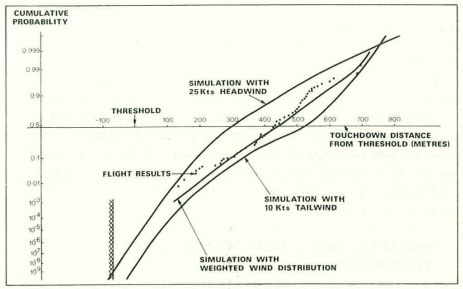

Figure 11. Simulation and flight tests: Concorde.

TABLE II
Simulation results versus requirements: Concorde

Parameter	Concorde simulation	TSS 1–2 requirement
Specific risk at 25 kt — nominal system	$3{\cdot}0\times10^{-8}$	3×10^{-6}
25 kt, system with tolerances	$1{\cdot}6\times10^{-7}$	
Average risk Distributed wind, nominal system	$1{\cdot}8\times10^{-10}$	**Typical apportionment** 2×10^{-8}
Average risk Distributed wind, system with tolerances	$1{\cdot}7\times10^{-9}$	
Missed approach rate Due to beam deviation	2%	<5%
Glide error deviation 500 ft to 100 ft	60% of TSS 1–2	
Vertical speed deviation Average 200 ft to 100 ft	1·5 ft/s	2·8 ft/s
Limiting wind	1·8 ft/s	3·0 ft/s

Parameter	Simulation Mean	Simulation Standard deviation	Aircraft Mean	Aircraft Standard deviation
Glide beam deviation At 100 ft	0 μA	28 μA	−3 μA	25·3 μA
Vertical speed At 100 ft	−12·5 ft/s	1·5 ft/s	−12·3 ft/s	1·54 ft/s
Touchdown range Past threshold	1480 ft	295 ft	1540 ft	310 ft
Touchdown impact rate	2·6 ft/s	1·0 ft/s	1·7 ft/s	0·5 ft/s

TABLE III
Simulation and flight test performance results: Concorde

A plot of the cumulative probability of landing distance for the aircraft flight test results against the Monte Carlo variable simulation results show good agreement (Fig. 11).

The simulation results reflect the combined probabilities with windspeed of normal variations and spreads of all relevant control system, aircraft and ILS parameters.

The figure shows the envelope for all parameter tolerances except windspeed for windspeeds of 25 kt headwind and 10 kt tailwind and also the results for a distributed windspeed. Superimposed are the individual results of the aircraft landings.

3. PROGRESS WITH NEW TECHNOLOGY AND TECHNIQUES

The earlier assessment methods for safety critical control systems have been outlined, as these form the background against which new methods will be developed for systems using new technology and new control techniques.

Several aspects of the design and assessment of new systems are of interest.

3.1. New redundancy architecture

Active control systems operating over a whole flight, as distinct from the limited period required in an automatic landing, must have a high availability as well as the ability to survive failures. Also unlike automatic landing systems they cannot rely on a reversion capability such as overshooting.

This will lead in many applications to the use of quadruplex or triplex monitored architectures which, subject to the integrity of the basic design implementation, can give the capability of surviving at least two independent failures.

3.1.1. *Quadruplex*

The quadruplex concept is illustrated in Fig. 12.

Data from sensors is fed to the associated computer and then transmitted across lanes to the other computers. The computers vote on the data and form the output commands. These output commands are then consolidated in the actuation system. This arrangement gives good lane segregation since the only signals crossing the lanes are restricted to those on the interlane links. Failure monitoring depends only on the detection of differences between like signals. It uses identical computers with identical software and this reduces the costs of design, manufacture and support.

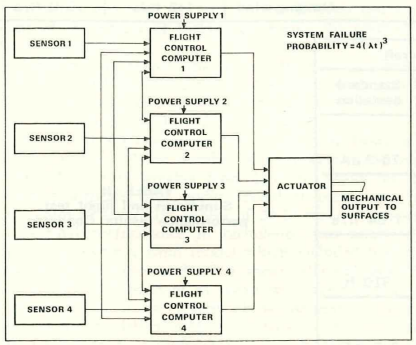

Figure 12. 'Quadruplex' system.

If an individual lane in a quadruplex system has a probability of failure in time t of $\lambda \times t$ and there are no common failures, then the probability of losing three lanes, and therefore the total system, is $4 \times \lambda^3 \times t^3$. For a one hour flight with a lane failure rate of 0·001 per hour the probability of losing the system is

$$4 \times (10^{-3})^3 \times 1^3 = 4 \times 10^{-9}$$

which is more than satisfactory for most requirements.

