

12

LEVEL II

RADC-TR-81-220
Final Technical Report
August 1981



AD A108752

ANALYSIS OF BUILT-IN-TEST FALSE ALARM CONDITIONS

Hughes Aircraft Company

John G. Malin
Richard W. Highland

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

DTIC
ELECTE
DEC 22 1981
S B

DTIC FILE COPY

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, New York 13441

24

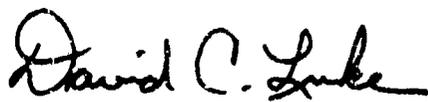
This report has been reviewed by the RADC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-81-220 has been reviewed and is approved for publication.

APPROVED:


DANIEL GLEASON, Capt USAF
Project Engineer

APPROVED:


DAVID C. LUKE, Colonel, USAF
Chief, Reliability & Compatibility Division

FOR THE COMMANDER:


JOHN P. HUSS
Acting Chief, Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (RBET) Griffiss AFB NY 13441. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document require that it be returned.

SUMMARY

Field support of military avionics and ground-based systems is generally based on the concept that such systems can be maintained by minimally trained technicians supported by a sophisticated, built-in-test (BIT) capability. In fact, BIT systems have not performed as efficiently as expected. This has placed unanticipated demands on maintenance personnel and has resulted in expenditure of excessive maintenance resources such as manpower, support equipment and spare parts.

The primary contributor to BIT inefficiency can be described under the generic term "false alarms" and this study was designed to address this problem by investigating false alarm experience for three representative systems. The study accomplished its objectives of investigation and determination of:

1. The causes of BIT false alarms and the relative frequency of occurrence of each such cause.
2. Design guidelines to minimize the occurrences and effects of false alarms.
3. False alarm rate prediction factors that provide for the evaluation of alternate BIT designs to determine their susceptibility to false alarms.

A major difficulty encountered in the investigation was simple identification of false alarms. It is intuitively obvious what is meant by the term false alarm (a BIT indication that fault-free equipment has failed) and the term has been used in specifications for years. However, the task of measuring false alarm rates is extremely difficult. The measurement difficulty is compounded by the fact that many actual failure events can masquerade as false alarms, such as intermittent faults which occur only under certain operational conditions. The measurement difficulty has been overcome in this study by supplementing theoretical definitions of false alarms with a consistent set of ground rules for breaking the impasse created when a BIT indication could be either true or false. Some error is introduced in this process but must be accepted as being unavoidable when analysis is limited to existing field data. By using this pragmatic approach, we have been able to quantify the problem and to break it down into its component parts. Having accomplished this, it was then possible to propose solution approaches.

Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

Table S-1 provides a capsule description of the BIT false alarm study. The major conclusion of the study is that the problem is amenable to solution, with the hypothesized solution being referred to as "optimal BIT." Theoretical performance of optimal BIT is compared with ordinary BIT in Figure S-1. As illustrated, the main defect with ordinary BIT is that extremely high probabilities of fault detection and isolation (demanded by military specifications) are only achievable by accepting a high incidence of false alarms.

Although specifications put limits on the allowable false alarm rate, such specifications are generally meaningless because it has been impossible to prove or disprove that the specification is being met (primarily because of the difficulty of identifying false alarms). The usefulness of BIT is seriously degraded by the presence of false alarms and it is hypothesized that most current BIT designs are operating beyond the point of greatest usefulness indicated in Figure S-1(b), i.e., the usefulness is in the region of diminishing returns. This should be compared with the hypothetical usefulness of the optimal BIT, illustrated in Figure S-1(d). In the latter case, the usefulness continues to improve with increasing BIT thoroughness. (A measure of BIT usefulness is the percentage of field problems resolved by using BIT. A measure of BIT thoroughness is the percentage of the system, weighted by predicted failure rates, that is tested by BIT.) As illustrated, this characteristic is achieved by suppressing the false alarm rate. The question of feasibility of optimal BIT thus translates into the feasibility of suppressing false alarms. More specifically, does the technology exist for accomplishing false alarm suppression and, if so, do we know how to utilize this technology to accomplish our purpose? This study answers both questions affirmatively. Microprocessors, expanded memories, sensors, components, circuits, etc. required to implement optimal BIT either exist or are in an advanced state of development. The problem of false alarms is sufficiently well understood to establish an overall approach at solving the problem and preliminary design guidelines have been generated.

The problem of false alarms is predominantly the result of BIT specifications and BIT designs being tailored to an ideal (noise-free) world. If all failures were in the form of hard, catastrophic faults and all systems performed precisely as they are theoretically supposed to, and if all environments within which systems have to operate were within specified boundaries, and if there were no external sources of RF interference, then BIT performance would be truly superb. But in contrast, the real world is extremely complex. All types of peculiar failure modes exist, many of which are intermittent in nature. Fault-free systems exhibit a wide range of variability and are prone to exhibit moments of abnormal or "anomalous" performance. Unique operational

TABLE S-1. CAPSULE DESCRIPTION OF THE BIT FALSE ALARM STUDY

<p><u>WHAT WE SET OUT TO DO</u></p>	<p><u>HOW WE WENT ABOUT IT</u></p>	<p><u>WHAT THE RESULTS WERE</u></p>
<p><u>THE PROBLEM:</u> The usefulness of BIT is seriously affected by the presence of false alarms. False alarms can degrade mission effectiveness of systems and contribute to the expenditure of excessive maintenance resources such as manpower, support equipment and logistic supplies.</p> <p>a. To define the problem. b. To quantify the problem c. To determine root causes of the problem d. To define approaches for minimizing the problem. e. To define prediction factors applicable to new systems.</p>	<p>a. Selected 3 systems for research. b. Accumulated and formatted data. c. Established analysis techniques. d. Analyzed the data. e. Reviewed past studies, including:</p> <ul style="list-style-type: none"> o Anomalous Maintenance Study o Missile-on-aircraft-test (MOAT) False Alarm Study o MORPEP Study <p>f. Reviewed selected studies from the literature. g. Interviewed field engineers and design engineers.</p>	<p>a. Problem clarified and defined. b. False alarm rates established for the three systems investigated. c. Root causes of false alarms itemized and relative frequency of occurrence determined. d. Design guidelines established to minimize occurrence and effect of false alarms. e. False alarm rate prediction factors for evaluating BIT designs determined.</p>
<p><u>CONCLUSIONS :</u></p> <p>a. Current BIT designs can be described as being suboptimal. They have been designed for an ideal (noise-free) world and are excessively sensitive for the real (noisy) world. b. The technology for solving the problem either exists now or is in an advanced state of development. c. The solution approach includes (1) memory and (2) filtering of false alarms via "smart" processing of stored data.</p>		

