Fault-Tolerant System Reliability Modeling/Analysis

C. Johan Masreliez* and Bo E. Bjurman†
Boeing Commercial Airplane Company, Seattle, Wash.

A formulation of a reliability analysis approach for a redundant channel system was a primary task of the NASA-sponsored Airborne Advanced Reconfigurable Computer System (ARCS) study. Major design objectives for the ARCS were to achieve two-fail-operational capability for the majority of failures in a triplex system by degradation from triplex to duplex to simplex operation, and to provide transient fault survivability. This paper presents the Markov model technique used for the reliability assessment of the application model flight control system, including sensors, computer system, servos, and hydraulic power. The application model included a fly-by-wire control wheel steering mode and an all-weather autoland mode. Primary reliability parameters evaluated were functional survivability and functional readiness as a function of time since last verification of a fault-free system.

Introduction

FAULT-TOLERANT systems are operating in many aircraft today. Military aircraft rely on redundant flight control avionics to achieve specified handling qualities. The latest generation of wide-body commercial transports are all offered with fail-operational category III landing systems. Future applications of active controls technology in commercial transports to achieve lighter airframe structures, and to expand the flight envelope into supersonic or STOL regimes, will depend on the use of fault-tolerant flight control systems. Digital computer technology will play a significant role in the cost-effective achievement of fault tolerance and in the maintenance of fault-tolerant systems.

The Boeing Company was involved recently in a study of an Airborne Advanced Reconfigurable Computer System (ARCS) sponsored by NASA Langley Research Center (NASA Contract NAS1-13654). Although the ARCS study concentrated on flight control system applications, the overall objective was enhancement of fault-tolerant computer system technology for commercial transport avionics in general. A major design objective for the ARCS was to achieve two-fail-operational capability for the majority of faults in a triplex system by degradation from triplex to duplex to simplex operation. Recovery from transient fault conditions typical for the aircraft environment was another design objective.

A primary study task was the formulation of a reliability analysis approach for a modular redundant system and the application of this analysis method to a baseline ARCS concept and to configuration alternatives defined during the study. The objective of the analysis was to bring out the economic aspects of redundant system reliability, i.e., function availability, as well as the safety aspect, or function survivability. Since no existing reliability analysis tools could be applied to the type of systems involved, a new computerized analysis tool was developed on the ARCS program.

This paper presents the ARCS reliability analysis technique and the new computerized tool. As a first step the evaluated reliability parameters are defined, followed by a postulated set of operational criteria to which a fault-tolerant flight control system might be designed. An identification of the type of system architecture that is evolving with the digital technology completes the list of criteria specifying the analysis capabilities required of a new tool. The balance of the paper then describes the features of the new reliability analysis tool and, by showing some analysis results, highlights its versatility as an aid in the synthesis and analysis of redundant configurations.

Presented as Paper 76-1931 at the AIAA Guidance and Control Conference, San Diego, Calif., Aug. 16-18, 1976; submitted Sept. 8, 1976.

Index categories: Guidance and Control; Analytical and Numerical Methods; Reliability, Maintainability, and Logistics Support.

*Research, Flight Controls Technology; presently Principle Development Engineer, Honeywell Marine Systems Division, Seattle, Wash.

†Research, Flight Controls Technology.

Survivability and Availability

In the convention used in this paper, a triple modular redundant (TMR) system by design requires two out of three correct signals at each voting node to continue functioning. Since a TMR system is based on majority voting, it fails (passively) upon occurrence of two like failures, as shown in Fig. 1. A triplex-duplex-simplex (TDS) system, in contrast, is designed for degradation to simplex upon occurrence of two like failures, as shown in Fig. 2.

The probability of system failure for the TMR system is the combined probability of first and second like failures. TDS system failure probability is the combined probabilities of

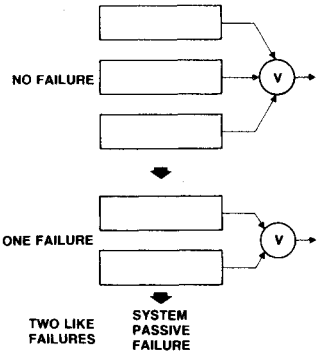

Fig. 1 Triple modular redundance (TMR).