However, a pure quadruplex system needs four sets of computers, sensors and power supplies and six bidirectional links. To reduce the total hardware required the triplex monitored system is being developed to meet the same overall requirement. This system uses self-monitoring as well as interlane voting.

3.1.2. *Triplex monitored*

A block diagram of a triplex monitored system is shown in Fig. 13. The total system failure probability is approximately $3(1-K)(\lambda t)^2$ where K is the probability of detecting correctly a failure in a lane by using self-monitoring. In a modern system a value of $K = 0{\cdot}9$ is easy to achieve.

Therefore the probability of system loss in a one hour flight for a system with a lane failure rate of 0·001 per hour is approximately

$$3(1-K) \times 10^{-3} \times 10^{-3} = 3 \times 10^{-7}.$$

This is higher than for a quadruplex system but still satisfactory for many applications.

The depth of monitoring is the key feature of the triplex solution and digital implementations have made it possible to monitor with far less hardware than is needed for control.

The form of monitoring now used is illustrated in Fig. 14.

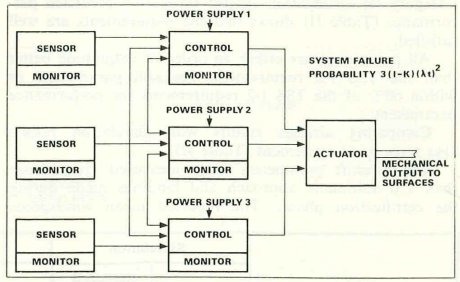

Figure 13. 'Triplex monitored' system.

3.1.3. *Comparison of quadruplex and triplex monitored systems*

The major advantage of quadruplex systems is that it is possible to achieve a very high degree of separation between the redundant lanes, which minimises the probability of common failures, the avoidance of which is paramount in successful redundant system design. Transmission of data across lanes can be achieved without losing segregation by using fibre optic links or by local optical isolators.

A disadvantage is that failure detection relies entirely on voting between similar lanes and such systems are susceptible, albeit remotely, to common mode failures.

The triplex monitored system offers a choice of failure detection capability to the system designer. For example it can use interlane voting for first failure detection and so match the integrity concept of the quadruplex system. Isolation of a second failure would normally be achieved by 'in-lane' monitoring to determine which of the two lanes is still valid. This is a dissimilar method of failure detection for the second failure which is a distinct advantage over a quadruplex system in the ability to avoid common mode failures.

Many analogue systems use in-lane monitoring to detect failures in high integrity applications. The first were in the VC10 and Concorde duplicate monitored automatic landing systems and these were followed by the DC10 and A300, although the computing configurations in these aircraft tend to be 'dual dual' rather than 'duplicate monitored'. In the latter type of system the separate monitoring lane can be much simpler than the control lane as is the case on the VC10 and Concorde. With the advent of the digital computer, designers have exploited the time multiplex nature of the computers to interleave control and monitoring tasks. For certain levels of integrity this is adequate but the possibility of common failures of the hardware or software is the ultimate limitation of this technique. In a recent civil aircraft application monitoring has been implemented with dissimilar microprocessors since the failure analysis of the microprocessors is considered too complex at this time to justify certification where reliance is placed on similar redundancy alone.

It is possible to configure triplex monitored systems with the same degree of physical isolation as quadruplex systems but if the number of line replaceable units (LRUs) and data links is to be minimised then the control and monitor must be packaged in one LRU. This makes the system more prone to common failures due to environmental hazards. Care must be taken to control these hazards and failure analyses of relevant factors such as wire segregation, battle damage, heat input, must be examined critically.

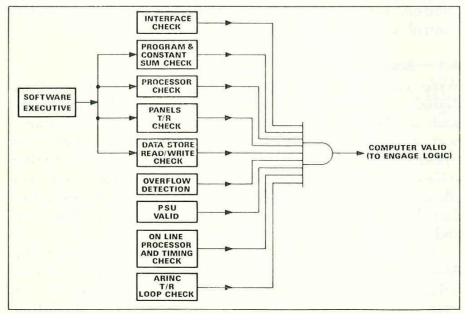

Figure 14. Monitoring schematic.

3.2. Future failure analysis techniques

The hardware failure analyses carried out on the VC10 and Concorde were long and tedious studies and consumed a large proportion of the total effort devoted to systems analysis by experienced project engineers. As active control systems become more complex it will be necessary to automate this process.