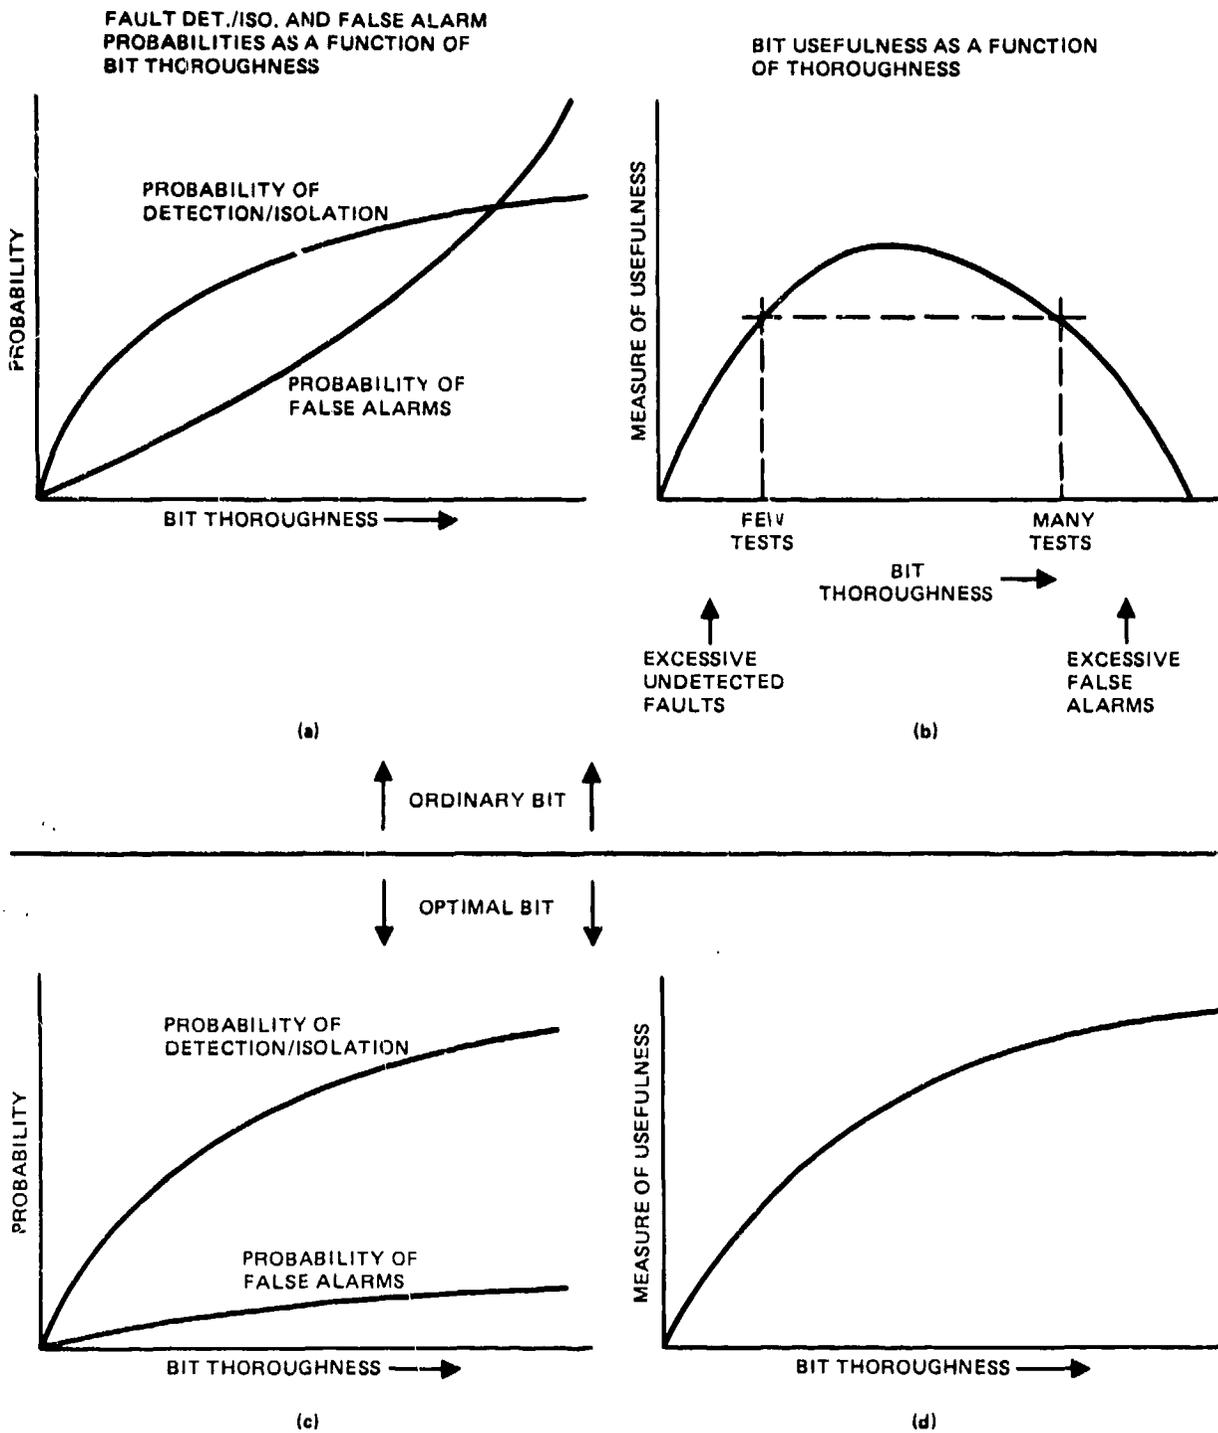


Figure S-1. Optimal Bit Compared With Ordinary Bit

conditions can cause fault-free systems to perform in a manner that is easy to misinterpret as an indication of a failure. Also, the real world environment (thermal, shock, vibration, humidity, power transients, etc.) can be sufficiently stressful, to cause systems to fail momentarily to meet performance standards. Understanding this phenomenon of "failure without a fault" is the key to understanding the bulk of the false alarm problem. During these incidents, "BIT does not lie," in the sense that BIT accurately detects anomalous performance. When BIT indicates a momentary signal excursion outside of test limits, the operator can have reasonable confidence that indeed the signal did exceed limits. However, more often than not, such anomalous performance is not a manifestation of a fault and it is a mistake to take a maintenance action based on the indication. Thus, by designing BIT for an ideal (noise-free) world, we have, to a great extent, created the false alarm problem. What is needed is a BIT that will not display a warning flag every time a momentary anomaly occurs, but one that will filter out those anomalies that do not warrant taking a maintenance action. The key to an optimal BIT might be described as supplementing existing (highly sensitive) BIT with a "smart box." Alternatively, each unit could have the "smarts" built into it.

The challenge of designing optimal BIT can be subdivided into three main areas:

1. System analysis, to define the intelligence that needs to be built into BIT. (How can false alarms be recognized? How can intermittent faults be separated from false alarms? A key issue is the type of filtering to be implemented: time thresholding, amplitude thresholding, relative frequency of occurrence, trend analysis, statistical testing, rate of occurrence.)
2. System design, to establish the functional definition of the required microprocessors, memories and other equipment. (What processing capability is required? How much memory is required? How can compatibility with intermediate-level maintenance be achieved?)
3. Equipment design, to include specific definition of the equipment that is required to implement the system functions.

Much of the analysis performed in support of this study can be used as an example of the type of "machine-analysis" capability that needs to be programmed into an optimal BIT. By relieving the maintenance person of the bulk of the interpretation task, the support concept of smart machine/ minimally trained technician becomes viable. It must be recognized, however, that occasionally situations will occur that have been totally unanticipated. At these times, the skilled maintenance person is invaluable. It is very unlikely that the man-in-the-loop concept can ever be eliminated.

PREFACE

This technical report presents the results of a study to investigate and determine (1) the causes of built-in-test (BIT) false alarms and the relative frequency of occurrence of each such cause, (2) design guidelines to minimize the occurrences and effects of false alarms and (3) false alarm rate prediction factors that will provide for the evaluation of alternative BIT designs to determine their susceptibility to false alarms. The study was performed for Rome Air Development Center (RADC) under Contract F30602-J0-C-0074. This report is prepared in accordance with CDRL item A002 and data item description DI-S-3591A/M.

Capt. Daniel Gleason was the Air Force monitor and the Support Systems and Maintainability Engineering Laboratory of Hughes Aircraft Company, under the management of Mr. R. A. Vande Steeg, was responsible for program execution. The program manager was Mr. E. C. Hamilton. Mr. John G. Malcolm was the principal investigator. Dr. R. W. Highland was the primary consultant. R. E. Davison and R. J. Dunlap, as well as other Hughes engineers, contributed to the report.

TABLE OF CONTENTS

	<u>PAGE</u>
List of Illustrations	xi
List of Tables	xii
1. Introduction	1
2. Background	3
2.1 Objectives of Study	4
2.2 Scope of Study	4
2.3 Organization of Report	6
3. Technical Approach	9
3.1 Description of Selected Systems	10
3.2 Data Characteristics	18
3.3 General Approach	20
4. Analysis Methods	23
4.1 Collection and Assembly of Data	23
4.2 Ground Rules for Identifying False Alarms	26
4.3 Special Analyses	30
5. Analysis Results	31
5.1 Rate of Occurrence of False Alarms	31
5.2 Causes of False Alarms	34
5.2.1 Category I False Alarms	34
5.2.2 Category II False Alarms	37
5.2.3 Probable Causes of Occurrence of Specific DPs.	47
5.3 False Alarm Rate Prediction Factors.	49
5.4 Distribution of BIT Callouts Among Aircraft.	66
6. Design Guidelines	71
6.1 Guiding Principles	71
6.2 System Design Guidelines	75
6.3 Procedural Guidelines.	79
6.4 Other Considerations	79
7. Conclusions	85
8. Recommendations	89

TABLE OF CONTENTS (Cont)

	<u>PAGE</u>
Appendix A. Analysis of Anomalous Maintenance Incidents	A-1
Appendix B. Considerations in Setting Electronic System Test Tolerances	B-1
Appendix C. Missile-On-Aircraft-Test (MOAT) False Alarm Study	C-1
Appendix D. Maintenance, Operation, Reliability and Performance Evaluation Program (MORPEP) Study	D-1
Appendix E. System 1 Supporting Data, Sequence 3 (Radar Confidence Test) DP Statistics and Rankings .	E-1

LIST OF ILLUSTRATIONS

<u>FIGURE</u>	<u>PAGE</u>
S-1 Optimal BIT Compared with Ordinary BIT	vi
3-1 BIT Structure for System 1	12
3-2 System 3 Trailer BIT Organization	17
3-3 Analysis Roadmap	22
5-1 Makeup of Intermittent/Inconsistent/Non-repeatable BIT Indications	38
5-2 Scatter Diagram, Number of Times DP Occurred Versus Number of Different Aircraft on Which DP Occurred	68
D-1 Simplified Diagram Illustrating MORPEP Inputs and Outputs	D-2
D-2 Analysis Road Map.	D-2
D-3 System Viewpoint of the Maintenance Process.	D-3

LIST OF TABLES

<u>TABLE</u>		<u>PAGE</u>
S-1	Capsule Description of the BIT False Alarm Study .	v
5-1	Summary of False Alarm Rates	34
5-2	Tabulated Data for Computing False Alarm Rates--System 1	35
5-3	Tabulated Data for Computing False Alarm Rates--System 2	36
5-4	System 2 Radar Failures Squawked by Pilots.	39
5-5	Possible Causes of System 1 Category II False Alarms	42
5-6	Number of Occurrences of System 1 False Alarms, by Root Cause and by Cluster Designation.	43
5-7	20 Top System 1 DP Failures Ranked by Number of Occurrences	48
5-8	Technique for Analyzing Filter Effectiveness Factor	54
5-9	Technique for Comparing Systems' Category I False Alarm Characteristics.	61
E-1	Occurrence of BIT DPs Among Aircraft	E-1
E-2	DP Listing, Ranked by Number of Occurrences.	E-2

1. INTRODUCTION

It is a commonly recognized fact that built-in-test (BIT) systems utilized in an operational environment do not perform as effectively as the military customer desires or the contractor expects. Symptoms of this ineffectiveness generally take one of two forms:

- (1) A fault indication when the tested equipment has not failed.
- (2) Improper isolation of an equipment fault; a fault-free unit is called out as being faulty when the fault is in another unit.

