AD-A066 697 CHARLES STARK DRAPER LAB INC CAMBRIDGE MA F/G 17/2
A REVIEW OF THE 3M DATA BASE FOR FAULT-TOLERANT SYSTEM INCENTIV--ETC(U)
JAN 79 A L HOPKINS N00014-76-C-0502
UNCLASSIFIED R-1244 NL

| OF |
AD
A066897

END
DATE
FILMED
5-79
DDC

LEVEL II

R-1244

# A REVIEW OF THE SM DATA BASE FOR FAULT-TOLERANT SYSTEM INCENTIVES

by

Albert L. Hopkins, Jr.

January 1979



The Charles Stark Draper Laboratory, Inc.
Cambridge, Massachusetts 02139

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

## REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

| | |
|---|---|
| 1. REPORT NUMBER | 2. GOVT ACCESSION NO. |
| 3. RECIPIENT'S CATALOG NUMBER | |
| 4. TITLE *(and Subtitle)* (6) A REVIEW OF THE 3M DATA BASE FOR FAULT-TOLERANT SYSTEM INCENTIVES. | 5. TYPE OF REPORT & PERIOD COVERED (9) Final Report July 1977 - Dec 1978, |
| 6. PERFORMING ORG. REPORT NUMBER (14) R-1244 | |
| 7. AUTHOR(s) (10) Albert L. Hopkins, Jr | 8. CONTRACT OR GRANT NUMBER(s) (15) N00014-76-C-0502 |
| 9. PERFORMING ORGANIZATION NAME AND ADDRESS The Charles Stark Draper Laboratory, Inc. Cambridge, Massachusetts 02139 | 10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS (11) Jan 79 |
| 11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research 800 North Quincy Street Arlington, VA 22217 (12) 37p. | 12. REPORT DATE December 1978 |
| | 13. NUMBER OF PAGES 34 |
| 14. MONITORING AGENCY NAME & ADDRESS *(if different from Controlling Office)* | 15. SECURITY CLASS. *(of this report)* UNCLASSIFIED |
| | 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE |

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Fault-Tolerant Systems, Avionics, Maintenance

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

This report covers the second phase of the subject contract, whose first phase was the study of a hierarchical form of a fault-tolerant data communication network. The second phase task involved an examination of the Navy's "3-M" (maintenance material management) data base to see if any evidence was readily available or easily extractable to affirm or refute the hypotheses underlying fault-tolerant system design.

(over)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

S/N 0102-014-6601

408 386

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

A large volume of aggregate data was examined for three aircraft types, the E-2C, the P-3C, and the S-3A. Several broad conclusions were reached with respect to the fault-tolerant system design hypotheses. Some observations and inferences are presented to conclude this report.

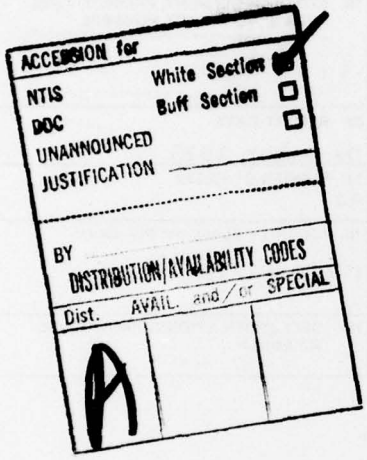

R-1244

# A REVIEW OF THE 3M DATA BASE FOR FAULT-TOLERANT SYSTEM INCENTIVES

by

Albert L. Hopkins, Jr.

January 1979

Albert L. Hopkins, Jr.

Norman E. Sears

The Charles Stark Draper Laboratory, Inc.
Cambridge, Massachusetts 02139

## ACKNOWLEDGEMENT

This report was prepared with the support of the Office of Naval Research under Contract N00014-76-C-0502.

The author is grateful to Eldon Hall for his advice and assistance in the data review and in the preparation of this report.

The assistance of Dr. Marvin Denicoff and Dr. Leonard Haynes of the ONR is also gratefully acknowledged. Special thanks are due to Messrs. Larry Minnaugh and James Kapp of the Maintenance Support Office Department of the U.S. Navy for their cooperation.

A 015 920

# 1. Introduction

The U.S. Navy's "3-M"(maintenance material management) reporting system contains the records of millions of field maintenance actions for ships and aircraft. The Draper Laboratory proposed a search of this data with the objective of seeking confirmation of some of the hypotheses upon which the design of fault-tolerant systems has been based. A considerable volume of aggregate data was searched, and a number of broad conclusions were reached. This report chronicles the effort's rationale and history, and concludes with several observations and inferences.

The remainder of this section is a narrative history of the program. Section 2 reviews the concept of fault-tolerant system design and identifies its underlying hypotheses. Section 3 describes the aggregate data obtained from the Fleet Material Support Office in Mechanicsburg, Pennsylvania. Section 4 describes the project's findings.

## 1.1 Background

Since about 1965, the Draper Laboratory has been studying and developing fault-tolerant digital control system technology. In the period between September, 1974, and August, 1975, the Office of Naval Research sponsored a demonstration and evaluation of a fault-tolerant input/output network at Draper via Contract N00014-75-C-0050. That work is documented in Draper report R-918, "Demonstration and Evaluation of a Fault-Tolerant Input/Output Network" by John F. McKenna, Jr., dated September, 1975. Both airborne and shipborne applications are appropriate for such networks, which can provide both damage tolerance and fault tolerance. They are intended to be managed by fault-tolerant computers, (although in non-critical applications they could be managed by less-systematic redundant organizations of computers).

In March, 1976, under ONR contract N00014-76-C-0502, the Draper Laboratory studied and demonstrated a hierarchical version of the fault-tolerant network, as documented in an MIT Thesis by Lt. James R. Seeley, USN, "Development of a Hierarchical, Dynamically Reconfigurable Input/Output Network," Draper Report T-644, June, 1977.

In the meantime, Draper had submitted an unsolicited proposal to ONR on the application of the fault-tolerant network concept to shipboard systems. In the course of discussions pursuant to this proposal, the existence of the Navy's 3M data base came to light. Draper subsequently submitted a second unsolicited proposal, this one for the present interrogation of the data base. The latter proposal was accepted by ONR, and the second contract was amended to include this work. It was originally intended to confine this work to shipboard systems; but, for various reasons, only aircraft systems were considered.

## 1.2 Summary of the Proposal

The proposed effort was an investigation of the hazard environment for some proposed fault-tolerant system architectures. The plan was to study available data on operational failures in order to arrive at an initial assessment of the potential value of fault tolerance in an operational combat system.

It has been established that systems can and should be designed in such a way as to enhance the testability and maintainability of their constituent elements and of the systems themselves. Moreover, some life-cycle cost reductions and enhancements of effectiveness can be accomplished by measures short of fault tolerance. Nevertheless, once the down-payment has been made to enhance testability and maintainability, the additional costs to produce a fault-tolerant system may be low in many cases. The benefits of fault tolerance are discussed in Section 2.

The general picture of the distributed, fault-tolerant system is one where a healing process takes place automatically after a failure, using surviving equipment in a different algorithmic configuration. To some, this would sound futuristic and hopelessly complex, but in fact the system structure is extremely tractable. We have shown, for example, that a reconfigurable communications network requires less than 2K x 16 bits of memory for its entire management program and data complement. We have moreover demonstrated a fault-tolerant multiprocessor that can manage a distributed hierarchical system, local processing complexes, and algorithmic software that covers the contingencies of individual

failed sensors and effectors. Therefore, the technology exists for such a system to be carried to its next phase of implementation.

Nevertheless, there is some degree of risk in attempting a radical change in system architecture. The more that any such risk can be offset, the more likely it will be that fault-tolerant system architecture will be able to be adopted into near-future system developments. One of the principal elements of risk concerns the validity of the assumed operational environment.

The distributed fault-tolerant system approach makes certain assumptions about the hazards that are encountered in operational life. In particular, the following assumptions are made:

1. Digital hardware has no greater, and perhaps a lower failure rate than analog electronic hardware or electromechanical hardware.
2. Random failures are uncorrelated.
3. Most induced failures from damage and electrical accidents have a limited extent, and are propagated only by system architectural shortcomings that can be overcome.
4. Some physical areas of a combat system are less vulnerable to damage and/or induced failures than others.

The data bases maintained by the Navy at Mechanicsburg, Pennsylvania, appeared to provide an opportunity to acquire at least a first order appraisal of whether

a. the above assumptions are reasonable, and
b. a distributed, fault-tolerant, integrated system architecture would have been of substantial benefit.

The proposed project is outlined below.

Proposed Work Statement

1. Two to four man-trips to Mechanicsburg, Pennsylvania, to access data bases containing data on operational failures and damage. Two trips would be required for familiarization with data retrieval methods available. Later, one or two trips might be needed to critique preliminary searches and specify subsequent searches.

2. Examination and reduction of the data thus obtained to find approximate orders of magnitude for correlations, spatial interactions, and any other hazard parameters that are deemed to be applicable.

3. Estimate the importance of masking, distribution, reconfigurability, hierarchy, integration, and damage tolerance in future systems assuming these same hazards will be present.

4. Prepare a final report to ONR.

5. Two man-trips to Annapolis to present results in the combat system survivability session at the Combat Systems Symposium.

## 1.3 Mechanicsburg Visit

In actuality, only one man-trip was made to Mechanicsburg by Draper personnel, and the fifth item became "not applicable," since the work began after the time of the Symposium.