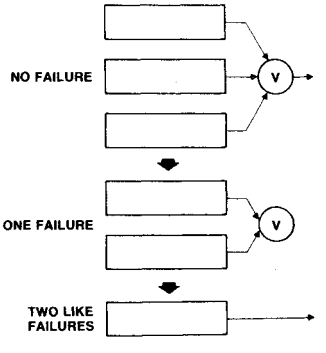

Fig. 2 Triplex-duplex-simplex (TDS).

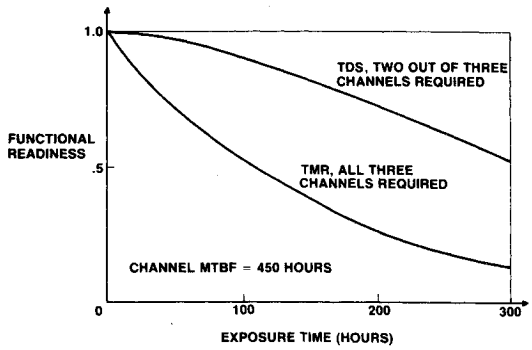

Fig. 3 TMR and TDS functional readiness.

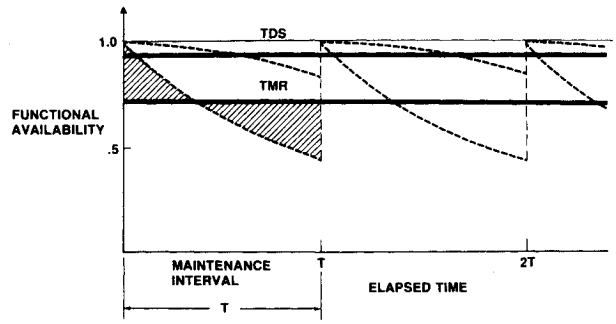

Fig. 4 Functional availability (principle).

first, second like, and third like failures. System survivability is one minus the probability of system failure. Note that a TDS system is not truly two-fail-operational, since some two-failure combinations may exist which cause system failure.

We define functional readiness as the probability of retaining a function at a certain time assuming the system is initially fully operational. Since the TDS system exhibits enhanced fault tolerance compared to the TMR system a higher functional readiness for the TDS system can be expected.

Consider, for instance, an autoland function, which is required to be fault-tolerant (fail-operational) at the point of decision (alert height), 100-50 ft above the runway, before the pilot can commit to a category III landing. The functional readiness of a TMR autoland system and of a TDS autoland system of equivalent complexity (equal component and channel failure probabilities) are plotted in Fig. 3 as a function of exposure time. The considerably higher functional readiness made possible by the TDS concept is one justification for our interest in the TDS approach.

Functional availability is the probability that a function is available for use at any point in time, given a certain maintenance schedule. Maintenance for this purpose implies that the system is checked, and repaired if necessary, within given time intervals. In this context, a fixed time interval is necessary only to provide a basis for averaging probabilities. Availability, with our definition, is the average functional readiness over the operating span considered, in this case the maintenance interval, as shown in Fig. 4. Our computerized reliability analytis tool must be capable of computing the survivability and functional readiness of the type of redundant systems that have evolved primarily for flight control applications.

Operational Scenario

The flight control system was selected as the design application for the ARCS to obtain a well-defined set of requirements from one technology in which the need for fault-tolerant processing was clearly established. This does not, however, restrict the use of the developed method or the general validity of the ARCS program results from other avionics applications, since the most severe system design requirements of a jet transport do exist in the flight control area.

Two types of flight control functions serve to illustrate considerations involved in the reliability analysis of fault tolerant flight control systems: a full-time flight crucial augmentation function, and a category III autoland function. These considerations result in design criteria addressing the reliability and the redundancy configuration required to perform the associated functions.

Let us first examine a control wheel steering (CWS) function providing control and stability augmentation for a controls configured vehicle. We postulate that the function is flight-crucial, i.e., the aircraft can not be flown without this function. The following assumptions, addressing the reliability assessment of this type of function, were deemed pertinent:

1) The aircraft will not be dispatched on a revenue flight unless the CWS function is completely fault-free.