Fault indications under these conditions constitute false alarms. As an introduction to the false alarm problem, the following paragraphs are taken from the statement of work for this study ("Analysis of Built-In-Test False Alarm Conditions").

"The negative impact of false alarms on maintenance policies and support costs has been documented on a wide range of systems. The extent of these false alarms contributes to the expenditure of excessive maintenance resources such as manpower, support equipment, and logistic supplies. BIT systems that experience high levels of false alarms may be rendered ineffective due to the lack of confidence in the integrity of the failure diagnostic information. False alarms can seriously degrade the mission effectiveness of systems that incorporate BIT to perform system monitoring functions. Erroneous indications of a system's capability may result in an unnecessary mission abort depending on the criticality of the system under test.

"The basis of false alarm conditions rests in unanticipated design deficiencies. Providing designers with guidelines to anticipate and remedy these deficiencies will result in a BIT end product with high levels of operator confidence in the validity of the test results. Prediction factors will provide insight concerning the extent of the false alarm problem and allow for the structuring of maintenance policies to minimize the impact of false alarm conditions."

2. BACKGROUND

Maintenance and logistic costs of fielded weapon systems are almost invariably higher than anticipated. This situation can no longer be tolerated. The evidence suggests that much of the blame for excessive levels of unnecessary maintenance is assignable to BIT. Direct evidence of BIT inefficiency is the almost universal maintenance practice of putting more reliance on operator observation (real and imagined) than on BIT, i.e., maintenance actions are generally initiated only if an operator "squawk" has been generated. BIT is simply used to "confirm" the problem, typically being ground confirmation of in-flight squawks. This is totally at odds with the original concept of BIT, which was based on the ideas that (1) the best test is one performed with the system operating in the environment for which it was designed and (2) critical parameters can only be measured and assessed by BIT, not by the operator. The concept envisioned a BIT that was so credible that an in-flight detection/isolation of a fault could be accepted at face value without ground confirmation. Theoretically, the failure information could be relayed to the ground crew via RF communication and a replacement unit made available at the time of aircraft landing. In the real world, little credence is generally given to airborne squawks without ground confirmation. (In spite of the obvious defect that faults which only exhibit themselves at altitude and under operational conditions are going to be invisible to the ground crew.) A reasonable inference is that maintenance personnel have found, by trial and error, that maintenance performed solely on in-flight squawks and in-flight BIT indications is totally impractical. This is another way of saying that BIT fault indications generated in an operational environment are generally not believable. In contrast, when a BIT pass has been achieved, such an indication has extremely high credibility.

Perhaps the preceding discussion explains why the problem of BIT false alarms has been tolerated for such a long time. Operations people and maintenance personnel have been able to develop "work-around" techniques. (In a jocular vein, one field person indicated that BIT becomes a very effective tool when you learn to ignore it.) And, of course, BIT truly is a superb maintenance tool for "well behaved" faults (hard, catastrophic faults occurring singly). Thus, from an operations point of view, the problem of false alarms has been more of an annoyance than a catastrophe. This has been under peacetime conditions, however. Because the problem of false alarms has never really been defined in precise terms, let alone quantified, nobody can really say what the true cost is. Perhaps, by tolerating the problem, real problems are being masked which would surface very quickly in an emergency situation. In any event, the price being paid in the

maintenance/logistics world justifies research into the extent of the problem and the root causes.

2.1 OBJECTIVES OF STUDY

This study had as its objectives the investigation and determination of the following items.

- o The causes of Built-in-Test (BIT) false alarms and the relative frequency of occurrence of each such cause.
- o Design guidelines to minimize the occurrences and effects of false alarms.
- o False alarm rate prediction factors that will provide for the evaluation of alternative BIT designs to determine their susceptibility to false alarms.

2.2 SCOPE OF STUDY

The approach taken includes a means to identify and verify false alarm occurrences and to determine their causes. The methodology avoids, to the maximum extent possible, incorrectly designating as false alarms those situations which are actually intermittent failures occurring only under operational conditions. Event repeatability was assumed to be a key distinguishing characteristic. If equipment is truly defective, although there may be moments when the equipment is functioning properly, the same failure mode will tend to repeat. On the other hand, if the equipment is truly fault-free, BIT may occasionally generate a random, false failure indication but such indications will generally not repeat. These generalizations were translated into pragmatic ground rules for classifying failure events as being false (probably) or valid (probably). Some error will naturally result but it is believed that the classification scheme is reasonably accurate.

The analysis distinguishes between two false alarm categories, designated as follows.

Category I - False alarms induced by a prime system failure where a BIT system designates a failed item which, in fact, is operating properly instead of, or in addition to, designating the true failed item.

This definition was not intended to encompass ambiguous fault isolation permitted by specification. Thus, if BIT is permitted to isolate a fault to one of two units, one of the two is obviously going to be fault-free. Callout of the fault-free unit is by design and is not considered a false alarm.

This exclusion principle was not an issue in this study because the BIT systems investigated were generally designed to isolate faults to single units. Neglecting the relatively small number of instances where this was not the case simplified the analysis and introduced insignificant error.

Category II - False alarms that occur when no prime system failure exists, where a BIT system designates a failed item which, in fact, is operating properly.

To the uninitiated, this definition may appear to be quite straightforward. In fact, it is quite controversial. The controversy revolves around the classic conundrum, "What constitutes a 'failure'?" One school of thought favors the extreme position that there is no such thing as a false alarm since a BIT indication that a system has failed means just that (assuming BIT has been properly designed and is fault free). Even momentary anomalies of fault-free systems are considered valid failures. For example, assume that radiation from an adjacent radar has interfered with proper functioning of a radar system under test and BIT has sensed this and has generated a fail indication. Although it is certainly true that the system has "failed" in a functional sense (momentarily), the system has not failed in the sense that something has "broken." Given the circumstances that exist at the time, the detected performance "malfunction" is perfectly "normal." From a maintenance point of view, the BIT indication is a false indication of need for maintenance and therefore is a false alarm, or at least a "maintenance false alarm."* In this study, we have generally taken the maintenance viewpoint. The main theme of the study is that current BIT systems, although excellent detectors of momentary system anomalies, have to some degree become "maintenance generators" because they fail to distinguish between "normal" system anomalies and anomalies which are manifestations of faults. It is concluded that future generations of BIT can resolve the problem of maintenance false alarms by incorporating the "smarts" for filtering out normal anomalies. In any event, the issue of semantics and definitions must not be allowed to cloud the basic issue of excessive maintenance. By whatever name, the problem is real and needs to be resolved.

* Another example is the type problem which is "fixed" simply by resetting the system, e.g., by turning the power off and then on. (Computer "hangups" are an example.) The term "recoverable failure" is useful in describing this type problem.

The results of the study are based on and are applicable to avionics and ground-based systems. Design guidelines and false alarm rate prediction factors have been developed which are applicable to the early design phase and the detailed design phase of the BIT system. The term "built-in-test" includes those test systems which perform prime system monitoring, prime system checkout, and prime system fault detection and isolation, and which are an integral or associated part of the prime system.

2.3 ORGANIZATION OF REPORT

The report is organized in the following manner:

- o Section 3, Technical Approach, describes what we did to attain the objectives of this study. It describes the systems selected for research and the basis for such selection.
- o Section 4, Analysis Methods, describes the analysis techniques used for identifying false alarms and classifying them as either Category I or Category II.
- o Section 5, Analysis Results, presents the relative frequency of occurrence of Category I and Category II false alarms and discusses root causes of false alarms. False alarm prediction factors are presented and analytical procedures for evaluating alternative BIT designs to determine their susceptibility to false alarms are reviewed.
- o Section 6, Design Guidelines, describes specific approaches that can be utilized to minimize occurrences and effects of BIT false alarms.
- o Section 7, Conclusions, summarizes the major conclusions of the study, especially the conclusion that current BIT designs have not been optimally matched to system performance in the real world and the solution of the problem of BIT false alarms lies in an optimal BIT which has the "smarts" to distinguish between "normal" anomalies and anomalies which are manifestations of faults.
- o Section 8, Recommendations, provides suggestions for future research into the subject of false alarms. It is recommended that such research encompass the subject of intermittent faults as well as false alarms.