The trip to Mechanicsburg was made on 9 November 1977, by the author, accompanied by two ONR staff members*. A meeting was held with Mr. Larry Minnaugh of the Maintenance Support Office Department, also attended by Mr. James Kapp of that organization. The principal outcomes of that meeting were the following.

1. Decision to access aircraft data exclusively.
2. Decision to use standard report formats.
3. Briefing on the 3-M system given.

* M. Denicoff and L. Haynes

The following subsections enlarge on these topics.

### 1.3.1 Decision to Access Aircraft Data

The reporting, filing, and distribution of aircraft maintenance data has been automated to a greater extent than for shipboard data. Therefore the cost of accessing shipboard data would have been substantially greater than that for aircraft data. Although the original intention was to investigate the shipboard environment, there were two factors which influenced the general readiness to alter the objective to aircraft. One was the fact that the Navy at that time was considering a long-lead time development of advanced V/STOL aircraft in which the possibility of an integrated fault-tolerant flight-control and avionics system was an important issue. The other was the absence of any prior field history study of the military aircraft environment for fault-tolerant systems. Prior studies were largely confined to spacecraft and civil transport aircraft.

### 1.3.2 Decision To Use Standard Report Formats

The vast volume of aircraft maintenance data, corresponding to some 115,000 WRA (Weapon Replaceable Assembly) removals per month, is aggregated in numerous ways in the many existing standard report formats. Although the possibility exists of creating new formats for search and aggregation, the special programming required contributes a significant cost and delay impact. It was therefore agreed that standard formats would be used, unless for some reason they were wholly inadequate, in which case the subject could be reconsidered.

### 1.3.3 Briefing On the 3-M System

Messrs. Minnaugh and Kapp presented a briefing on the reporting and dissemination methods used by the Maintenance Support Office Department. Some 41 basic formats are presently used, with an additional 27 variants. The reporting frequencies range from monthly to annual, with some reports being purely on-demand. Several of the reports are classified Confidential.

The format information is summarized in a document entitled "Catalog of Aviation 3-M Information Reports." As furnished to Draper, it bears the additional designation FMSOINST 4790.1B 1 Feb. 1977, and is issued by: Department of the Navy, Navy Fleet Material Support Office, Maintenance Support Office Department, Mechanicsburg, PA, 17055. This document also contains information on how to prepare requests for information reports, and shows the various source document formats.

Further trips were not needed, thanks to the catalog and to information received via several telecons between Mr. Kapp and the author. This willing collaboration on the part of MSOD was extremely helpful.

## 1.4 Summary of the Investigation

After a study of the catalog, Draper requested and received several report documents concerning the E2-C, the P3-C, and the S3-A aircraft. The rationale was based on the mission similarity between these aircraft and the proposed V/STOL-A aircraft, and the relatively large dependence of these aircraft on digital devices.

The documents consistently show the high volume of maintenance activity that has had much general discussion in recent years. Electronic assemblies appear to have relatively high failure rates, but this may well reflect their complexity as much as their fragility. Complex mechanical assemblies, such as engines, likewise have high failure rates. There is a clear challenge to the designer to reduce these rates by one means or another.

The mission abort rate is fairly low, except for engine failures. Redundancy is clearly necessary for any flight critical function, though it contributes to maintenance and supply problems. This is the arena in which fault-tolerant systems can hope to compete well.

Beyond the patent finding of high failure rates and few aborts, useful inferences were hard to find. The more detailed observations are discussed in Section 4.

## 2. The Fault-Tolerant System Concept

The purpose of this section is to portray the type of system architecture which motivated this data base study. The subject is perhaps best introduced by the appended extract from an article in Astronautics and Aeronautics, September, 1978, by Deyst and Hopkins, entitled "Highly Survivable Integrated Avionics." (See Appendix.)

In military aircraft there is a great deal to be gained by a high degree of automation. This is first because of the need to overcome performance barriers set by the reaction speed and accuracy of the human crew. Such performance barriers may affect flight control modes, aircraft architecture, countermeasures, and system reconfiguration and recovery. Second, automation allows crew training to be devoted more toward the managerial aspects of the mission where the human capability excels. In particular, the automation of malfunction responses can greatly reduce the training burden if it preserves the nominal transfer function of the aircraft as perceived by the crew.

In order for such automation to be possible, it is probably necessary to give full authority over the aircraft to the electronic system, i.e. the aircraft becomes a CCV (Control-Configured Vehicle) ipso facto. In general, the more life-critical functions that are excluded from automation, the more difficult it will be to realize economic benefits. As an example, consider the implementation of automatic landing on a V/STOL aircraft for small-ship operation, which has the potential for saving some minutes of hover time capability, and therefore some tens of thousands of pounds of gross weight. This enormous saving accrues from automation of several life-critical functions. If we consider those mission electronics with no life-critical functions, no such startling an example is evident, although automation can improve performance, mission success, and maintainability to some degree. A life-critical mission function such as terrain following, however, benefits greatly in performance and/or safety by automation, assuming that such automation is possible with the requisite reliability and economy.

The existence of a fully flight-critical fault-tolerant electronics system would offer the design latitude to automate any function whatever, whether critical or not. Another potential source of system economy and reliability arises if the system is wholly integrated. (cf. Appendix). For one thing, this permits the pooling of all sensor information, which in turn makes it possible to allocate sensors in such a way as to obtain maximum fault coverage for a given number of sensors, or alternatively to use a minimum complement of sensors to obtain a given degree of fault coverage. The first of these two alternatives is of greater interest, because of the leverage afforded by the integrated system to yield large benefits from few additional sensors. The significance of this leverage is not so much in terms of flight criticality as it is in dispatch reliability in a reduced material condition.

The fault-tolerant integrated system, in a nutshell, offers a framework in which to achieve numerous alternative contingent operation modes with the fewest possible components, and with the greatest possible commonality of information processing resources. As such, it helps to address problems of supply. It is also capable of performing fault detection and diagnosis in real time without using dedicated built-in test equipment, by means of routine comparisons of functionally redundant data. It further offers the possibility of automatic record keeping for system configuration and diagnostic data for the entire aircraft.

Thus, at least on the surface, the wholly-integrated fault- tolerant system approach is believed to possess great potential for future weapons systems. A number of questions are raised, however, which would be appropriate to address at this time, and which helped to motivate this project.

1. A fault-tolerant system is possible only if its components have adequately low failure rates. Can digital airborne hardware exhibit failure rates that are adequately low? Can other hardware?
2. Correlated failures severely undermine redundancy. Are field failures correlated or not?

3. Are induced failures from damage and electrical accidents prevalent such as to defeat redundancy? Is damage tolerance necessary or not?

4. Is damage vulnerability constant over the physical extent of the aircraft, or are there relatively safe areas?

5. Are there problems with respect to diagnosis, supply, and maintenance complexity, that would defeat a fault-tolerant system? Alternatively, would the system itself tend to have fewer such problems than existing systems have?

The remainder of this report addresses the data base access and the conclusions based on essentially these questions.

## 3. Accessing the Data Base

The five topic questions at the end of the preceding section are reiterated here in brief form:

1. Failure rates.
2. Correlation.
3. Induced failures.
4. Safe places.
5. Diagnosis, supply, and maintenance.

Failure rate information is directly available through the reports for every identifiable subsystem, module, or component for which there exists a unique work unit code in any given aircraft type. Thus, for example, the AN/AQA7 Sonar Computer Recorder Group in the P-3C is found to have exhibited 13.0 mean flight hours between failures in 25,593 flight hours from January through March, 1978 (MSO4790.A 2142-04).

Correlation information does not appear to be available from the data base. One might conceivably think of making inferences from the number of aborted missions, but this would not be particularly well advised in the absence of further data.

Data that identifies induced failures is not carried in the data base, nor is it supplied from the field. Likewise, there is nothing that would help to distinguish safe locations from vulnerable ones.

There is a good deal of data that bears on diagnosis and supply. Maintenance complexity is inferred from the source document formats and the total volume of maintenance actions.

The following subsections briefly describe the aggregate report formats that were selected for review.

### 3.1 Reliability and Maintainability Trend Analysis Summary MSO 4790.A 2142-04

These reports contain subsystem-level data on the number of maintenance actions, the mean flight hours between maintenance actions (MFHBMA) the mean flight hours between failures, (MFHBF), and various other aggregate data including unscheduled maintenance man-hours per flight hour. (For the P3-C's T56 engines, the latter

value exceeds unity, while all other subsystems have values less than one). These reports also show the rank of the subsystem's failure rate per flight hour. It is interesting that nearly all of the MFHBF's lie very near a straight line when plotted against failure rank. For the P3-C data reviewed, the MFHBF may thus be approximated by 25 + 2.3R, where R is the failure rank. The first and second ranked subsystems (Sonar Computer/Recorder and Engine) deviate from this approximation with values of MFHBF of 13.0 and 16.3. Thirty-two subsystems had MFHBF's of under one hundred hours on the P-3C in this period.

The failure rank seems to show electronics in a bad light, and particularly the electronics that is part or all-digital. Equipment with failure ranks of 1,3,4,5, and 6 appeared to be in this category (AN/AQA7, AN/ASN84, AN/APS115, AN/AYA8, and AN/ASA70 respectively).

The mean flight hours between maintenance actions (MFHBMA) is in nearly all cases substantially less than the MFHBF, indicating that a false removal problem may exist.