2) The exposure time on which to base prediction of average system failure probability, or average hazard risk, will be the average flight duration of jet transports today, or approximately one hour. With this exposure time loss of system functions must be extremely improbable [FAR 25.1309(b)].

3) The exposure time on which to base the prediction of specific system failure probability for a long flight will be ten hours. A reasonable design goal for the *specific* hazard risk for a long flight was set at the level of the *average* hazard risk of present day operation.

4) The system must provide failure status indication in time to allow for a minimum 30 min diversion.

"Extremely improbable" is a term that has not been defined in any published regulatory document. In recent aircraft certification programs, however, a number equal to or less than 1×10^{-9} has been imposed upon manufacturers by the FAA to represent the probability of an event designated as "extremely improbable." The average hazard risk for a one-hour flight therefore should not exceed 1×10^{-9}.

The present-day hazard risk, as represented by the fatal accident data from all causes for 1968-73 world-wide jet transport operations, is 5 per 10^6 flights. A 20% contribution by the flight control system was considered to be a reasonable allocation of risk budget for this subsystem of the aircraft. The specific risk for a ten-hour flight therefore should not exceed 1×10^{-6}.

The category III autoland function represents a case in which, in contrast to the CCV-CWS function, the exposure time in the critical final landing phase is very short (45 sec). This is beneficial to the hazard risk, but in terms of function availability, the elapsed time still is on the order of tens of hours, since system repair will not be conducted in flight, and many times may not be available, even between flights.

The FAR regulations are ambiguous relative to the autoland function survivability requirement. FAR 25.1309(b), dealing with aircraft system design, states in effect that failure or *failure combinations preventing continued safe operation* must be shown to be extremely improbable. Advisory Circular 120-28A, dealing with category III autoland only, states in effect that system *function failure* must be extremely improbable. Since unsafe operation after function failure primarily would occur only in combination with category III conditions, the latter statements results in a survivability requirement two orders of magnitude more demanding, as shown in the following, and, from a safety point of view, in an overdesigned system.

Autoland function failure (complete loss of function) is assumed to be hazardous under two conditions: 1) it occurs during the 45 sec following alert height passage in category III conditions; or 2) it occurs during the same 45-sec phase in category II or better visibility, and the pilot fails to recover control. The combined probabilities of autoland failure and category III weather, or autoland failure and pilot failure to

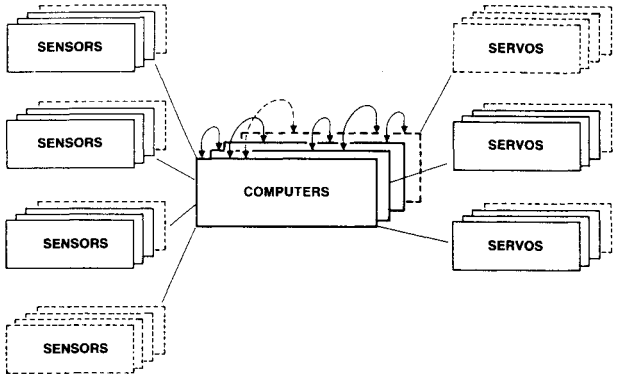

Fig. 5 Fault-tolerant system architecture.

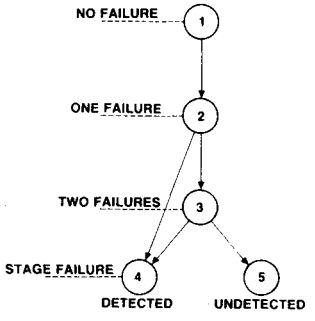

Fig. 6 Stage Markov model.

recover, respectively, must be shown to be equal to or less than 1×10^{-9}.

Allowing for a 1% probability of category III weather conditions, which appears to be a conservative assumption for a world-wide average based on existing weather statistics, the result is a required autoland function failure probability equal to or less than 1×10^{-7} per landing. Probability of pilot failure to take over in category II or better visibility of one in 1000 has been accepted and used for certification purposes; this results in a required autoland failure probability of 1×10^{-6} or less, in other words a less demanding design condition than category III. In contrast, traditional interpretation of AC 120-28A results in a required autoland function failure probability of 1×10^{-9} or less.