- o Appendices include brief summaries of some other studies that are particularly pertinent to this study. For example, the Missile-On-Aircraft-Test (MOAT) False Alarm Study is discussed in Appendix C. Computer printouts and other detailed tabulations generated during the study have not been included in this report in order to avoid cluttering the document with details of little interest to the general reader.

3. TECHNICAL APPROACH

The general technical requirements of the study were analysis, investigation and development of design and prediction methodology pertaining to false alarm occurrences. Specific tasks itemized in the study statement of work are listed below.

- o Identify and verify the occurrences of Category I and II false alarms. This task shall require the acquisition and analysis of a statistically sound data base. The task shall also require the identification and removal of data which are incorrectly designated as false alarms which, in fact, are intermittent failures which occur only under certain operational conditions.
- o Determine the cause of false alarms for each false alarm category. Investigation of the causes shall include but not be limited to examination of the specifics of inadequate test design, BIT hardware/software failures, environmental operating factors and BIT sensor tolerance levels.
- o Determine the relative frequency of each of the causes that produce false alarms for each false alarm category.
- o Provide design guidelines and procedures that shall minimize the occurrences of false alarms. The guidelines and procedures shall include but not be limited to the choice of BIT sensors and sensor measurements, BIT hardware/software interfacing and selection, and BIT circuitry design.
- o Provide prediction factors to estimate false alarm rates for each false alarm category as a function of BIT design. The prediction factors shall be based on prime system circuit design and complexity, and associated BIT system detection/isolation, specifications and design.
- o To ensure that the results of this effort are representative and sufficiently comprehensive, the contractor shall utilize data pertaining to BIT systems which shall include but not be limited to BIT design specifications, BIT design circuitry, BIT design techniques, and BIT operational performance.

How we went about accomplishing these tasks is the subject of this section. Essentially, attainment of all study objectives was accomplished by analysis of data representing three selected electronic systems.

3.1 DESCRIPTION OF SELECTED SYSTEMS

A statistically sound data base is a first prerequisite for meaningful analysis. Three representative military systems were selected for research for this study primarily on the basis that the available data for these systems, when taken as a whole, represented a meaningful and extensive data base. A second consideration was that the subsystems of interest (radars/weapon control) in these systems were designed by Hughes. This gave us the added advantage of being intimately familiar with the past history of this equipment and being able to supplement field data supplied by the military with data generated by internal testing and data generated by numerous in-house studies. The fact that the three systems are of differing "design ages" and have been designed, respectively, for three different military branches (Army, Navy, Air Force), means that the combined false alarm experience is probably very representative of most military electronic systems and that conclusions reached should have general application for these types of systems.

With respect to the completeness of the data base, there were some inadequacies. This is inevitable when using field data. These deficiencies were annoying but not insurmountable in performing false alarm analysis. The precise nature of the analysis performed on each system was tailored to the type data available. For convenience, the systems will simply be referred to as systems 1, 2 and 3. System 1 is a complex Navy radar/weapon control subsystem in a two-seater aircraft, system 2 is an Air Force radar subsystem in a one-seater aircraft, and system 3 is an Army artillery-locating ground radar.

Our basic technical approach is to integrate the separate results from analysis of the three systems into a single body of design guidelines. It is recognized that some risk is involved in drawing general conclusions from a limited data base. In order to minimize this risk, the current study has been supplemented by a review of many other studies, both internal and external.

SYSTEM 1. System 1 is a radar/weapon control system in a two-seater, Navy tactical aircraft that went through the first carrier deployment approximately six years ago. This avionic system contains 28 weapon replaceable assemblies (WRAs) per system. BIT detects system faults either in flight or on the ground and displays the most likely failed units (or two units, when there is an equal likelihood that either of two units contains the fault). BIT software contains about 55,000 words and less than 5% of the avionic hardware is dedicated to BIT.

BIT includes a series of automatic tests that are initiated by operator command. Additionally, a set of functions are continuously monitored and a series of special tests provided that are both manually operated and operated with computer assistance, for use by maintenance personnel. Approximately 60% of the avionic failure rate is subject to continuous monitoring. The operator-initiated confidence test is comprised of over 500 functional tests. (The radar portion of the confidence test is the focus of research performed for this study and is comprised of approximately 250 functional tests.) These tests are grouped in four sequences, with each sequence individually controlled by the operator. The confidence test is supplemented with four sequences of fault isolation tests. The BIT structure is illustrated in Figure 3-1. The tests within each sequence are computer controlled with test results displayed on a tactical information display (TID). The four confidence-test sequences can be performed in approximately 5 minutes. Information stored by the computer on the pass or fail status of each test is used to indicate on the display that part of the system that is faulty and to provide a degraded mode assessment. After each test sequence passes, a check mark appears in a BIT box on the TID. In the event of a failure, an "X" appears in lieu of the check, and the failed test number appears in the box. In addition, the faulty WRA number appears on the TID beneath the box. As each test sequence is completed, a check mark, a degrade symbol, indicating a mode is degraded, or an X, indicating a mode is lost, appears over the appropriate mode abbreviation on the TID. The completed display permits the operator to assess mission capability.

Certain functions in the system are continuously monitored throughout the mission. The functions monitored were selected on the basis of mission necessity, monitoring feasibility and whether or not the loss of the function would otherwise provide an indication of the condition to the operator. Failures detected during continuous monitoring are indicated by the appearance of a two- or three- letter symbol near the bottom of the TID. The letters are chosen to provide a key to the failure. The operator may initiate the confidence test to determine the tactical capabilities retained.

The confidence test and fault isolation tests together include approximately 1000 separate tests. Tests are identified by a "decision point" (DP) number. It should be noted that to identify a particular test, it is necessary to state both the DP number and the test sequence number (since different sequence tests may carry the same DP number). The availability of DP numbers in addition to WRA callouts provided us with a very unique opportunity for in-depth analysis. Accordingly, more time was dedicated to investigating system 1 than was expended on either of the other two systems.

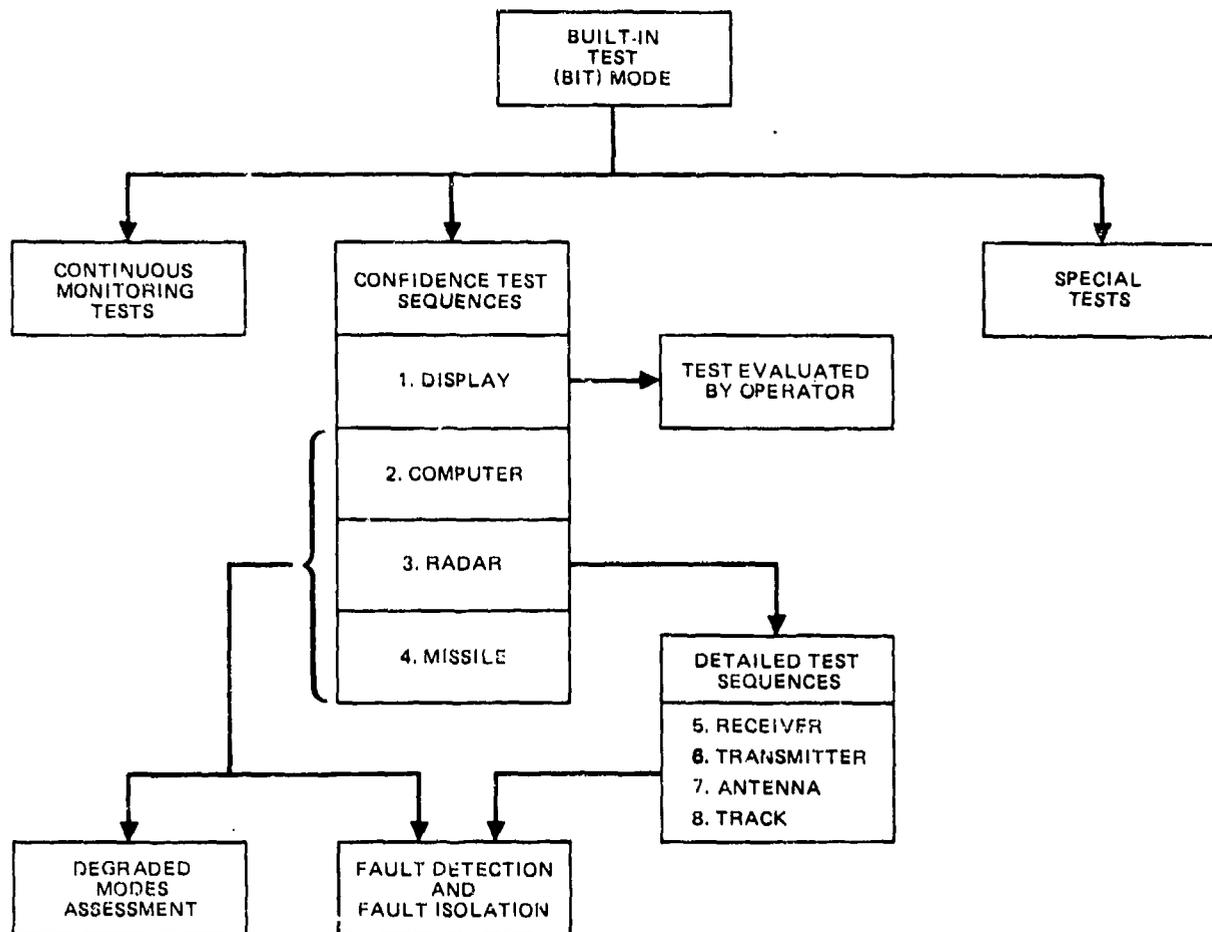


Figure 3-1. BIT Structure for System 1.