### 3.2 Fleet Failure Summary (MS04790.A 2107-01)

These reports break out monthly failures for a possible trend analysis, and also show the incidence of failures below subsystem level. This data considerably mitigates the apparent indictment of digital electronics by the reports described in subsection 3.1. Digital modules do not here appear to have an exceptionally high failure incidence compared to other electronic modules and various connectors and mechanical devices. This data also highlights the inherent complexity of some of these high failure rank items by their long lists of subsidiary modules and components.

### 3.3 Aircraft Degradation Ranking Summary (MSOD 4790.A 3364-01)

These reports rank subsystems according to which ones are most responsible for non-operational readiness (NOR) and reduced material condition (RMC) of the given aircraft type. These reports provide a surprising contrast to the two already discussed. For

the P-3C aircraft, the highest NOR/RMC ranked subsystem is the T56 engine, which was ranked second in failure rate. The next five subsystems are electronic. Surprisingly, however, their failure ranks are 23, 32, 38, 133, and 40, respectively. Their problems are primarily with supply, rather than with maintenance. The number-one failure-ranked sonar computer/recorder ranks only 18th in NOR/RMC. Of the 1,959 NOR/RMC hours attributable to this unit, 87 per cent were due to supply.

### 3.4 3M Aviation Type Equipment History Inquiry (MSO 4790. A 2097-01) Parts 1 and 2 Only

These reports are voluminous, as they show for each work unit code, for each aircraft type, the number of instances of each malfunction type code. The format breaks these down into when-discovered and action-taken categories.

Again using the AN/AQA7 as an example, for the period from January through March, 1978, there were 3,441 "discoveries" of a need for maintenance. Of these, approximately 25% were classified as improper alignment, 12% improper display, 12% no defect, 11% connection defect, 10% no output, 5% incorrect output, 4% internal failure, 2% broken, and the remainder (about 17%) distributed among some 75 other malfunction categories. About 21% were discovered in flight by the crew with no mission abort, 20% between flights by the ground crew, 22% during preflight or postflight inspection, 26% during an in-shop repair and/or disassembly for maintenance, and the remainder (about 11%) in 14 other when-discovered categories.

In the same period, there were 6,207 actions taken. Of these, about 24% are listed as "no defect," but this includes 17% of cannibalization items.

The malfunction codes are not mutually exclusive, so that it is likely that like faults are reported under various codes. Neither are the codes unambiguous, so that it is also likely that numerous different faults are reported under the same code, e.g. incorrect output, or internal failure. Nevertheless one gets the impression from this and other reports that quite a wide variety of malfunctions occurred, not attributable to a single weak point,

or to a recurring causal event.

One noteworthy aspect of these reports is that they indicate a low number of mission aborts, with some exceptions. Engines account for the most aborts per unit failure. In the P-3C, the inertial navigation set accounted for 17 aborts in 762 aircrew-detected failures in preflight and flight. All components other than engines and inertial navigation caused only 28 aborts. Engines alone caused 44.

### 3.5 3-M Aviation No Defect Item Analysis Summary (MSO 4790. A 2551-01)

These reports are designed to summarize the impact of false removals of equipment, i.e. when a defect is suspected, but not found. For the period October 1977 through March, 1978, for example, 12% of all no-defect items for P-3 family aircraft were attributable to the AN/AQA7. But these items represented only 10% of all items concerning this particular subsystem. This was the first-ranked item for the P-3 with respect to no-defect items. Second was the inertial navigation unit, AN/ASN42, with 4% of the P-3 no defect items, which represented 15% of this subsystem's items. Overall, about 9% of all maintenance items were no-defect items. This number apparently does not include cannibalization.

There is a distinct tendency for electronics to exhibit more no-defect item percentages than for mechanical parts, as one might expect: Engines, for example, had a 1.5% no-defect rate.

### 3.6 3-M Aviation Cannibalization Analysis Summary (MSOD 4790. A 3855-01)

These reports show high cannibalization items, and manpower data connected with cannibalization actions. They do not indicate a particularly high level of cannibalization, which seems somewhat surprising in the light of the significant problems of NOR and RMC due to supplies. The cannibalization rate for any one assembly code rarely exceeds one action per hundred flight hours.

## 3.7 3-M Aviation Reparable Component Installation and Removal Report (MSO 4790. A 2303-01)

These reports summarize the service times for components aboard aircraft between installation and removal. For the P-3C aircraft overall in CY1977, there were 24,290 installations and 16,234 removals. Of these removals, 5,250 were for cannibalization. This would appear to contradict the conclusion from the preceding subsection, and be more in line with what one would expect. The data for individual work unit codes is rather difficult to interpret.

## 4. Observations and Conclusions

Subsystem failure rates in the subject aircraft seem to be quite high, particularly in the more complex subsystems. If one were to try to achieve fault tolerance by redundancy at the subsystem level, it is difficult to see how critical functions could be safely handled by systems where so many elements have MFHBF's below one hundred hours. The failure rate data strongly suggests that fault tolerance can be achieved only if it goes below the subsystem level for its replaceable elements.

Although many of the high failure-rank items are at least partly digital, the evidence would not support an indictment of digital hardware as unreliable. Indeed, experience in other contexts suggests that digital hardware is more reliable than its analog equivalents.

Correlation of failures can not be judged from the data reviewed. It is, perhaps, marginally relevant to note that the mission abort rate is kept quite low by the redundancy and on-board management philosophies that are embodied in these aircraft. Data contained in A2142-04, A2097-01 and A3364-01 for the period January 1978 through March 1978 are of interest in this respect. In addition, the P-3C Orion Weapon System Description LR27955 dated March, 1977 helps to provide an understanding of the system and the maintenance philosophy employed. This maintenance concept depends upon on-board capability to detect and isolate failures to the individual module level. Where feasibile, fault isolation is computer-initiated by using a diagnostic program. Where not feasible, equipment has been provided with Built-In Test Equipment features, or readily accessible test points to accommodate standard test equipment affording fault isolation to the individual replaceable module. On-board test equipment, backup systems and spare modules are provided to support a comprehensive in-flight maintenance which was designed to minimize the number of mission failures or aborts. The success of this maintenance philosophy is reflected in the data from the 3-M reports.

From the data available in A2097-01 on the P-3C, it is clear the maintenance concept is reasonably successful in reducing

the number of aborts. Table I summarizes this data for the period January 1978 through March 1978 on the engines and several avionics or mission related systems. Table II provides a similar summary for the E-2C. It is Draper's understanding that this aircraft has in-flight maintenance capabilities similar to the P-3C.

Diagnosis does not appear to present a serious problem, according to the no-defect data in A2551-01. The percentage of false removals would appear to be below 10%, which is a far better situation than prevails in commercial airlines practice, for example. This is a somewhat surprising finding.

Supply is a widespread limitation on operational readiness, as evidenced in A3364-01, although one should be careful to distinguish NOR from RMC here. A fault-tolerant system would tend to address the NOR problem, but would make a reduced material condition both more prevalent and more acceptable.

It may be concluded from these observations that the airborne failures of components presently pose a serious problem in terms of cost and performance of naval aircraft, but not of safety. The growth of automation can be expected to be severely limited unless one or more current problem areas are redressed.

Integrated fault-tolerant systems have considerable potential in several respects. They permit maximum pooling of data, which tends to optimize the tradeoffs of performance, economy, and safety. This would appear to be highly important if it prevented component proliferation while yielding combat superiority. The additional benefits of diagnosability and uniformity of information handling components would favorably impact maintenance as well as other elements of life cycle cost.

This review has partially succeeded in its purpose by roughly identifying a prevailing environment in one class of aircraft. Although it might be useful to review similar data for other classes, e.g. fighters, it should be recognized that the 3-M data base is designed for purposes other than the present one. Questions remain concerning correlated malfunctions, damage exposure, and other induced malfunctions, which should be dealt

TABLE I

P-3C FAILURE DATA FROM A2097-01 P. 550 to 1264

| WUC | ABORTS | FAILURES[1] | FAILURE RANK | EQUIP. NOM. |
|---|---|---|---|---|
| 2230 | 44 | 123 | 2 | Engine |
| 4211 | 28 | 73 | 44 | Primary AC Power |
| 5646 | 0 | 28 | 147 | Air Data Comp. |
| 5738[2] | 1 | 304 | 34 | Auto Flt. Co. |
| 612M[2] | 0 | 355 | 17 | HF Radio Set |
| 632K[2] | 2 | 347 | 22 | UHF Radio Set |
| 6422 | 5 | 489 | 11 | Intercomm. |
| 6912 | 1 | 115 | 23 | Data Terminal |
| 7116 | 2 | 57 | 74 | Direction Finder |
| 7143 | 4 | 184 | 35 | Tacan Nav Set |
| 726A | 2 | 566 | 4 | Radar Set |
| 732B[2] | 3 | 319 | 6 | Tact. Data Display |
| 734F | 17 | 762 | 3 | Inertial Nav Set |
| 7366 | 1 | 262 | 5 | AYA8 Data Anal. |
| 7367 | 1 | 212 | 13 | ASQ114 Dig. Comp. (Central) |
| 7378[2] | 3 | 785 | 1 | AQA7 Sonar Comp. |

[1]Failures discovered by flight crew during preflight checkout or inflight.

[2]System is dual to provide backup.