We propose adopting the more relaxed interpretation of the FAR's as a more reasonable approach which, without compromising safety, will allow more cost-effective designs. The stricter requirement fails to recognize the functional availability aspect; an autoland system of present technology complying with AC 120-28A can be predicted to be unavailable for category III operation during a substantial portion of the aircraft operating time.

Requiring an autoland function failure probability of 1×10^{-7} or less results in accidents due to autoland function failure being extremely improbable; less than one accident in one billion automatic landings. To put this in perspective, present-day jet transport accident data show approximately two fatal landing accidents (from all causes) per one million landings.

The general class of fault-tolerant systems which has evolved in response to the type of requirements just described, and as a result of digital computer technology advancement, can be represented by the block diagram in Fig. 5. Capability to model all system architecture characteristics of consequence to the reliability properties of the system function is a requirement on the reliability analysis tool. The tool must handle redundancy levels of at least four throughout sensors, computers, and servos. It must consider all voting nodes throughout the system and the type of redundancy management constraints imposed by the design (TMR, TDS, etc.). It must recognize dependencies within channels. Finally, and most important, it must be able to take into account parameters that depend on the ability of the system to detect, isolate, and recover from faults.

The Need for Improved Reliability Analysis Tools

The technology required to implement a system based on the TMR redundancy management concept is well established today, as exemplified by production systems like the 747 automatic landing system, which is analog with a single force voting node at the autopilot servo output. The complexity of this system to a considerable degree represents the feasibility limit of an analog system. The digital flight control computer has, however, opened the door towards enhanced fault tolerance by the use of more sophisticated redundancy management strategies. One such enhancement technique is the introduction of several internal signal voting nodes, which will increase the success probability by increasing the number of success paths through the system. Other approaches make use of the superior logical processing power of a digital computer to improve the fault tolerance. The TDS system concept is an example of this, in which the fault tolerance of a TMR system is enhanced by using the digital flight control system processor to detect, isolate, and recover from failures occurring at duplex redundancy level by providing for degradation to simplex operation. An important factor here is the capability of detecting computer failures by self-test, and the possibility of implementing sophisticated reasonableness tests for sensor and servo signals.

The increased sophistication of digital flight control systems of today have made many previously used reliability analysis tools obsolete, since several new unique requirements have emerged. One such requirement is the ability to assess system functional readiness, as well as system survivability. Of particular interest is, for example, the functional readiness of the stability augmentation function required for STOL landing and the survivability of this function during a landing, given that the function was available at the go-ahead decision point.

A second requirement is the ability to taken into account the probability of surviving a second channel failure, i.e., the probability of successful degradation from duplex to simplex operation following a second channel failure in a triplex system. This probability is denoted "the second failure coverage," where coverage is defined as the conditional probability of detection, isolation, and recovery given a failure. It will be assumed that voting between three redundant signals will guarantee a unity first failure coverage in the TDS system, i.e., that the system is strictly fail operational.

The third requirement concerns the modeling of dependencies between the various system modules which demands special consideration when these dependencies act across signal voting nodes internal to the system. Other desirable features of the reliability analysis tool are transient fault modeling and the ability to assess two different failure mode probabilities corresponding to a detected and an undetected (latent) failure.

CARSRA: A New Reliability Analysis Tool

Ten different reliability analysis programs, six from Boeing and four from NASA, were evaluated relative to the requirements outlined. Since none of these were applicable, a new program was developed. This program, which uses approximately 800 Fortran statements, is denoted CARSRA for Computer Aided Redundant System Reliability Analysis. The main features of the CARSRA analysis approach are the following.