SYSTEM 2. System 2 is a fire control radar set in a single-seat, Air Force air superiority fighter aircraft that has been operational for about six years. There are 9 line replaceable units (LRUs) per system, with 8 of them being tested by BIT. BIT detects system faults either in flight or on the ground and identifies the unit which is most likely to have failed. A small reference table is available to maintenance personnel for identifying second and third choices.

BIT software during the period for which data was collected consisted of about 4000 words. More recently a programmable signal processor has been added which boosts the BIT software to 12,000 words. BIT software performs the functions of scheduling tests, configuring the radar system for the tests and evaluating test results.

Approximately 1.6% of the avionic hardware is dedicated to BIT. The BIT hardware provides the various test circuits and signals which are controlled by the software. There are approximately 150 tests, with the faulted LRU indicated by a fault flag. About 33% of the tests are contained in individual LRUs. The operator-initiated BIT can be performed in three minutes.

BIT test failures are recorded in two BIT matrices, one pertaining to a continuous monitoring BIT and the other to an operator-initiated BIT. The continuous monitoring BIT contains tests which can be performed without removing the radar from its normal tactical mode. The operator-initiated BIT is performed by taking the radar out of its tactical mode and placing it in the initiated BIT mode. In the initiated BIT mode, the BIT software controls the radar system in such a way that the required testing can be performed.

Each BIT matrix includes 144 cells. Certain of these cells are used to identify faulty units and the remaining cells to identify which BIT tests were failed. BIT fault isolation is accomplished in either of two ways:

1. By failing a particular test which directly isolates to an LRU.
2. By use of a deductive process in which a pattern of test failures is used as a basis for isolating to a given LRU.

In the case of fault isolation of the first type, the failure of certain tests can isolate a failure to a particular unit regardless of other failures which may also be indicated. Most BIT tests are of the second type. For tests of this type, proper operation of several LRUs is required for a test PASS. When the results of such a test are FAIL, isolation to the LRU which has failed is possible only by use of deductive logic.

In addition to the fault isolation information contained in the BIT matrices, BIT also includes fault isolation annunciators on the various LRUs. The annunciators can be used for direct fault isolation without use of a matrix.

SYSTEM 3. System 3 is an Army artillery-locating ground radar, consisting of an Operations Control Group housed in a shelter carried on a 1-1/4-ton vehicle and an Antenna-Transceiver Group mounted on a trailer. When shifting to a new location, the trailer is towed by a 5-ton truck. This vehicle also carries the generators which supply system power. The system is designed to achieve high availability, with 90% of all repairs performable in the field by the maintenance person normally assigned to the crew. Mean-time-to-repair (MTTR) is 30 minutes. The system features on- and off-line diagnostic software, built-in test equipment and automatic fault isolation to the replaceable unit level. The BIT system is much more sophisticated than those of the other two systems. This is, in part, a consequence of the fact that design constraints (e.g., weight, volume) are less severe for a ground system than for airborne systems. Some of the design features are identical with those that we hypothesize should be contained in an optimal BIT.

In the context of this study, it is of interest to note that operational success of system 3 is contingent on suppressing radar (i.e., non-BIT) false alarms. The radar uses new clutter-rejection techniques in its signal processor to filter out ground noises, enemy jamming and adverse weather conditions. Additionally, each track is tested against a series of discriminants by the signal and data processors to filter out unwanted returns from birds (feathered variety) or aircraft. These measures give the system an extremely low false location rate. (Perhaps the same kind of dedicated effort will be required to minimize the BIT false alarm problem that is the subject of this study.)

The shelter contains digital electronics, a signal processor, a computer, a printer, an operator console and a B-scope display. The trailer contains a large antenna, the radar transmitter/exciter and receiver and other analog equipment. BIT controls for the total system are provided in the shelter.

The shelter and trailer BIT tests are essentially independent of each other. However, the trailer BIT tests are based on the assumption that the shelter is fault-free. This assumption is necessary because the shelter is used for trailer BIT data collection. Fundamental differences exist in the structure of the shelter and trailer BIT tests, reflecting differences in the type equipment (digital versus analog). The trailer BIT is a more normal type of testing where collected data is processed and assessed on the basis of stored

tolerances. In contrast, shelter BIT employs the concepts of redundancy circuits, parity checks and other techniques appropriate to digital circuits. A key feature is the concept of running digital circuitry diagnostics off "signature data bases" in the computer. The reference "signature data" (measurements of signal transitions and timing) is derived from a system that is known to be good. The computer performs tests by injecting test signals and comparing output signals within the signal processor to the stored signature data base. Using "real world" performance data as a reference appears to avoid the problems created when a test standard is based on anticipated performance which is theoretically accurate but in fact is not representative of the performance of real equipment.

Pertinent failure information for both trailer and shelter BIT tests is communicated to the maintenance person via a printed message. The concept of reproducing the results of BIT testing in the form of a printed message (including the time and date of the test) appears to be one approach to eliminating some of the weaknesses leading to false alarms which exist in the airborne systems that have been reviewed. In implementing the concept of the printed BIT message into the system, great care was given to the human aspects of the problem. The design goal was to provide optimum convenience to the system operators and O-Level maintenance personnel. Messages were designed to be read and interpreted by personnel with relatively low skill levels. Minimum reference to technical manuals is needed to interpret BIT messages. As a backup -- but only as a backup -- system 3 includes features to be utilized by the more sophisticated user (for a more detailed assessment of system performance than is normally required).

BIT dynamically tests all major units of the system and consists of the following major types: (1) On-line System Self Test, (2) Off-line Status Test, (3) Off-line Fault Isolation Test (FIT), and (4) Integration Aid (for use by upper echelon maintenance personnel).

The system test (on-line BIT) does not require an operator action. It is automatically performed during the actual operation of the radar. Functional units of the radar are tested periodically by scheduling test beams at specified intervals and comparing the data collected to expected results. Appropriate error messages are generated if a fault is encountered. If faults are ignored, the displays will not be tied up with repeated notifications of the fault condition. Once a fault has been detected and declared, all subsequent declarations of the fault are inhibited until the "SYSTEM FAULT" button has been pressed twice in succession with no intervening faults. However, computer faults such as "parity error" are considered to be non-recoverable faults and halt the computer. The system self-test is best described as cyclic

with anomalies declared as faults. An "M"-occurrences-out-of-"N"-opportunities criterion is included for faults considered to be recoverable before a fault is declared and printed. On-line BIT provides a first step in fault isolation by indicating which off-line BIT program should be used to isolate the failure.

The two basic off-line BIT tests are the Status Test and the Fault Isolation Test (FIT). The primary purpose of the Status Test is to provide the operator with a level of confidence of system operability. The operator is provided with the capability of (1) running specific tests, (2) continually cycling a given test and (3) receiving a printout of report data not meeting test criteria. The Fault Isolation Tests are utilized to isolate faults to the lowest possible number of cards or assemblies. In addition, it is possible to use FIT as an exhaustive status test. The signal processor portion of the Status Test is embedded in the FIT.

The BIT message formats have been structured for clarity and simplicity. For example, the basic STATUS message consists of two lines. The first line identifies the test in which the fault was detected, identifies the general area of the fault (using a fault branch number decimal code) and indicates the time of detection (hours, minutes, seconds). The second line directs the operator as to the action he should take, typically identifying the specific (shelter or trailer) fault isolation test that should be run. The basic fault isolation test message is a multi-line printout. The first line identifies the failed test and the specific fault and the time of occurrence. The following lines provide fault isolation directions with replacement units listed in order of decreasing probability (although directions may be modified to include removal of a less likely faulty unit first because of ease of removal).

The basic organization of system 3 trailer BIT test is illustrated in Figure 3-2. Both the Status Test and FIT are layed out on a modular construction basis, with each major test module dedicated to a major functional unit of the trailer. The major test modules of the Status Test and FIT are illustrated. The major test modules are performed in the ordered sequence illustrated in order that the complete trailer may be systematically checked. This sequence dependency exists because an operational function of one major unit may be an integral part in the testing of another major unit and this operational function must be checked by its own major test module prior to its use in any other test module.