TABLE II

E-2C FAILURE DATA FROM A2097-01 P. 98 TO 369

| WUC ASS. NO. | ABORTS | TOTAL[1] FAILURES | FAILURE RANK | EQUIPMENT NOMENCLATURE |
|---|---|---|---|---|
| 2230 | 5 | 13 | 7 | Engine |
| 5143 | 3 | 4 | 135 | Pressure Indicator |
| 562D | 1 | 16 | 64 | Air Data Comp. |
| 5648 | 0 | 3 | 153 | Air Data Comp. |
| 5686 | 1 | 37 | 37 | Stall Warning Sys. |
| 5732 | 1 | 47 | 27 | Auto Flt. Cont. |
| 6117 | 0 | 23 | 67 | Radio Set |
| 611C | 0 | 13 | 94 | Comm. Set |
| 611G | 0 | 35 | 53 | Radio Set |
| 6315 | 1 | 102 | 8 | Radio Set |
| 7244 | 1 | 21 | 49 | Alt. |
| 726D | 3 | 161 | 2 | Radar |
| 726E | 1 | 20 | 40 | Radar |
| 726G | 0 | 65 | 21 | Radar |
| 726H | 0 | 31 | 34 | Radar |
| 726J | 0 | 206 | 1 | Control Indicator |
| 7288 | 0 | 38 | 28 | Comp. Indicator |
| 728D | 1 | 5 | 125 | Comp. |
| 728E | 2 | 121 | 4 | Comp. |
| 734H | 1 | 137 | 6 | Inertial Nav. |
| 761E | 0 | 64 | 3 | Countermeas. Rec. |

[1] Detected by flight crew during preflight checkout and flight.

with by first establishing a baseline fault-tolerant system design. Then, perhaps, a renewed search of the data base will show something more substantial, although the magnitude of the effort would probably need to be a good deal greater than the present one.

Finally, the 3-M data base involves the generation, transmission, and processing of a great deal of data. A future integrated fault-tolerant digital system would be able to expedite much of this by automatic collection and summary of on-board diagnostic information. The possibility also exists of furnishing each WRA with a unique machine-readable identifier code, so that systems could largely track their own configurations and report on their own operational readiness for a variety of missions.

# APPENDIX

Reprinted from *Astronautics & Aeronautics*, September 1978.


# Highly Survivable Integrated Avionics

**By JOHN J. DEYST, JR., and ALBERT L. HOPKINS, JR.**
The Charles Stark Draper Laboratory, Inc.

A number of studies and flight tests have indicated the potential benefits of a control-configured design approach for many types of aircraft. The Air Force F-16 fighter is the first operational example of a vehicle which takes advantage of control configuration, and its success testifies to the practicality of the concept. Projected V/STOL aircraft will likely require stability augmentation to achieve their full potential. Many prototype transport-aircraft designs incorporate some form of configuration control, and there seems little doubt that future transports can benefit from it.

Control studies and experiments have shown the adequacy of current procedures for designing effective feedback control laws for control-configured vehicles (CCV); and contemporary sensors and actuators can fulfill most performance CCV requirements. Digital fly-by-wire systems have been flown and are finding their way into operational aircraft. Thus, much of the technology for control-configured design can be directly applied to advanced aircraft of several types.[1]

For control-configured vehicles, especially transport aircraft, the flight-control system becomes "flight critical"—requires a level of reliability commensurate with the hardness of the basic airframe structure itself. To attain high reliability, conventional design replicates large independent subsystems (redundancy). But this technique cannot produce a practical design to meet a full-time "flight critical" requirement. Typically, it yields an extremely complex and unwieldy configuration, entailing excessive numbers and types of components, with associated problems of availability and maintenance.

Thus reliability, maintainability, and redundancy management of avionics systems become the primary issues that must be simultaneously addressed before the full potential of advanced control technology can be applied to modern aircraft.

Stephen Osder of Sperry's Flight Systems Div. has characterized a "complexity divergence" in redundant avionics.[2] He finds that the redundancy "overhead" in today's triplex, quadruplex, or dual-dual redundant flight control systems runs almost an order of magnitude greater than the operation (or function) itself. Current autoland systems, entailing large overheads, are certificated for full-authority control for flight segments of *less than one minute* in commercial transport aircraft. For full-time, full-authority operation of double fault-correcting autopilots (fail-op, fail-op) with today's technology, Osder projects an overhead of up to twenty times the basic functional hardware.

The digital computer has the potential to reduce these overheads. Some experimental digital avionics have demonstrated it. Nevertheless, most existing approaches to redundant digital avionics have fallen substantially short of their potential, largely for one or more of these reasons:

—Independent subsystem organization with limited cross-utilization of information, resulting in inefficient use of resources.

—Incomplete fault coverage, owing to single-point failure modes in redundant computer complexes, redundant buses, and redundant sensors and effectors.

—Large centralization of computation, with attendant high cost of multi-tasking, fast-response software, a uniform redundancy level for all functions, and large losses of system capability or redundancy due to single-point failures.

The tradition of separate subsystems, each covered in some fashion for the eventuality of malfunction, leads to undue reliance on human responses, overlaps in authority, and inordinate complexity in design, integration, validation, and certification. On the other hand, providing an extremely dependable information system permits the design of a highly survivable integrated avionics system with substantially less overhead than needed by existing avionic systems whose survivability is limited to one or two failures.

In what follows we explain our concept of the type of system that we feel will be necessary to attain the performance advantages that lie tantalizingly before us. We cannot offer total solutions and clearly there are unresolved questions that can only be answered by additional development and testing. Some progress has been made, however, and we hope to give the reader a reasonable understanding of what has been accomplished and what remains to be done.

Some recent approaches to critical systems have been based on total integration, in the sense that information generated or developed anywhere within the system can easily be made available anywhere else in it.[3-5] This characteristic directly increases performance, economy, and survival, all at the same time. There are several traditional reasons, meanwhile, why such systems have been avoided in the past:

—Desire for functional separation, so that failure in a non-critical function does not cause outage of a critical function, and so that maintenance is not impeded by ambiguities as to the place of a failure.

—Need to eliminate single-point failures.

—Desire to tolerate correlated, or similar, failures.

These traditional guidelines are important in formulating any integrated system concept.

First, the necessary functional separation can be realized with a combination of hardware and software partitions within the integrated system. A

single system can securely support numerous concurrent functions related to one another through well-defined, moderate-bandwidth interfaces. For example, flight control and inertial navigation can both be supported by the integrated system and both functions can use information from common inertial sensors. Inadvertent attempts to violate interfaces by a software element must not adversely affect "innocent" functions, whereas the "guilty" element's function may be lost. Also, the system's hardware and software partitions must make it possible to detect and diagnose hardware and/or software malfunctions to the level of the replaceable module, or the partitioned program. Meanwhile surviving resources are allocated to the most critical function on a priority basis.

Second, single-point failure modes must be wholly eliminated, which first requires that some means exist for detecting any and all failures, and moreover requires that the source of an error be rendered incapable of affecting the system's tolerance to subsequent failures.

Third, the system must be purged of correlated failure modes. The alternative of using diverse hardware and software modules is rejected here, much as it is in the case of aircraft engines. The penalty for providing dissimilar engines is greater than the penalty for purging correlated failure mechanisms, and similarly with computer hardware and software.

The design of integrated systems for high survivability is a relatively new field of scientific and engineering endeavor. Current emphasis is on fault tolerance, for which a considerable foundation has been laid in the computer field.[6] Concepts of fault tolerance at the system level are based on the ability of digital computers to execute decision algorithms of the complexity required to maintain system integrity despite the occurrence of non-massive failures or damage. The generation and use of the necessary algorithms is feasible today for systems whose overall failure rates are in the area of $10^{-6}$ to $10^{-7}$ failures per hour, and the state of the art is moving on toward $10^{-8}$ to $10^{-9}$ failures per hour, at which point it will be feasible to think of putting systems of this character in passenger aircraft.

Today we lack not so much the technology of designing such systems as the way of "validating" them to the requisite confidence levels.[7]

Before proceeding further, we should clarify terms and provide a few definitions. In design practice the term "integrated system" is often used ambiguously to denote virtually any degree of combination of functions that were separate at one time or another. To be more exacting, we find it useful to characterize *a very high level of integration of functions*, in which it is taken for granted that all avionic functions fall within the system scope, and where flexible communication paths exist throughout; this we call *wholly integrated*. It is also important to distinguish systems composed of diverse, dissimilar elements (e.g., different computers) from those using identical elements in modular fashion. The latter we call *homogeneous systems*.

The following paragraphs discuss a number of the attributes required in *homogeneous, wholly-integrated systems* to achieve the performance, economy, and safety needed for flight-critical avionics.

The underlying system concept: the system as a whole needs to acquire information via sensors, to interface it, communicate it, to process it via computers, and to deliver it to actuators (including displays). Having *alternative* means for fulfilling these needs, in the event of any reasonably likely sequence of malfunctions, makes survival possible. Here we consider the attributes of the various system constituents from the viewpoint of survival, in turn treating sensors, digital information processing, and actuators.

Sensors are specific to their own locations and functions. The alternative to a failed sensor is a second sensor, which may or may not be either identical to or co-located with the first. In the wholly integrated system, sensor redundancy is chosen on the basis of projected failure rates and modes, on local damage probability, and on the ability to extract redundant information by the analysis of dissimilar or remote sensors. In any event, reliable digital information processing must be available to diagnose a sensor failure. Choosing sensors to serve multi-functional roles and making their outputs generally available throughout the system increases cost-effectiveness.