For the purpose of the reliability analysis, the system is partitioned into stages and modules, where a module is set of elements performing a specified function and a stage is a collection of identical, redundant modules. Each stage is

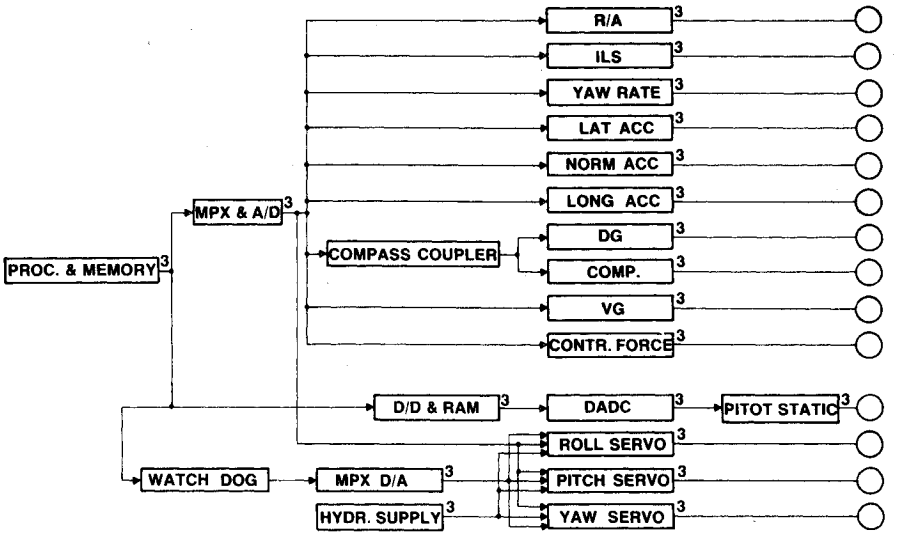

Fig. 7 Typical dependency tree.

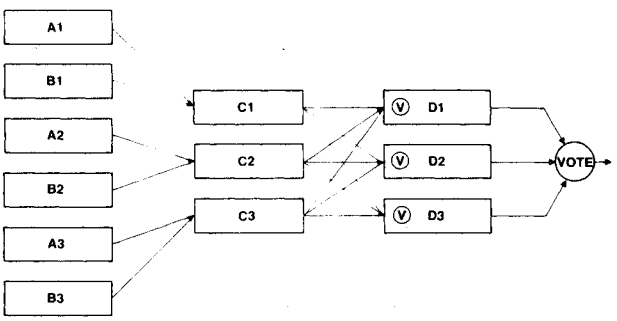

Fig. 8 Triplex system.

modeled by a Markov model, describing the different redundancy states of the stage. The two last states (assigned the highest numbers) in the Markov model are stage failure states. The next to the last state corresponds to a detected failure and the last state to an undetected failure. Up to nine states in the Markov model are permitted for each stage. An example of a stage Markov model is given in Fig. 6.

State 1 represents the no-failure state, where three modules are operational. Upon the first module failure, a transition is made from state 1 to state 2. From state 2, which corresponds to duplex operation, a transition is made either to state 3, i.e., degradation to simplex operation takes place, or to the detected failure state 4. In a system based on the TMR concept, all transitions out of state 2 go to state 4, whereas in a TDS system, the majority of all failures, i.e., all covered failures, end up in state 3. Upon a failure in state three, the system will fail, detected (state 4) or undetected (state 5).

The model of Fig. 6 is a simple representation of a triplex stage. Higher resolution may be incorporated by also introducing latent (undetected) failure states at the one- and two-failure levels. The ability of the system to reject transient faults usually will depend on the operational redundancy level in that the effect of a transient occurring in one of three redundant channels usually is less damaging than a transient occurring at duplex redundancy. This difference will be reflected in failure rates, which will depend on the operational redundancy state. The Markov model is ideally suited to model these transient effects, since the transition rates between the various Markov states may be specified arbitrarily.

Depending on the way in which the various stages have been defined, some stages generally will have the property that failure of one module in the stage will cause loss of function of a module in one or several other stages. This introduces a statistical dependency between the various stages which has to be taken into account in the analysis. A module that when failed will cause loss of function of another module in a different stage will be called a "dependency" module, and the corresponding stage a "dependency stage."