Much progress is being made to this end. System description languages applicable at an architectural level are under development, and in some cases in use for process control, and development of such a capability should be pursued more energetically in the field of active controls.

Similarly, at the next level, component or 'device' failure, diagnostic software is widely used for computer analysis in the automatic testing field. Such programmes could be implemented to cover some aspects of a failure mode and effects analysis (FMEA) of active control systems, but further development will be necessary to give a fully comprehensive analysing capability for systems operation. Moreover a high speed FMEA capability could enable failure analysis to be conducted in parallel with systems design and thus give a better interaction between basic design and failure assessment at an early stage.

The failure analysis of software is perhaps the most difficult problem facing the designers of future digital safety critical systems. At present there is no known method of quantifying the risk of a software error, and current safety assessment approvals are based upon a knowledge of the adequacy of procedures used for writing and controlling software.

Software now enters into design in many areas which were once the preserve of hardware implementation only, such as in complex control strategies, accurate complex schedules, sophisticated redundancy management and powerful selftest. The impact on design control is widespread.

3.2.1. *Software control*

The software development process is illustrated in Fig. 15. The figure shows how the process is broken down into distinct stages and how each of these stages is checked by testing.

The techniques used by most flight control system designers is to segregate the software into separate modules which can be visually checked and thoroughly tested. The level of segregation is determined by the ability to comprehend each task and it can result in a large number of modules.

The process of producing safe software may be likened to that of building safe airframe structures, where detailed inspection is necessary at each critical stage of construction. Also because of the real time aspects the software engineer is unable to split the program into functional elements. Instead he must split it into time frames with only part of a functional task complete in each frame. This complicates the organisation of the program and leads to very rigid structures to ensure that tasks are performed at the correct time.

Because of the problems of ensuring safety, most flight critical systems are programmed in assembler language.

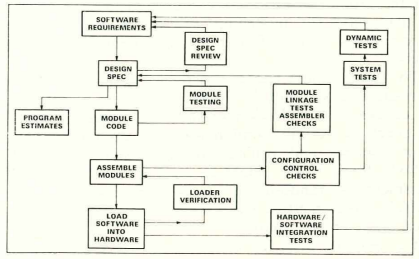

Figure 15. Software development structure.

This ensures a one-to-one correspondence between the source code and the machine code, and enhances store and run time efficiency. However, it does lack visibility to the systems engineer. This has been overcome by extensive use of 'comment' in a high level language such that the source code is immediately preceded by the high level equivalent.

Host software used in the development process is also critical since it can introduce common errors. It must also therefore be 'visible', structured, fully tested and fully documented.

3.3. Progress in performance analysis

The performance analysis methods to be used in the future should follow the basic pattern used in the automatic landing programmes. However, the wider use of ACT for improved handling, gust and load alleviation, ride control and other similar applications will, in the author's opinion, demand more rigorous and sophisticated models of both the aircraft and the atmosphere in which more confidence can be placed. This is not to say that the data used for assessment of current design is inadequate but rather that the introduction of ACT, with relatively high gain feedback control, can vastly change the pattern of dynamic response of aircraft performance and structural loading and the models involved must be 'understandable' in greater depth.

Dynamic modelling of the aircraft for many ACT applications will need to cover the full flight envelope and include not only all the appropriate weight, CG positions etc, but also aircraft structural modes across the whole of the range of frequencies likely to react to the ACT.

In the aerodynamics and structure areas it will be necessary to establish techniques for the inclusion of such aspects as flutter and flying controls deformation effects in the model.

The ACT systems characteristics themselves will need accurate representation in performance analyses simulations, including nonlinear effects.

The work will be of a cross-discipline nature requiring close liaison between the people working in the different areas to ensure adequate understanding of the interfaces and optimisation of the overall design solutions.

A common approach to the analytical problems in the different disciplines would be an advantage. Aerodynamic modelling in terms of transfer functions, now under study, could for example bring an understanding of the problems of aerodynamics and control systems closer together. The modelling of atmospheric turbulence also needs further development to meet the validation accuracy required in the assessment of ACT systems.

The wind models used for VC10 automatic landing assessment were deduced largely from measurements taken near the ground at various test sites and in the Concorde the data base for modelling was expanded by making measurements of aircraft response and deducing the wind effects.