The BIT tests described in the preceding paragraphs have been tailored for use by the O-level maintenance person. Additional test flexibility is provided for the use of higher echelon maintenance personnel in the form of the Radar

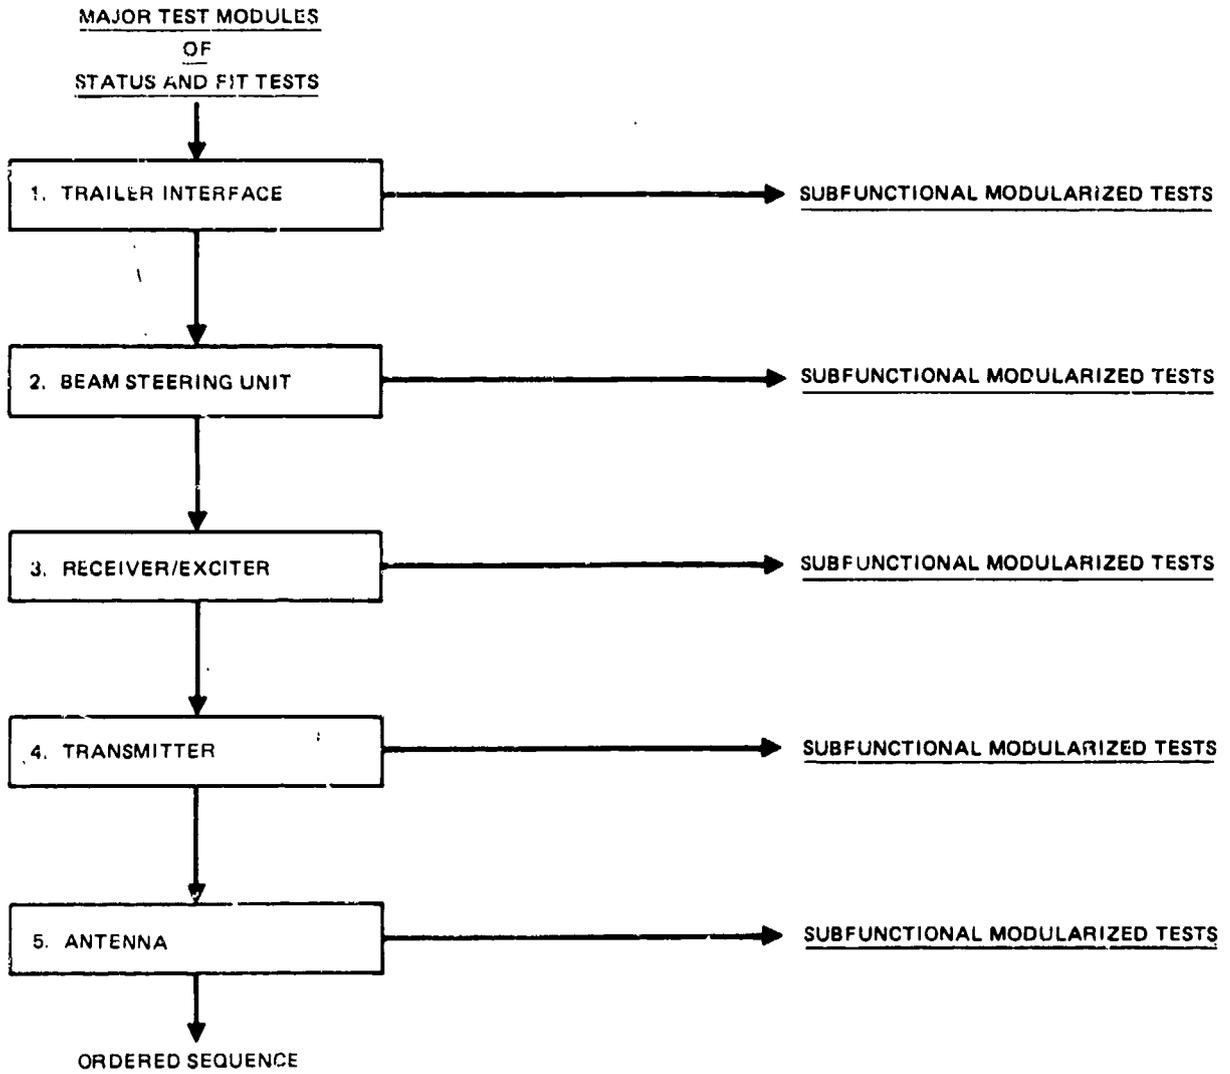


Figure 3-2. System 3 Trailer BIT Organization

Integration Aid (RIA) program. This program allows a user to set up special command table(s) and to repeatedly execute a command table or to cycle between several different tables of commands in a fixed sequence.

The program has many pre-stored command tables which may be used "as is" or adapted via function code to the specific needs of the user. The user may define special command tables and store them in designated spaces with the RIA program. These user tables may then be used in conjunction with the pre-stored table in any manner that the user chooses. The RIA program is especially useful for such special purposes as measuring power output, measuring noise figure measurement and measuring pulse characteristics.

3.2 DATA CHARACTERISTICS

It is important that the reader appreciate that the BIT-related data utilized in this study was not collected specifically for identifying BIT false alarms. The available data was collected for other purposes, such as monitoring operational reliability and maintainability. Therefore, it was necessary to evaluate the available data and determine how it could be analyzed in relation to BIT false alarm objectives. The best opportunity for achieving these objectives occurs where quantities of BIT-related data are large. This permits the data to be separated meaningfully on the basis of BIT false alarm criteria. The BIT data from systems 1 and 2 was ample in quantity and quality but the available data base from system 3 was very limited. This is because of the developmental stage of system 3 and because the data came from only two systems. In contrast, there are hundreds of aircraft-installed systems 1 and 2 and these systems have been operational for many years. This latter point should not be misconstrued to mean that the respective designs are "frozen." Quite the contrary, these designs are in a state of flux and are being continually upgraded. Thus, system 2 has recently added a programmable signal processor and system 1 is in the process of evaluating the performance of digital modifications for replacing many of the analog units. Thus, any false alarm rates arrived at in this document do not precisely reflect performance of latest configurations. Realistically, however, figures of merit of this type seldom exhibit sudden, dramatic improvements.

Types of information required to identify BIT false alarm conditions include the following:

1. BIT pass/fail data plus specific BIT tests failed.
2. LRUs/WRAs identified by BIT as being faulty.

3. Priority of LRU/WRA removal (if more than one unit is called out by BIT).
4. Organizational-level (O-level) maintenance action and apparent effectiveness.
5. Intermediate-level (I-level) maintenance action.

Most studies of the BIT false alarm problem are faced with a totally inadequate data base, typically based solely on maintenance action reports and often limited to I-level data only. Traditionally, O-level CND (cannot duplicate) rates and I-level no fault rates have been accepted as being virtually synonymous with false alarm rates. In fact, these parameters should be considered only as very coarse indicators of false alarm rates. Lack of credibility results from the fact that the "confirming" test is performed under a totally different environment than the environment in which the fault is initially detected. For example, clearly it is going to be impossible to "duplicate" on the ground avionic failure modes which exhibit themselves only under the stress of in-flight environmental factors. Also, there are usually so many differences between flightline and shop environments and between BIT and shop tests (e.g., system test versus unit test) that it should be expected that many valid BIT-detected faults are not going to be detected at I-level. Because of the deficiencies in the traditional approaches to investigating the false alarm problem it is not surprising to find wide disagreement whether the basic problem is really one of BIT false alarms or is actually one of hidden defects. Although there are many strong opinions as to the relative proportion of the two contributing factors, there is little objective data to support such opinions. The approach of this study is to use the traditional measures of false alarms--but only as a coarse guide--and to focus on the area for which there is the greatest ignorance: the general performance of BIT in an operational environment, with emphasis on the most significant characteristic, namely repeatability. (As will be seen, lack of repeatability is the key villain of the drama, with this characteristic more a reflection of system performance variability than a reflection of BIT circuitry. The solution lies in supplementing BIT with "smart" functions for recognizing normal variability.)

For purposes of studying the BIT false alarm problem, the data base compiled for this study is far superior to any data base that can be compiled from standard military maintenance data systems, for the following reasons:

1. The data for system 1 included complete MAF (maintenance action form) data, including narrative information. This narrative information is often very informative but is not available from the Navy 3M maintenance data system. Of

even more importance, BIT data were made available in the form of "BER" (BIT evaluation report) cards, on which the radar operator records all BIT indications, including tests which failed, units called out as being faulty, and failed modes. This information can be correlated with maintenance actions, by comparing the dates on the BER and MAF forms. The BER data are not included in the Navy 3M system. The BER cards frequently show operator comments which are useful in understanding peculiar circumstances.