Thus, there are two types of stages: dependency and nondependency stages. Some stages may be both dependency and nondependency stages. The dependency structure of a system may be described by a dependency tree diagram, an example of which is displayed in Fig. 7 for a particular implementation flight control system. In this system, the processor/memory stage is a dependency stage, the MPX and A/D stage is both a dependency and a nondependency stage, and each sensor function is a nondependency stage.

The lines connecting the different stages indicate the dependency structure in the sense that a failure in, for example, the processor channel A, will cause loss of function of the channel A sensor and servo modules. The digits at the upper right-hand corner of each block indicate the number of redundant modules in each stage.

The signal consolidation, or voting nodes, to the right in the figure represents functions needed for system survival. A loss of a combination of these functions will cause system failure. This combination may be sepcified by a particular entry to the program.

CARSRA treats dependency between stages by an approach that may be denoted "exhaustive conditioning." The essence of this approach is to make the nondependency stages independent via conditioning upon the failure status of the dependency stages. This approach is explained most easily by presenting a simple example. Consider the triplex system outlined in Fig. 8, consisting of two sensor stages A and B, a multiplex stage C, and a computer stage D. The sensor signals are multiplexed and cross-strapped into the computers, where signal selection (voting) and failure detection is performed in software.

TMR operation is assumed, which implies that two out of three signals are required at each voting node. For simplicity, the assumption is made that the output voter has zero failure rate.

Note that the sensor stages and the multiplex stage are mutually dependent in the sense that a multiplex failure will cause a module failure both in sensor stage A and sensor stage B. The dependency is unidirectional, since a sensor failure will not prevent the multiplex module from acquiring data from a sensor in another channel. Figure 9 displays the corresponding dependency tree.

To explain the approach, the following theorem will be needed: Let E_i, $i = 1, 2, \ldots n$, be disjoint events with

$$\sum_{i=1}^{n} P(E_i) = 1$$

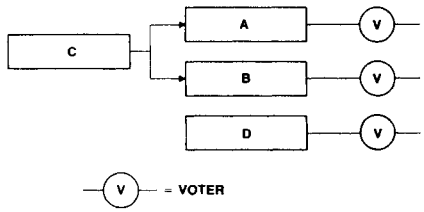

Fig. 9 Example dependency tree.

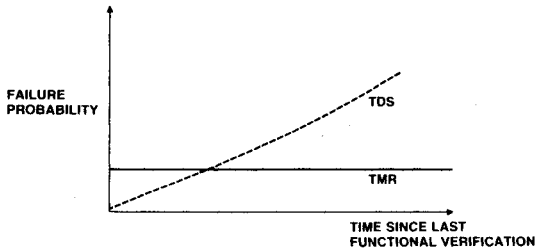

Fig. 10 TDS and TMR failure probability.

Let F be an arbitrary event. Then,

$$P(F) = \sum_{i=1}^{n} P(F|E_i) \cdot P(E_i)$$

The success probability for the system of Fig. 8 now may be found by defining the events E_i as follows: E_1 = no multiplex module failed; E_2 = one multiplex module failed; E_3 = two multiplex modules failed; E_4 = all multiplex modules failed.

The probability of system success may, according to the previous theorem, be expanded:

$$P(S) = \sum_{i=1}^{4} P(S|E_i) \cdot P(E_i)$$

The advantage of this representation is that the probabilities $P(S|E_i)$ usually are easier to find than finding $P(S)$ directly. For example,

$P(S|E_1)$

= P{(stage A survives) and (stage B survives) and (stage D survives)}

becomes, with R = module reliability and $Q = 1 - R$,

$$P(S|E_1) = (R_A^3 + 3R_A^2 Q_A) \cdot (R_B^3 + 3R_B^2 Q_B) \cdot (R_D^3 + 3R_D^2 Q_D)$$