For example, the primary sensors required for flight control are inertial—i.e., gyros and accelerometers. Inertial instruments also serve fire control, navigation, autoland, and pilot-display functions. Each of these functions, however, has a different set of requirements in terms of performance and reliability. The performance requirements on flight-control gyros and accelerometers are quite loose compared to those for the other functions. Typically, gyro drift rates of the order of 1000°/hr and accelerometer errors of 0.05 $g$ are acceptable for aircraft flight control and air-to-air fire control. By way of contrast, the navigation and certain air-to-ground missile fire-control functions imposed on the aircraft avionics system require gyro drift below 0.01°/hr and accelerometer biases within $10^{-4}$ $g$. Conversely, the failure rate for flight control may be as stringent as $10^{-9}$ per hour, whereas a $10^{-1}$ per

hour failure rate for inertial navigation is reasonable. Pilot display and autoland requirements lie between these two extremes, requiring gyro drift rates of about 1°/hr and failure rate to $10^{-4}$/hr.

To fulfill these requirements, the wholly-integrated system employs multipurpose inertial sensors. Skewed strapped-down gyros and accelerometers clustered together on a rigid structure provide the inertial-navigation sensor function. These instruments must be quite accurate to provide navigation-grade angular velocity and acceleration, and they have correspondingly high cost. The skewed configuration allows a minimum number of instruments to attain a specified level of reliability. Clustering is necessary to assure the precise alignment between the sensors, required for inertial navigation. Typically, four gyros and four accelerometers suffice to provide the required navigation reliability.

The same set of sensors will also serve the pilot display, autoland, and flight-control functions. However, the group of four gyros does not provide enough redundancy to meet flight control reliability; and, because of the clustered configuration, the group is quite vulnerable to local damage. Hence it takes additional instruments, of lesser performance, to meet the more stringent reliability and damage tolerance requirements for the other functions. These could consist, for example, of four lower-grade gyros in a skewed strapdown configuration. These gyros would be physically separated from each other and from the navigation instruments to give damage tolerance. The performance and cost of the additional instruments would be considerably less than for the navigation sensors and they would serve the flight-control, pilot-display, and autoland functions.

The set of eight gyros, all skewed relative to each other, constitutes a sensor array that can yield the high level of total reliability and damage tolerance necessary for the pilot-display, autoland, and flight-control functions. An appropriate redundancy-management system can sustain as many as five random gyro failures without loss of the critical angular-rate information necessary for safe flight. Furthermore, because of the dispersion of four gyros, damage to the entire navigation cluster can be sustained and an additional gyro failure can occur without detriment to flight control. Thus, the eight-gyro set yields a very high level of survivability even in the event of severe local damage. By integrating sensor functions this complement of eight gyros can replace about 20 instruments that would be employed in a conventional system to attain equivalent overall reliability.

Effective failure detection and identification keys reliability and survivability in the sensor system. The most powerful method for achieving the required level of performance appears to be through application of recent results in *analytic redundancy*.[8-10] This method employs the dynamic and kinematic relationships between the quantities (measured by various instruments) to compare them for redundancy management.

For example, the method provides a systematic approach to the design of algorithms to detect and identify failures in dispersed gyros. The problems of relative motion between sensors, due to aircraft structural flexibility, are accounted for, and false alarms are suppressed to appropriate (specification) levels. Similarly, gyro data can be compared with

F-1 DISTRIBUTED INFORMATION-PROCESSING HIERARCHY

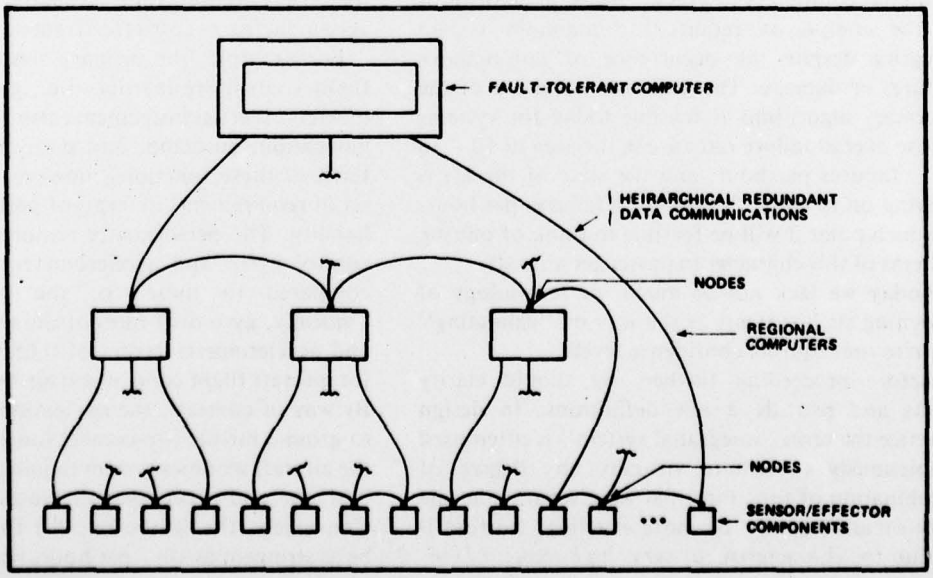

air-data measurements, through the governing kinematic relationships, to help identify gyro failures. In this fashion failures of five gyros can be sustained and the surviving triad identified. Additional relationships between inertial instruments, air-data instruments, and radio navigation aids can be applied in systematic fashion to make maximum use of all available sensor redundancy.

The sensor configuration must also be chosen with appropriately sized line replaceable units (LRUs). Single-point failures in large units cause excessively large losses in performance. An inertial navigator is an example of an excessively large LRU, because a single-point failure within it can mean the loss of all information from six precision inertial instruments. By the same token, very small failure units overburden the redundancy-management system. Appropriately sized units appear to be individual inertial sensors, navigation radio receivers, air-data pressure transducers, and the like. Furthermore, replication of identical sensors to serve several roles, such as identical gyros to serve navigation, flight control and fire control functions, meets the principle of homogeneity discussed previously.

The substantial reduction of sensor components through analytic redundancy has a hidden cost. The failures in the components involved must be independent, and the system must be designed to guarantee this. Part of the independence will reflect structural, physical dispersion.

Another involves the immunity to correlated events such as lightning strikes.

Still another is the independence of power sources.

Finally, we require the unconditional survival of data paths.

These topics—important design constraints on any system—are particularly severe in this case, which depends critically on the availability of a large subset of all the sensors.

Power-source independence and data-path survival are natural by-products of the proposed homogeneous system. Power-source independence accrues from dispersed remote-controlled breakers. Effectively, the system's redundancy is wholly dependent upon its ability to communicate among every surviving component of the system. Contemporary approaches, by contrast, do not provide this ability with the degree of confidence that will support flight-crucial functions.

Attaining the richness of communication necessary to manage sensor redundancy, as well as the redundancies of all other avionic elements, requires a high degree of interconnection of elements. A number of basic approaches need to be considered in choosing the interconnection technique. Dedicated paths that run from a central computer to each sensor and effector site (sometimes referred to as a "star" connection) are not immune to limited physical damage, and they contribute to problems of weight and complex connections. They are therefore increasingly being abandoned in favor of some form of multiplexing. One-way buses, known as "broadcast buses," offer a partial solution to these problems, but they require a separate bus for each sensor, and do not appear to address the problems posed by limited damage events. Redundant two-way buses with remote couplers can be made damage-tolerant, although they are still vulnerable to nodes that generate uncontrolled bus transmissions, or 'babble,' on all redundant bus copies.

Data communications in the wholly integrated system must be completely dependable in the sense of preserving partial functioning despite likely failures and damage. At least one approach exists to this kind of data network.[11] Now in the prototype stage of development at the Draper Laboratory, it is based on the ability of the interface elements to route data around the site of a failed component or a damaged area. The resulting network is immune to the problem of "babbling" components. This makes it highly likely that all surviving components will continue to be available to the system.

This approach distributes computational functions to numerous dispersed computational locations. Certain of these sites perform small tasks which are not unto themselves particularly critical. Logically, these functions are organized into a hierarchical structure, depicted in F-1.

Functions localized to only the lowest level of the hierarchy tend to be the least critical ones. For example, a small computation site may be dedicated to each rate gyro; this site does some dedicated computation to enhance the performance of that instrument and handles the digital transmission of gyro data to higher level functions that use the data. Due to redundancy within the overall system the single site is not critical. Loss of the gyro can be compensated for by the higher-level functions.

Functions situated at higher levels of the hierarchy tend to be more critical. For example, flight control, a centralized supervisory function, gathers information from numerous lower-level functions and processes it into useful forms. Its outputs are sent back to the locales dedicated to an individual actuator or groups of them. While the individual elements of the flight-control system are not critical, its overall function is. High redundancy must guarantee its operation.

The physical dispersion of sensors and actuators, and the reliability and survivability of the various functions, as well as the hierarchical organization of the information processing, all give form to the

hardware mechanization of the proposed computational approach. Each computational site of the system can be either simplex, with no internal redundancy, or internally redundant to increase the reliability of that site. The degree of internal redundancy will depend on the reliability requirements for the site. Similarly, the computational power of a site can be matched to the computational requirements. These computation requirements may vary from being very slow, using a single microprocessor, to very fast, requiring advanced minicomputer performance.