The VC10 and Concorde were assessed for the effect of discrete gusts and later the Concorde performance was checked in a selection of the most critical cases for the effect of certain combinations of discrete gusts of various shapes and timing. It is considered that this method of assessment might be used in conjunction with assessment of random turbulence (described by power spectral density methods) with fixed pattern wind shear to cover the average situation.

The extreme occurrences are ones which may lead to an accident, and cannot be neglected at the low probability risk levels being considered for ACT performance.

An understanding of atmospheric turbulence is important to the whole field of aeronautical design, and a great deal of new information is now being collected from in-flight measurements. However there is as yet no universal agreement about the models which are best suited to particular circumstances.

The author considers that some internationally agreed range of models should be decided as a guide for safety assessment analyses of various types of ACT system rather than leaving the choice of model to specific system designers, and to the judgement of different approval authorities. This would then concentrate the future collection of data into a specific area of interest and give the widest practical backing of experience to future operational systems.

4. NEW ACTIVE CONTROL SYSTEM APPLICATIONS

In recent years a range of new applications of ACT have been considered and this has yielded a relatively large number of experimental and prototype programmes. A few of these have been evolved to production standard.

In the military field there are a number of aircraft containing safety critical terrain following systems and a number containing fly-by-wire. The large amount of published material on the latter leaves little further to be said other than that the design and safety assessment techniques involved with fly-by-wire systems follow the same lines as previously discussed in this paper in relation to automatic landing.

There is one major difference in that the time on risk is increased to the full flight time, rather than the few minutes required for an approach and landing.

The next step after fly-by-wire, when this is adequately developed, will undoubtedly be into so-called 'fly-by-light', using fibre optic data transmission for increased protection from electromagnetic interference.

There are two safety critical developments which are known in some detail by the author which are worth a specific mention from the viewpoint of safety of design and methods of assessment. These are the secondary controls for the Airbus Industrie A310 and the overall control system of the Boeing YC-14.

4.1. Secondary controls

Wide consideration is now being given to the use of digital controls for the operation of secondary surfaces, such as flaps and slats, on both military and civil transport aircraft. Such systems will employ computer controlled electrical drive to the mechanical cross-coupling system between symmetrical pairs of surfaces and will allow a more complex automatic scheduling of operation than hitherto, and a more comprehensive failure detection and control.

In one such system at present under design for the Airbus Industrie A310 the possibility of common mode failures in software has been avoided by using dissimilar microprocessors, with different software, in a dual-dual computing architecture. Both dual system outputs are consolidated together to drive the surfaces. In the event of a failure being detected in one of the dual computer systems, its output demand will be frozen and the second system will continue to drive the output, but at reduced rate. In the event of any surface control mechanical transmission failures or actuator runaways the system can freeze the surfaces in a symmetrically balanced state.

The use of dissimilar software will considerably ease the software design control problem and the ensuing failure modes and effects analysis. This is a relatively simple system.

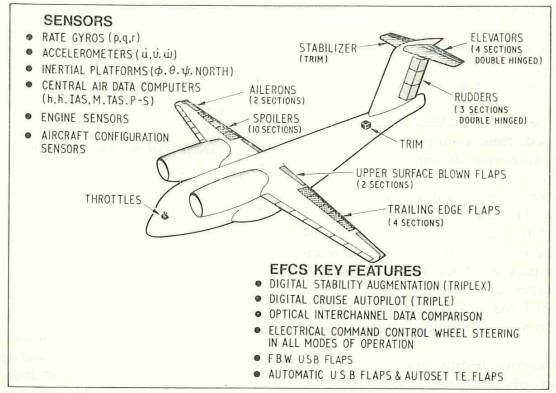

Figure 16. YC-14 AMST EFCS 1975.

4.2. Integration of primary, secondary and engine controls

At the other extreme, it is inevitable that in complex STOL and VTOL aircraft types some integration of the normal separate control and engine functions will be required to achieve a common command system. One example of this in recent years was the Boeing YC-14 which had integrated control of the normal flying controls, upper surface blown flap controls and engines.

Figure 16 shows the flying control surfaces of the aircraft which are commanded by large authority electrical signalling from the electrical flight control system (EFCS). Full time mechanical signalling is also provided for all surfaces except the upper surface blown flaps (USB) which are solely fly-by-wire.