Furthermore,

$$P(E_1) = R_C^3$$

The next term in the expression contains the factor $P(S|E_2)$, i.e., the probability of system survival, given that one multiplex module is failed. In this case, both of the remaining modules in sensor stages A and B have to survive for system survival:

$$P(S|E_2) = R_A^2 \cdot R_B^2 \cdot (R_D^3 + 3R_D^2 Q_D)$$

$$P(E_2) = 3R_C^2 Q_C$$

Finally, $P(S|E_3) = P(S|E_4) = 0$. Summarizing, the system reliability becomes

$$P(S) = (R_A^3 + 3R_A^2 Q_A)(R_B^3 + 3R_B^2 Q_B)(R_D^3 + 3R_D^2 Q_D) R_C^3 + R_A^2 \cdot R_B^2 (R_D^3 + 3R_D^2 Q_D) \cdot 3R_C^2 Q_C$$

Note that this expression differs from what is obtained if the stages are assumed to be independent:

$$P(S) = (R_A^3 + 3R_A^2 Q_A)(R_B^3 + 3R_B^2 Q_B) \cdot (R_C^3 + 3R_C^2 Q_C) \cdot (R_D^3 + 3R_D^2 Q_D)$$

As was mentioned previously, the voted (or signal consolidated) outputs from the nondependency stages constitute the functions required for systems survival. There are, however, situations in which these functions themselves may be redundant, for example the redundancy between aileron and spoiler control surfaces of certain aircrafts. CARSRA will model this situation by accepting a success event tabulation covering all nondependency stage combinations equivalent to system success.

Summarizing, the computation is performed in three different steps: Markov modeling for each stage, treating dependencies between stages via exhaustive conditioning, and specifying the functions needed for success by success configuration tabulation.

Functional Readiness Feature

Fault-tolerant systems will continue to perform their functions, even after incurring one or several module failures. Since this is an inherent system capability, it will be of interest to assess the probability of experiencing a certain redundancy degradation within a prescribed time interval and, furthermore, to be able to assess the system failure probability given that a certain degradation has taken place. This information could be used to establish functional readiness criteria, i.e., the system failure states that will not cause deferment of a particularly critical phase of a mission, for example, a landing in low-visibility weather conditions.

CARSRA accepts a selected functional readiness criterion specifying the combination of modules which could be failed, and computes the probability of having any of these modules failed as a function of time. It also computes the system failure given a functional readiness criterion as a function of time in separately specified time frames. Two different system failure modes may be specified, e.g., detected or undetected system failure.

The following relations are used. Let $PFR_i(t_1), (i = 1, 2, \ldots N)$, denote the probabilities associated with the different specified functional readiness configurations at a time t_1, and let $PFP_i(t_2)$ be the conditional probability of a certain failure mode at an exposure time t_2, given functional readiness configuration i. CARSRA then computes:

$$PFR = \text{functional readiness} = \sum_{i=1}^{N} PFR_i(t_1)$$

PFP = failure probability

$$= \left[\sum_{i=1}^{N} PFP_i(t_2) \times PFR_i(t_1) \right] \cdot PFR^{-1}$$

In addition to the previously mentioned features, two levels of computational accuracy may be specified. The resulting computational roundoff errors are indicated in the computer printout.

Comparison between the TDS and TMR Redundancy Management Concepts

As was mentioned previously, the improved fault tolerance of a TDS system relative to a TMR system is obtained by providing for further redundancy degradation from duplex to simplex upon a second module failure in a stage. The reduction in system failure probability obtained by a TDS compared to a TMR system is roughly proportional to $c_2 = (1 - c_2)$ where c_2 is the second failure coverage. If the

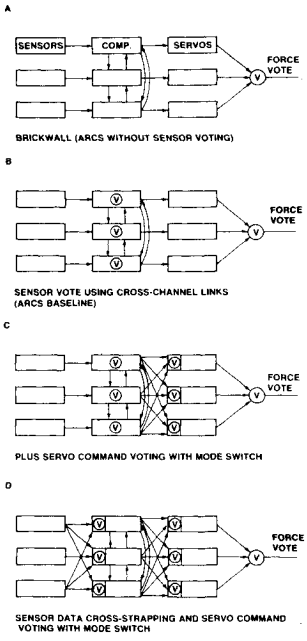

Fig. 11 Voting configurations.

second failure coverage is equal to 0.9, the reduction will be by a factor ten and, if it is equal to 0.95, by a factor twenty.