2. The data for system 2 was collected by contractor personnel on site at the various Air Force bases. They made a special effort at capturing data as completely and with as much accuracy as possible, and, possibly more important, they made a complete record of BIT results and correlated maintenance actions with BIT results by recording both types of data (when related) on a single data-collection card. It is a straightforward matter to compile the cards on an aircraft by aircraft basis, so that we have a fairly complete historical record of BIT performance over an extended period of time for a fairly large number of different aircraft.
3. System 3 was undergoing a reliability demonstration test and so all pertinent data, both BIT and maintenance data, were being carefully collected and recorded. This test was being conducted by contractor personnel so there was no problem in acquiring the data and utilizing the results of analyses being made for reliability assessment.

3.3 GENERAL APPROACH

This section will describe in general terms how we went about the task of performing false alarm analysis. The next section (Section 4, Analysis Methods) is dedicated to a detailed discussion of the analysis methodology.

After accumulating as much pertinent data as possible for each system, we reviewed this data base with the idea of deriving a general understanding of the type intelligence that could be derived. System 1 data provided a large amount of detailed information on specific test failures. System 2 data provided much information pertaining to O-level "cannot duplicates" and I-level incidents of units checking no fault. System 3 provided information on advanced BIT techniques. As we reviewed and analyzed the data, we developed procedures for false alarm identification, for determining false alarm frequency and for identifying false alarm causation factors. At the same time, we gained insight from the data as to false alarm prediction factors and design features which would lead to reduced false alarm rates.

Figure 3-3 is an idealized description of the process described above. An important first step was to eliminate data that was irrelevant to the subject of false alarms. But even after maintenance events were recognized as being potential false alarms, it took considerable judgment to sort out the events that were deemed to be actual false alarms. In some cases, there was sufficient doubt as to prevent such classification. In order to facilitate classification of false alarms into either category I or category II, we explored a number of different analytical techniques. We eventually derived a set of ground rules (described in Section 4) for simplifying this process.

After all BIT callouts had been classified as being either valid or false alarms, it was then a straightforward process to compute the rate of occurrence of each type false alarm. For the category II false alarms (i.e., fault indications when there is no fault), the most significant index was assumed to be the fraction of total BIT indications falling into this category. After filtering out category II false alarms from the data, this left the true failure incidents. We then computed the fraction of true failure incidents that fell into category I. (It should be recalled that a category I false alarm represents a true failure incident. It is only false in the sense that the wrong unit has been called out.) Other percentages can be computed from the data provided, if the reader so desires.

In order to determine the root cause of the false alarms, our approach was to focus on the specific tests that failed most frequently and engineering analysis was performed on this subset of data. To facilitate this type of analysis, many different techniques were utilized, including statistical analysis of the data and detailed investigation of the way in which the offending tests were mechanized.

Having gained engineering insight into the factors causing BIT false alarms, it was then possible to draw up a set of design guidelines for minimizing the problem. Definition of these factors also led quite naturally to the development of false alarm prediction factors.

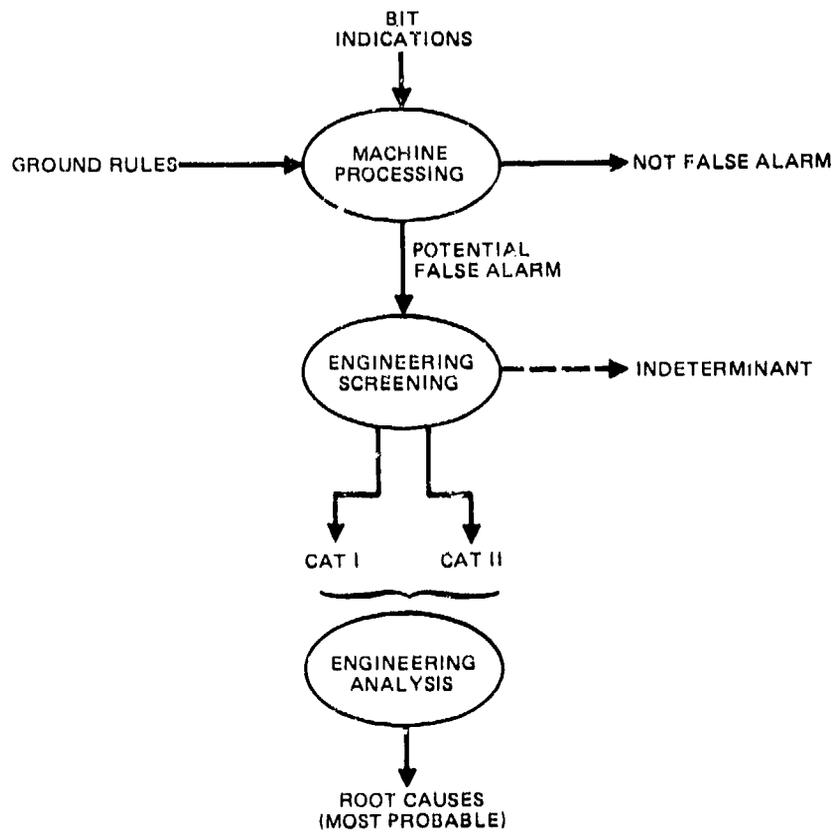


Figure 3-3. Analysis Roadmap

4. ANALYSIS METHODS

This section describes the methodology used to achieve our goals of establishing false alarm rates and to uncover root causes of the false alarm problem. The task of quantifying the problem is particularly challenging because of the elusive nature of the subject and because of inadequacies inherent in a data base constructed from field data. The analysis approach was guided by the philosophy that analysis must be partly unstructured, consisting of uninhibited research and guided by intuition. Much of our investigation can be described as being unstructured, especially for system 3, which had relatively little available data. On the other hand, it was recognized from the beginning that the amount of data for systems 1 and 2 was so extensive and so varied that the task of sifting this data in the search of false alarms could be overwhelming unless the task was carefully organized. What we set out to do, as a first step, was to screen the available data and to organize that part of the data pertinent to the false alarm study in a manner that would facilitate investigation. Our objective was the creation of a notebook for each system in which all pertinent data was organized (1) on a "per aircraft" basis and (2) on a calendar basis. This has been accomplished in the form of two, 100+ page notebooks, one for each system, representing a compilation of field data collected over a period of time of approximately one year (covering 1979 and 1980) and encompassing more than 30 aircraft per system.

4.1 COLLECTION AND ASSEMBLY OF DATA

We describe below how we went about achieving our goal of generating the system 1 Maintenance Data Notebook. Essentially, the same steps were taken for system 2.

1. COLLECTION OF RAW BIT DATA

A large supply of BER cards was available but a rapid inspection indicated that many of these had missing data or other deficiencies. It was noted that one squadron, a training squadron, did an especially conscientious job at filling in BIT data, including failed tests (DPs) and BIT unit callouts (WRAs). Accordingly, this squadron was selected for detailed investigation and the BER forms for this squadron were culled out (more than 4000 written against 31 different aircraft).

2. ORGANIZATION OF RAW BIT DATA

The BERs from the selected squadron were sorted in the following manner:

- (a) A file was set up for each aircraft, to contain all the BER cards written against that particular aircraft.
- (b) Within each file, the BERs were arranged in order of the calendar date on which the BER card had been generated. (Date is important because this is the parameter by which we are able to link BIT indications with subsequent maintenance actions, as described on maintenance action forms.)

3. COLLECTION OF PERTINENT MAINTENANCE DATA

We utilized a Hughes Maintenance Data System for this purpose. This data system is supplied with two sources of data:

- o Navy 3M MAF (maintenance action form) data, supplied by the NAVY on magnetic tape.
- o Narrative data written on the original MAF form but not included in the Navy 3M system. (A contractor representative on the base where the squadron is located, collects copies of the original MAFs. on a routine basis, and inputs the narrative data via a data terminal located on the base.)

A computer printout was generated for the 31 aircraft for which we had BER data, sorted in the same manner as the BER files, that is, by aircraft and by date.

4. GENERATION OF REFERENCE DOCUMENTS FOR CORRELATING BIT/MAINTENANCE DATA

The desired Maintenance Data Notebooks were generated simply by combining pertinent BIT data from the BER file with pertinent maintenance data from the MAF file (with the combined data sorted by aircraft and by date). As a minimum, the BIT data included failed tests and identification of the units indicated by BIT as having failed. Additional BER information was recorded if important to the false alarm study. For example, if it was noted on the BER card that the temperature warning light came on, this would indicate the probable cause of the failure condition detected by BIT. As a minimum, the recorded maintenance data included the maintenance action taken, if any. As backup information, it was also noted whether or not units checked faulty at I-level.

The data base utilized for this study can be summarized as consisting of the Maintenance Data Notebooks described above, the BER file (as a backup) and the maintenance data printout (as a backup), plus numerous other special-purpose printouts and data tabulations.