The centralized function of certain tasks, as well as the hierarchical nature of the information processing, suggests highly centralized design, However, what is required is a physical structure which can support a hierarchical functional organization but which is also dispersed and capable of restructuring and reassigning functions to different computational sites to compensate for failures. The process of restructuring and reassigning functions can, in fact, be looked upon as the highest level function of the hierarchy. This function, as well as most of the other high-level functions, must be transferable from one computation site to another.

High level or critical functions are performed by highly reliable fault-tolerant multiprocessors.[11,12] Because of the need for physical dispersion to achieve damage tolerance, there may be two or more fault-tolerant multiprocessors remote from each other. These multiprocessors, reflecting state-of-the-art, high-performance microprocessor technology, will have substantial computational power—able to perform flight control, navigation, autoland, and thrust control plus system configuration and restructuring functions.

A single one of these multiprocessors serves as the lead computational site of the hierarchy performing all system-wide configuration control and management functions, and possibly the most critical vehicle functions, such as flight control. Such a site can be called the *lead node*. The other fault-tolerant multiprocessors perform the other high-level functions, such as navigation, and occupy the functional rungs directly below the lead node. While physically similar or identical to the lead node, they serve a secondary role. Due to the large amount of internal redundancy and redundancy management, these computational sites will have an extremely low failure rate. The only significant failure mode is, in fact, physical damage.

In the event of physical damage, a secondary node assumes the functions of the failed lead node. The secondary node is in constant communication with the lead node. Essential information on the state of the system is routinely passed over to the secondary node, at sufficiently high frequency, so that the secondary node can take over essential lead-node functions in the event of damage. The information passed back and forth is highly encoded and the processors constantly check the codes. In the event of damage to the lead node, the damage is detected, to a high probability, by the secondary node by checking for errors in the highly encoded messages. The secondary node then takes over the essential functions of the lead node. Under worst-case assumptions this may involve curtailment of less critical functions so as to make computational resources available for the more critical tasks.

Lower-level positions in the hierarchy are generally serviced by dedicated microprocessors. Most of these functions are associated with equipment that has a failure rate considerably greater than a simplex microprocessor's. Loss of a dedicated local processor is equivalent to loss of the associated subsystem, component, or sensor. However, the failure rate of the combined dedicated simplex processor and subsystem or component is not significantly greater than that of the subsystem or component alone.

In some cases a low-level processing function may require high reliability. An example might be a redundant actuator for which the processor provides internal redundancy management. This function must be performed at a reliability that considerably exceeds that of the redundant elements within the actuator itself. Then the dedicated processors will be made redundant to provide the necessary level of fault tolerance.

Computational sites are joined by a network of links[13]—a link being a fully duplex communication path between any two computational sites, or nodes. The preferred technology will probably be fiber optics. Each node is interconnected to at least three other nodes. Each node also contains switching circuitry so that the links can be connected. Thus, if a link is viewed as an I/O bus segment, a node can, by making the appropriate internal switch-closures, extend a bus through itself or cause it to "Y." The lead node can build an I/O bus which reaches all other nodes by issuing commands that cause other nodes to set up a branching bus structure. Not all links are used in this process; some remain idle. In the event of physical damage or node failure, the lead node identifies the failure and bypasses the failed or damaged units by activating idle links and nodes. If the lead node is destroyed, another node takes over its function, by isolating the damaged node and regrowing the branching bus around itself as the new origin. F-2 illustrates a network with a bus grown to communicate with all elements.

To the extent to which additional idle links and

F-2 NETWORK WITH GROWN BUS

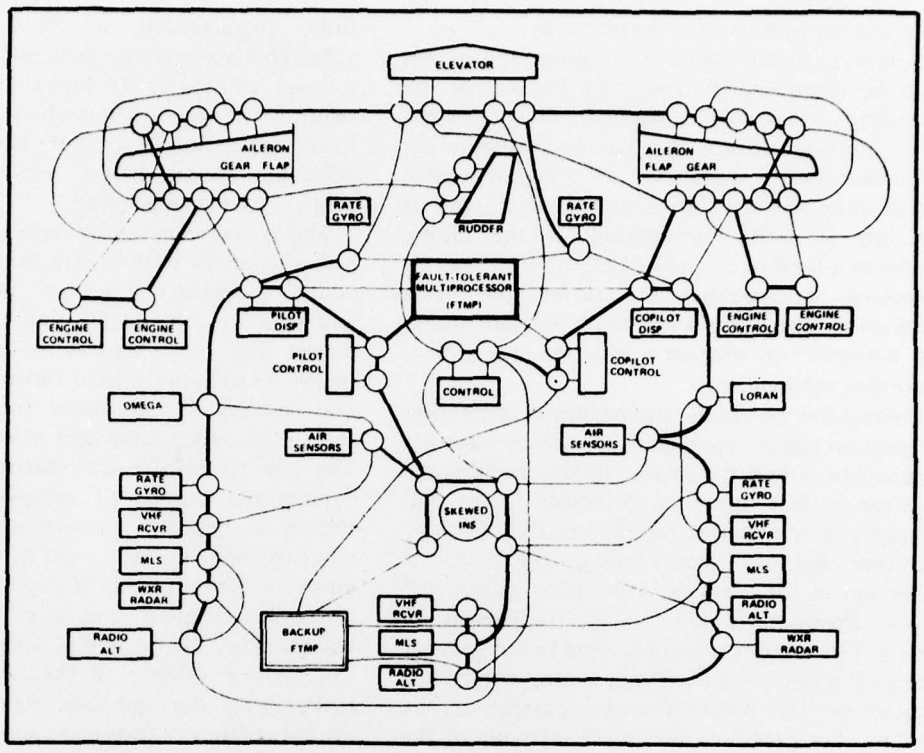


interconnection patterns permit, it is possible to provide greater and greater freedom on assigning functions to particular hardware as one moves up in the hierarchy. Thus, the lead-node role can be filled by any of the two or more fault-tolerant multiprocessors. Certain other secondary supervisory functions could also be provided by any of the fault-tolerant multiprocessors. At lower levels an alternate functional site of equal capabilities may not be available, but it may be possible for a higher-level site of greater capabilities to absorb the function. A reduced-performance mode, where crude control by a relatively busy supervisor replaces finer control by a dedicated subordinate when the subordinate fails, can thus be accommodated.

Thus far in our discussion we have examined the sensor and computational elements of a highly integrated avionics system. The third element consists of the components and subsystems that carry out commanded physical tasks—namely, actuators and displays. In terms of the management of resources in the highly integrated system, actuators differ from sensors and computational elements in that a failed actuator usually cannot be ignored, whereas a failed sensor or processor can. Hard-over control surface failures prove catastrophic in many flight situations and appropriate measures must be taken virtually to eliminate this possibility.

Segmentation of control surfaces offers an effective means to achieve control redundancy. The surfaces are partitioned so that control authority is more or less equally distributed over all the segments. Each segment has its own, independent actuation system, and these systems are all independently controllable. Failure of an actuator produces loss of the associated control-surface segment only. Segmentation of surfaces reduces the effect of the loss of a single one, and appropriate reconfiguration of control-surface commands can compensate for the loss.

Currently available redundant hydraulic actuators provide a high level of fault tolerance and serve as flight-critical elements in modern aircraft. Typically, these actuators attain fault tolerance by providing a high level of redundancy of critical elements within the actuators themselves. In the context of segmented control surfaces, however, it may be desirable to utilize simplex actuators with redundancy managed by input/output monitoring of each actuator. Comparing input command signals with feedback signals from transducers indicating actuator output can do this. By accounting for acceptable time lags in the response, it is possible to monitor the actuators and determine, with high confidence, when a failure occurs. In the highly integrated system the local dedicated processor assigned to the actuator will perform this monitoring task and indicate the presence of a

failure to superior levels in the hierarchy.

Following identification of a failure, the system must be reconfigured to account for loss of the offending control-surface segment. This is a high-order function; in fact, it requires reconfiguration of control laws to account for a net loss of control effectiveness and possible asymmetry as a result of the loss as well. Furthermore, if the failure produces a hard-over surface deflection, it may be necessary to neutralize the surface by some mechanism that releases hydraulic pressure within the actuator or otherwise destroys the force-producing mechanism.

Control-law reconfiguration involves increasing the gain on control commands to surface segments adjacent to a failed segment. In this fashion the adjacent surfaces produce sufficient forces and moments to compensate for the lost one. Clearly, this can only be effective up to the limits of deflection of the surfaces. Loss of a segment, of course, lowers maximum control effectiveness. In the event of successive failures, even to the point of losing all segments of a surface, control commands to other surfaces will be altered to compensate for the loss. For example, loss of all segments of the starboard elevator will be compensated by increasing the gain on control commands to the port elevator and employment of aileron control to compensate for roll moments due to the asymmetric elevator deflections. In a similar fashion, loss of an aileron can be compensated for by differential elevator control and the like.

We have thus far discussed the three major elements of the wholly integrated avionics system—sensors, computers, and actuator/controls. Carefully synthesizing these into a total system must be done in a way that satisfies all requirements and functions yet retains sufficient flexibility for growth and change without major redesign.

In system synthesis, the computation task in toto spreads over a rather large number of small computation sites, and the tasks performed at each remote site are relatively small and well defined. Thus, the hardware configuration immediately imposes a modularization upon much of the software design.