The improved fault tolerance of a TDS system, although desirable, would perhaps not by itself provide justification for implementation, since the survivability of a TMR system already may be adequate. However, this low failure probability of a TMR system may be realized only if all three channels are operational at the time the system function is needed. If this is not the case, i.e., if one channel is failed, the survivability is poor, since failure of any of the two remaining channels will result in loss of the system function.

The increased fault tolerance of a TDS system suggests the possibility of occasionally permitting a flight critical mission phase, such as a category III landing or a STOL landing, even with one channel failed in the system, since an adequately designed TDS system will survive the vast majority of all second channel failures.

Figure 3 showed the characteristic of system functional readiness as a function of time for the TMR and the TDS systems. Observe that the slope of the curve is zero at time zero for the TDS, but falls off rapidly for the TMR.

The failure probability of the TDS system, when operated with a functional readiness criterion that permits occasional category III or STOL landings with a degraded system, will depend on the relative frequency of these degraded system landings. This parameter will in turn depend on the maintenance process in that the longer maintenance is deferred, the higher the probability of actually having a degraded system during a landing, and thus the higher the probability of system failure.

For a TMR system, however, the probability of failure will be independent of the maintenance strategy, since all channels always will be operational before the functions may be used. This situation is illustrated in Fig. 10, which shows that the TDS system failure probability initially is lower than that of the TMR system, but increases with increasing time since latest complete verification of system operation. This verification could, for example, be done in flight by using the system, or by a thorough ground test following system maintenance.

The increased functional readiness, which conceivably could be realized by a TDS system, and the feasibility of implementing this philosophy in a digital flight control system, provides a significant incentive for seriously considering the TDS concept in applications in which high functional readiness is of importance.

Table 1 Failure probabilities

Configuration	Failure probability
A. Brick wall	26.0×10^{-7}
B. Sensor voting	5.4×10^{-7}
C. Sensor and output voting	5.3×10^{-7}
D. Sensor and output voting, cross-strapped sensor signals	2.3×10^{-7}

Application Example

A typical question to which CARSRA will provide the answer is: What is the most efficient way of arranging signal consolidation points in a given system?

Figure 11 shows four possible signal voting configurations in a triplex flight control system. The "brick-wall" system has one voting node at the servo outputs (force voting). The second system employs software sensor data voting, in addition to the servo voting, using digital data links for cross-channel communication. The third system introduces an additional voting node on the servo command signals between the computers and the servos. In a TDS implementation this requires additional logic in the servo interfaces for the monitoring and decision making required for duplex to simplex servo reconfiguration. The fourth configuration is obtained by adding sensor signal cross-strapping to the third configuration, so that each computer interfaces directly with all sensors. Cross-strapping requires additional interface hardware, but decreases the sensitivity to computer failures.

The resulting system failure probabilities for a command/stability augmentation function, using present-day sensors and a one-hour exposure time, are presented in Table 1.

In conclusion, this comparison of voting configurations shows that sensor voting significantly improves the survivability, but that the benefit of servo command voting is negligible. Direct cross-strapping of sensor signals into each computer results in a factor 2 improvement over cross-channel exchange of sensor data via digital data links. This benefit is obtained in spite of the added interface hardware required for the fourth configuration.

Concluding Comments

Application of digital technology to flight control systems opens the door to improved redundancy management concepts and techniques. The TDS concept described in this paper improves the fault tolerance of a triplex system by providing for degradation to simplex following second like module failure. Although a TDS system will not tolerate all possible second-failure combinations, its survivability would be sufficiently high to permit initiation of a critical mission phase, for example a category III landing, with one module failed. A substantial increase in function availability would result in comparison with the fail-operational, fail-passive TMR implementation.

A reliability analysis program, CARSRA, was developed to handle the analysis of fault-tolerant modular redundant systems. Significant features of this program are its ability to account for coverage parameters, to model transient fault effects, and to handle statistical dependencies between modules, thereby providing analysis capabilities not existing in other available reliability analysis tools. The versatility of CARSRA was demonstrated in analysis of flight control systems and rapid transit vehicle control systems. Work presently is underway to include the capability to model the effect of equipment repair.