The USB flaps are a key feature of the high lift system, acting as a thrust vector control and rotating the thrust to nearly vertical on STOL landing approach. They are also used for drag control, in conjunction with throttle thrust control, when automatically holding the selected approach speed. Increased lift is generated on the outboard wings by large circular-arc, double slotted, trailing edge flaps. Variable camber leading edge high lift flaps are also used, and their useful operating angle of attack is further increased by boundary layer control blowing. Optimum performance is produced by separately scheduling the leading edge, outboard flaps and USB flaps.

The triplex digital electronic flight control system (EFCS), (Fig. 17), is designed to be fail-operative for a single failure and fail-passive for a second failure, and provides:

(*i*) Excellent flying qualities, using conventional piloting techniques—no special 'STOL-mode' training is required.

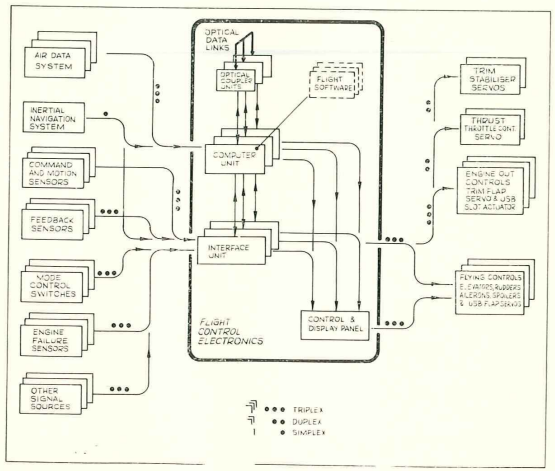

Figure 17. YC-14 EFCS block diagram.

(*ii*) Automatic compensation for powered lift STOL characteristics, control on the reverse slope of the power curve and axis cross-coupling.
(*iii*) Good engine-out performance by automatic balancing of the USBs on the remaining engine efflux.

The triplex redundant control enabled large control surface authorities to be used, thus avoiding any tendency towards marginal control authority during windshear or turbulent STOL operations.

A feature of the system is the elimination of many cross channel signal paths. The customary connection of all sensors into all redundant paths for voting purposes was avoided and instead the data was transmitted between channels on time multiplexed fibre-optic highways. In addition to the advantages of electrical isolation, the use of fibre optics eliminated the risk of external sources of electromagnetic interference corrupting these critical cross channel signals.

Design visibility and rigorous testing and analysis are essential for the avoidance of common design errors in both the hardware and software and these were key features in the system development. For the software, normal modular programming techniques were extended to provide a highly disciplined modular structure with minimum inter-modular communication and interaction; this approach, together with comprehensive testing and analysis supported by computer aids for assembled programme analysis guarded against programming and assembly errors as well as design errors.

Hardware testing was similarly meticulous, involving some 10 000 words of acceptance test software specifically designed to exercise the hardware over the extremes of operation.

More than 600 hours flight experience of the triplex digital control system was gained on two experimental prototype YC-14 aircraft.

An interesting comparison has been made between the YC-14 and Concorde automatic flight control systems. The Concorde system is perhaps the most complex fail-operative analogue AFCS ever built, and the YC-14 incorporated the first comparable fail-operative digital system.

A comparative analysis indicated that the computing and interface hardware of the YC-14 would need to be increased by only about 40% to perform the more extensive task presented by Concorde requirements. However the total weight of the system would still be only half that of the earlier Concorde analogue design.

5. CONCLUSION AND COMMENTS

Many conclusions can be drawn from the foregoing but in the author's opinion there are three most worthy of mention.

(*i*) Automatic landing systems were the first safety critical flight control systems to be approved for regular operations and the techniques evolved form a very adequate basis for the design and assessment of most safety critical systems.
(*ii*) The new capability brought about by microelectronic digital devices will encourage the use of new computing redundancy architectures such as triplex-monitoring and dissimilar redundancy in the form of dissimilar microprocessor hardware and software.
(*iii*) The assessment principles used in the past should be applicable to most new safety critical systems but more rigorous aerodynamic, structural and atmospheric turbulence models may be required for the assessment of the more advanced ACT systems such as ride control and load alleviation.

ACKNOWLEDGMENTS

The author acknowledges the assistance given in the preparation of this paper and the assembly of background data by Mr. D. I. Jackson and also by Messrs. J. Aplin, G. Belcher, R. George, K. Rosenberg and J. Taylor, all of Marconi Avionics Ltd. His thanks are also due to Aerospatiale, Airbus Industrie, The Boeing Company, British Aerospace, British Airways, the Civil Aviation Authority (Airworthiness Division), the Service Technique Aéronautique, and Smiths Industries.

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