The fault-tolerant multiprocessors, however, will serve many tasks, and will not possess this same natural modularization. Yet the system should exhibit the advantages of modularity to accommodate inevitable changes over its life cycle without the need for total revalidation or total recertification. The state of the art in computer security measures is approaching the maturity needed to provide the desired "modularity" (independence) among co-resident software modules. The two principal vulnerable points are lapses in the security hardware, and operating systems never fully characterized in behavior. The first vulnerability essentially does not exist in the fault-tolerant computer; its probability is sufficiently remote. The second vulnerability can be overcome by a combination of structure and simplicity in the operating system design, today achievable with proper system management.

The system concept we have described is based primarily on its reliability attributes, but it has some additional features of interest. The focus on sensor and effector components rather than subsystems means that functions can be added or deleted largely as software entities rather than as "boxes" that embody the equivalent functions with their own additional sensor and effector components. The line replaceable unit becomes smaller, i.e., components instead of subsystems. The digital system can be highly modular with a few different types of modules, each used in moderate to large numbers. Certain kinds of maintenance, especially on pooled elements, can be postponed to a convenient time. The system is necessarily configured to diagnose failures to the line-replaceable-unit (LRU) level, thus speeding maintenance. Finally, the redundancy of different parts of the system can be graded to suit different needs, rather than having every function executed at the highest redundancy level.

The system we have described shows great promise, but much research, development, and testing remains before its potential benefits will be available in operational aircraft. Progress has been made. For about a year, the Draper Laboratory has had in operation a small-scale simulation of a highly integrated flight-control system for a conventional transport aircraft. The system has these principal elements:

—Hybrid simulation facility with KC-135 flight simulation and graphics generator.

—Boeing 707 cockpit mockup from an early Curtiss Dehmel simulator-trainer.

—Six-node communications network interfacing with the flight simulation.

—Breadboard emulation of a fault-tolerant multiprocessor with redundancy management software for itself and the communications network. Contains a simplified control wheel steering fly-by-wire flight-control algorithm (digital autopilot).

—Redundancy-management software for a dual set of flight-control sensors. This software is executed in the fault-tolerant computer.

F-3 gives an overall diagram of the simulation. To date, the simulated information-processing elements have interfaced with the aircraft simulation through a single node, pending a revision to the facility, at which time the interface will properly consist of a multiplicity of nodes, as required in the

F-3 CONNECTION DIAGRAM OF SIMULATED CONTROL-CONFIGURED VEHICLE (CCV)

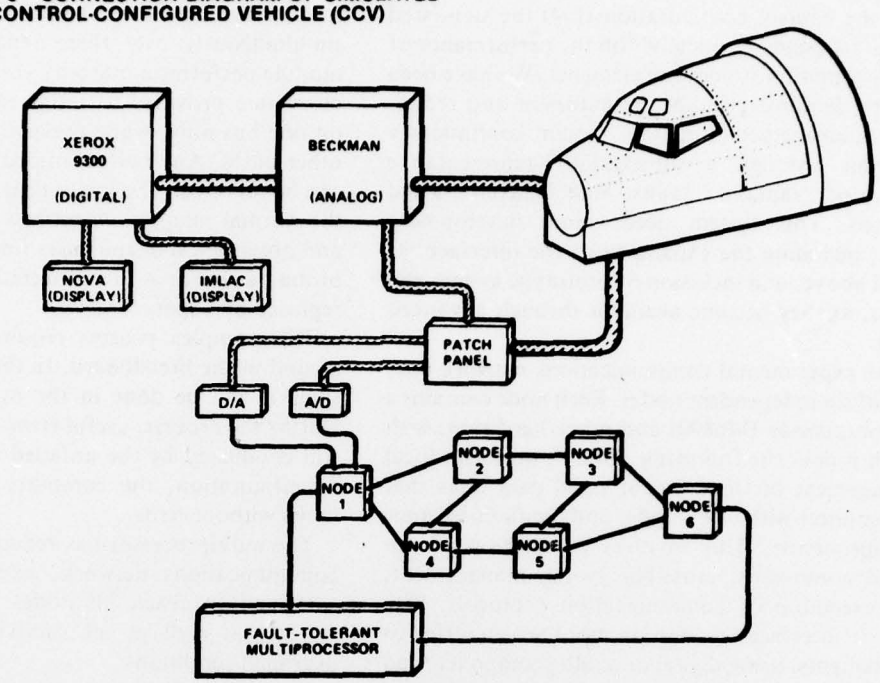


F-4 LOGICAL STRUCTURE OF CCV SIMULATION

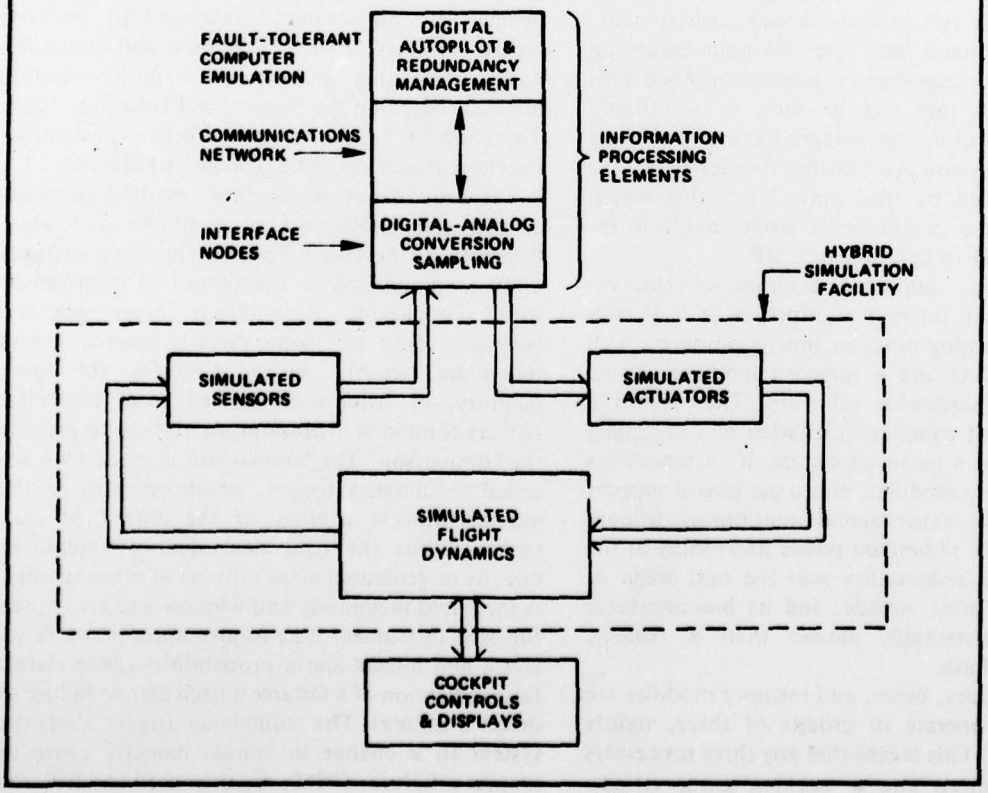

integrated system concept.

In the control configuration (F-4) the simulated aircraft depends "critically" on the performance of the information-processing elements. We have been able to demonstrate that the autopilot and redundancy-management functions remain continuously available through a substantially representative variety of "random" faults, both inadvertent and injected. This system needs more development work, including the expansion of the interface, as noted above, and inclusion of prototype system elements, as they become available through advanced R&D.

The experimental communications network consists of six independent nodes. Each node contains a microprocessor (M6800) and other hardware, with which it does the following. First, it performs local management of three digital serial data links that interconnect with other nodes and the fault-tolerant multiprocessor. This involves recognition of encoded commands, cross-bar switch management, and execution of communication protocols. Second, it interfaces to one or more sensor/effector components, using digital or analog components ad hoc. Third, it executes local programs, again ad hoc, to support the operation of those sensor/effector components.

Each of the experimental nodes uses approximately 40 semiconductor components, not counting its ad hoc program storage and interfaces. Each link is a full-duplex shielded, twisted pair, with a 40-kilobaud data rate. Without exceeding contemporary technology, a production node with a 1-megabaud rate can be built in a suitably compact form using chip integration and/or hybrid packaging. It is also possible that the electrical links can be replaced by fiber-optic links; this would greatly increase immunity to electromagnetic interference, such as lightning or EMP.

F-5 gives a simplified diagram of the experimental fault-tolerant multiprocessor.[12] It contains 14 processing modules (microcomputers with small memories) and 6 memory modules. These numbers are somewhat arbitrary. They reflect a partial stage of completion of what was originally intended to be a balanced system of 14 processors and 14 memory modules, plus a peripheral support computer. The experimental multiprocessor does not contain the redundant power and timing or the degree of bus redundancy that the next stage of development must include, and its bus interfaces are also substantially slower than a realistic prototype will use.

The processors, buses, and memory modules are designed to operate in groups of three, tightly synchronized. This means that any three processors can be organized into a working group (triad). Likewise, memory modules are organized into triads, and all intermodule transmissions occur simultaneously over three separate buses. Every module performs a majority vote on the buses, and means are provided to assign modules to transmit on one bus only, while preventing transmission on other buses. Any fault confined to a single module can be survived. The job in progress is completed in the normal manner. A process of reconfiguration and observation of the buses finds the fault and the probable source is disconnected from the bus and replaced by a spare module.

This complex process requires a fraction of a second in the breadboard. In the next development stage it will be done in the order of 10 millisec. During the process, useful error-free computation is still conducted by the unfailed modules; and after reconfiguration, the computer can again sustain faults without error.

The multiprocessor has redundant ports into the communications network, any one of which is sufficient to reach all nodes of an undegraded system, as well as all surviving nodes in most degraded conditions.

In parallel with development of the fault-tolerant multiprocessor and its associated data-communications network, we have devised effective sensor-failure detection and identification algorithms based on the method of analytic redundancy, mentioned earlier, which makes use of all pertinent kinematic, geometric, and dynamic relationships between sensed variables). Extensive analysis and simulation have shown that an efficient hypothesis-testing method, based on the Sequential Probability Ratio Test (SPRT),[14,15] proves very effective for comparing these sensed variables to identify failures.

The redundancy-management method proceeds in two steps—detection of a failure and identification of the failed element. Detection depends upon a comparison of the output of each sensor with comparable data from other sensors. Necessary data may come directly from a second redundant sensor, which measures the same quantity, or from data derived from dissimilar sensors combined in appropriate fashion to provide the comparison. The comparison is made by a so-called redundancy trigger, which operates on the moving window average of the output of each sensor minus the equivalent quantity measured directly or generated using outputs of other sensors. A threshold magnitude and window size are chosen for each instrument type to give appropriate false-alarm and missed-alarm probabilities (false alarm: false indication of a failure; missed alarm: failure to detect a failure). The redundancy trigger alerts the system to a change in status, namely, cause to suspect a failure. SPRTs are then used to verify the

alarm and identify the failed element.

As done in the simulator, the SPRT makes sequential observations of a process which is the difference between the output of the suspect instrument and the equivalent quantity generated from outputs of unfailed instruments. The SPRT gathers enough information to choose between two hypotheses: the two sources of information are compatible; the two sources are incompatible. Random sensor noise and other sources of error are included in the simulation as appropriate to represent typical instrument performance. SPRT provides a systematic test, well founded on statistical theory; it has proven highly effective for this application.[5]

The analytic redundancy methodology provides the basis for incorporating the additional sources of information necessary for unambiguous identification of failures. The simulation uses three types of analytic redundancy. Translational kinematic redundancy exists between the integrated output of the accelerometers, vertical gyros, and rate gyros and the outputs of the air-data sensors, (Mach meter, altimeter, and alpha and beta vanes). Rotational kinematic redundancy relates the integrated outputs of the rate gyros and the outputs of the vertical and directional gyros. Altitude kinematic redundancy exists between the altitude given by the altimeter outputs and the altitude computed as the double-integral of the accelerometer and vertical gyro outputs. Together these provide the basis for total coverage of all failures in the simulated sensor system.

Extensive "hands on" testing of the simulated system is in progress. Recovery from injected multiple sequential failures of sensors and computer elements has been repeatedly demonstrated during critical flight phases, such as simulated precision ILS approach. Effective fault detection, identification, and reconfiguration provide sufficient masking of faults to keep the crew virtually unaware of them. Once system checkout was completed, no missed alarms or false alarms were encountered in any of the simulated flights.

It seems unlikely that extrapolations of systems developed so far for aviation will ever yield the full-time flight-critical avionic systems that sophisticated and complex vehicles call for. Escalation of overheads in redundant elements plus the additional equipment required for redundancy management simply prohibit extrapolation of existing designs. But overheads can be greatly reduced by replacing independent subsystems with a *wholly integrated system approach*.

The technology to permit this step has been demonstrated at an early experimental level. It consists primarily of redundant sensor and effector architecture, a digital communications and control network (which also manages power distribution), a fault-tolerant computer complex containing at least one highly dependable redundant computer structure, validated software, and a family of algorithms

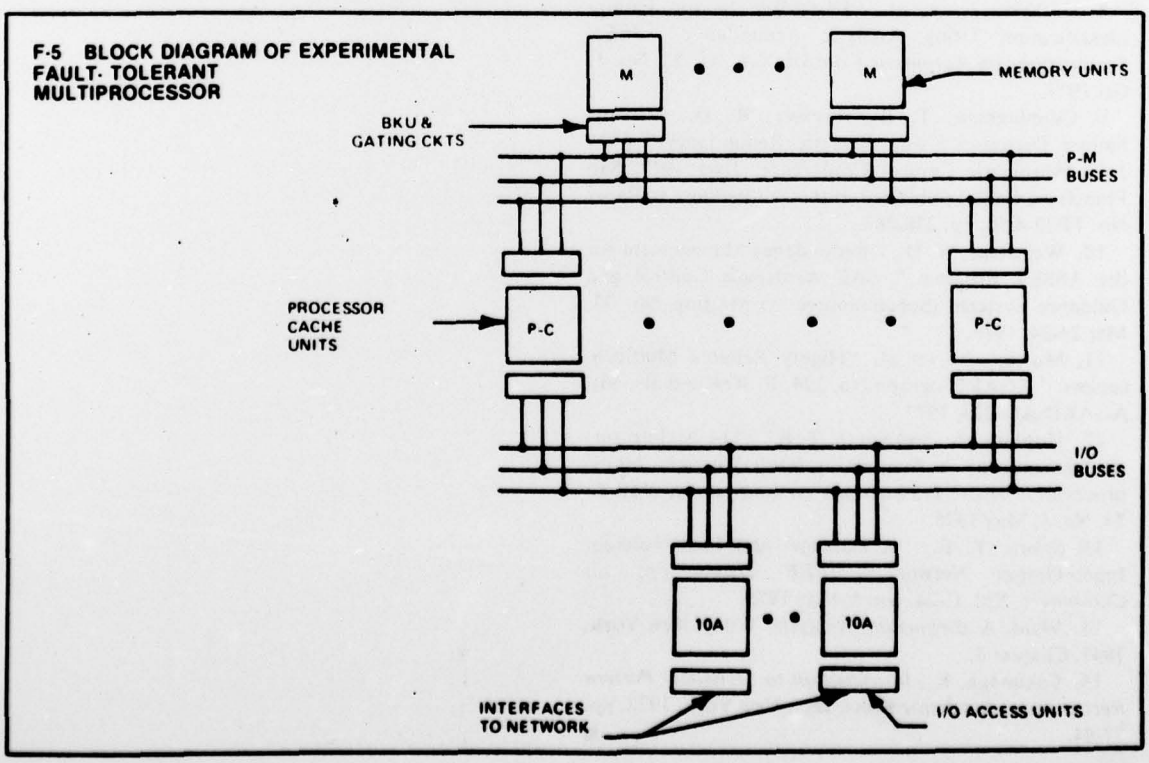


F-5 BLOCK DIAGRAM OF EXPERIMENTAL FAULT- TOLERANT MULTIPROCESSOR

that distinguish components that malfunction from those within specified limits. The development cycle that now treats these issues individually should produce a demonstrable merged system within a few years.

**References**


1. Ostgaard, M. A., and Swortzel, F. R., "CCVs: Active Control Technology Creating New Military Aircraft Design Potential," Feb 1977 *Astronautics & Aeronautics* (*A/A*).

2. Osder, S., "Chronological Overview of Past Avionic Flight Control System Reliability in Military and Commercial Operations," AGARD-ograph No. 224, P. R. Kurzhals, ed., AGARD-AG-224, 1977.

3. Bjurman, B., et al, "Airborne Advanced Reconfigurable Computer System (ARCS)," NASA CR-145024, Boeing Commercial Airplane Co.

4. Wensley, J., "SIFT—Software Implemented Fault Tolerance," *Proceedings* 1972 Fall Joint Computer Conference, AFIPS, Vol. 41, Part 1.

5. Hopkins, A., and Smith, T. B., "OSIRIS—A Distributed Fault-Tolerant Control System," *Digest Compcon*, IEEE Computer Society, San Francisco, Mar 1977.

6. Avizienis, A., "Architecture of Fault-Tolerant Computing Systems," Digest International Fault-Tolerant Computing Symposium, IEEE Computer Society, Paris, 1975.

7. Goldberg, J., "New Problems in Fault-Tolerant Computing," *Digest International Fault-Tolerant Computing Symposium*, IEEE Computer Society, Paris, 1975.

8. Deckert, J., et al, "F8-DFBW Sensor Failure Identification Using Analytic Redundancy," *IEEE Transactions on Automatic Control*, Vol. AC-22, No. 5, Oct 1977.

9. Cunningham, T. B., Poyneer, R. D., "Sensor Failure Detection Using Analytic Redundancy," 1977 Joint Automatic Control Conference, June 1977, San Francisco, Calif.; published in the *Proceedings* as Paper No. TP22-4:50, pp. 278-287.

10. Weinstein, W. D., "Redundancy Management for the ASSET Program," SAE Aerospace Control and Guidance Systems (Subcommittee A) Meeting No. 37, Mar 23-24, 1976.

11. Murray, N., et al, "Highly Reliable Multiprocessors," AGARD-ograph No. 224, P. R. Kurzhals, ed., AGARD-AG-224, 1977.

12. Hopkins, A., and Smith, T. B., "The Architectural Elements of a Symmetric Fault-Tolerant Multiprocessor," *IEEE Transactions on Computers*, Vol. C-24, No. 5, May 1975.

13. Smith, T. B., "A Damage- and Fault-Tolerant Input/Output Network," *IEEE Transactions on Computers*, Vol. C-24, No. 5, May 1975.

14. Wald, A., *Sequential Analysis*, Wiley, New York, 1947, Chapter 3.

15. Fukunaga, K., *Introduction to Statistical Pattern Recognition*, Academic Press, Inc., New York, 1972, pp. 77-84.


■