Having constructed the data base described above, it was then necessary to design an analysis plan for focusing on the key issues of (1) identifying false alarms, (2) determining their rate of occurrence and (3) determining their root cause. The first issue is inherently the toughest. Unfortunately, errant fault indications do not carry little flags saying "I am a false alarm." The traditional designators of O-level "cannot duplicate" (meaning that it has not been possible to duplicate an in-flight squawk on the ground) and I-level "no fault" (meaning that it has not been possible to duplicate a flightline squawk in the shop) are indicators of false alarms but are considered too coarse for our purposes. They are contaminated by many non-BIT aspects, such as skill levels of maintenance personnel and quality of shop equipment. Furthermore, many I-level "confirming" faults actually are totally unrelated to the BIT symptoms causing the unit to be removed from the system.

Considerations like these led us to believe that identification of false alarms should be based solely on consideration of O-level data. Of course analysis of O-level data also has its problems. With a fully instrumented system and if all potential sources were being monitored, O-level identification of false alarms might be straightforward. For example, if stray RF energy from an adjacent interceptor were detected simultaneously with a BIT indication of anomalous performance of a system's radar, the BIT indication could instantly be identified as a probable false alarm. Such capability does not exist in the real world, especially not in tactical systems.

So the challenge is to use some indirect method to identify false alarms, utilizing our available data base. Any such method must be based on some characteristic that is unique to false alarms. Non-repeatability is believed to be the key. True defects or flaws in a system--in contrast to false alarms--are (generally) permanent and can be characterized by repeatability of failure symptoms when BIT is run. In the case of hard faults, the failure symptoms will repeat every time that BIT is run.

In the case of intermittent faults, the recognition problem is more difficult because the failure symptoms may or may not be detected the next time that BIT is performed (depending upon whether or not the fault happens to be in a failed state). Nevertheless, if the intermittent fault represents a permanent flaw, the failure mode will eventually

recur. Thus, even with intermittent faults, repeatability is a key consideration. This can result in many different strategies for separating out intermittent faults from false alarms, but they all have the essential ingredient of setting a time window during which it is observed whether or not the same problem recurs. Recurrence is taken as indication of an intermittent fault, while lack of recurrence is taken as evidence that the initial failure indication can be written off as a false alarm. In our analysis, the time window was taken as two missions. In effect, we are giving the failure mode an opportunity to recur during about four hours of limited continuous monitoring plus perhaps about 8 runs of initiated BIT. Lack of recurrence is taken as evidence that the initial occurrence can be written off as a false alarm.

With these observations in mind, we created ground rules for identifying false alarms that, when applied to a large mass of data such as we were looking at, will positively distinguish between false alarms and hard faults and will tend to distinguish between false alarms and intermittent faults. Some error is inevitable. For example, there is a class of faults that is "self-healing," such as dirty contacts which are cleaned by the act of removing the unit. Per the ground rules, these will be incorrectly classified as false alarms. To compensate for this error, in pinpointing root causes of false alarms, we leaned heavily on engineering analysis, particularly analysis of those tests which failed most frequently.

4.2 GROUND RULES FOR IDENTIFYING FALSE ALARMS

The main criterion for a BIT fault isolation "success" is assumed to be disappearance of the BIT symptoms of a problem when the maintenance action called for by BIT is taken. (Note that this is totally independent of whether or not removed units check faulty at I-level.) Conversely, if the BIT indications remain unchanged following the maintenance action, the BIT indications can be classified as a CAT I "false alarm." In establishing the rate of occurrence of false alarms, it is, of course, necessary to establish how many discrete occurrences of false alarms (CAT I and CAT II) have occurred within the sample of data. One approach is to categorize every BIT indication as valid or false. However, since corrective action following a BIT indication is frequently postponed, this approach would generate meaningless statistics. To understand this, assume the existence of a single hard fault and assume that no maintenance action is taken over a period of time when 3 missions are performed. The record would show 3 separate BIT callouts, if the BIT test is performed once per mission. Even if the BIT callout is eventually determined to be a false alarm (category I), it would be grossly misleading to state that 3 BIT false alarms had occurred. If maintenance were delayed 6 missions, would we say that 6 false alarms had occurred? Or if the delay were

"N" missions, would we say that "N" false alarms had occurred? Clearly, the "N" statistic would be nothing more than a measure of maintenance delay and would contribute very little of a basic nature to our understanding of the false alarm problem. For the above situation, there was a single false alarm, in a generic sense. In other words, meaningful false alarm analysis must address the basics of the problem by filtering out simple repetitions of a single failure event. This is accomplished by grouping the repetitions together into a single "cluster" of events and then classifying the cluster. In our example of three separate occurrences of the same invalid BIT callout on three missions, this would be treated as a single cluster and would be counted as a single false alarm for purposes of computing false alarm rates.

For purposes of classifying false alarms, a cluster is defined as a sequence of three or more events all involving the same unit. (Our analysis focused on unit callouts and did not consider whether or not the callouts were the results of different test failures.) The events to be considered are (1) BIT callout of the unit, (2) removal/replacement of the unit (whether or not there is a recorded BIT callout) and (3) the recurrence of the BIT callout on the next mission. In the following discussion, any report of any of these events will be referred to as a "squawk." This convention is used because the bulk of the reported data comes directly from cards filled in by the pilot and/or radar operator. Each of the three events is identified with a single aircraft mission. To be considered a cluster, there must not be any long time-gaps when the unit is not being called out by BIT. More than half of the squawks generated during the period of time of the cluster must contain a BIT callout of the unit (or an indication of removal of the unit) and there must never be a "gap" of three or more squawks which do not contain such an event. Examples of clusters:

- o Two out of three sequential squawks indicate a BIT callout of an 011 unit.
- o Three of five sequential squawks include BIT callouts of an 031 unit.
- o Ten of 17 sequential squawks either call out an 031 unit or indicate that an 031 removal action has been taken (and where intervening squawks not dealing with the 031 unit only occur singly or in sequential pairs).

The ground rules for identifying and classifying BIT false alarms are summarized below.

1. Clusters are assumed to be CAT I false alarms if a unit removal action is followed by recurrence of the same unit callout, i.e., it is assumed that a real problem exists

but BIT is not properly isolating the problem. Any cluster encompassing multiple removal and replacement (R&R) of the same unit is also assumed to fall into this category.

2. If no unit R&Rs occur throughout a cluster, the cluster is interpreted as a valid detection/isolation under the conditions that the cluster terminates with a unit removal and the next two squawks are clear of callouts of the same unit.

To illustrate cluster analysis, we use the following definitions:

- B = BIT indicates unit faulty; no maintenance.
- (R&R) = unit is R&R'd (with or without BIT indication).
- O = BIT does not indicate unit faulty; no maintenance.

EXAMPLE

<u>Number of Successive BIT Callouts Before BIT Indication Clears</u>	<u>Symbolic Representation</u>
6	031: BBBBBB(R&R)00

These incidents are assumed to be the result of "delayed maintenance." As such, they are not false alarms and BIT has correctly detected and isolated the problem.

3. Clusters with no removal actions of any kind and where the BIT LRU callout eventually stops being generated are assumed to be CAT II False Alarms. Since maintenance personnel are not taking any action to correct the indicated problem, it is assumed that they understand the significance of the display and deem it not to be of importance relative to missions being performed, i.e., not a real problem.

EXAMPLES

<u>Number of Successive BIT Callouts Before BIT Indication Clears</u>	<u>Symbolic Representation</u>	<u>Number of Instances In System 2 Data Base</u>
3	BBBOO	9
4	BBBBOO	1
5	BBBBBBOO	3
6	BBBBBBBOO	1
7	BBBBBBBBOO	3
15	BBBBBBBBBBBBBBBOO	1
21	BBBB.....BBBOO	1

4. The cluster ground rules are also generally applicable to pairs of events. If the first item is a simple BIT callout and the second event is a unit removal and the problem goes away for two or more squawks, the pair is assumed to represent a valid BIT detection/ isolation and not a false alarm. If no removal action is taken and the problem still disappears, this is assumed to represent a CAT II false alarm.

EXAMPLES:

Non-False Alarm Events: B(R&R)00

CAT II False Alarms: BB00, BOB00 (21 occurrences of these types in system 2 data base)

5. The same ground rules are generally applicable to single maintenance events and single BIT callouts. A single BIT callout with no removal action is considered to be a random false alarm of the CAT II type if followed by two flights with no callouts (assumed to reflect random system performance variability, caused by system transients, momentary environmental stress, etc.). A BIT callout preceding a single removal action is considered a non-false alarm event when the removal action is followed by two flights with no BIT callouts. When there is no record of an operator BIT prior to the removal action, it is assumed (based on detailed analysis of a sample of System 1 field data) that BIT was in fact utilized by maintenance personnel in two thirds of such events. (Under field conditions, it is mandatory that in-flight BIT results either be put into memory or be manually recorded, for subsequent use by maintenance personnel. As a consequence, we were able to obtain an excellent record of in-flight BIT results. On the other hand, when maintenance personnel use BIT, there is no strong reason why they should record BIT results or even to indicate whether or not BIT was used. As a consequence, our data base reflects many unit removals where there is no indication of whether BIT was used or not. BIT would not be used, for example, if a defect was obvious by inspection or observation. Our ground rule is to use a weighting factor of 2/3. For example, in system 2, there were 430 isolated removals without any record of a BIT and so 287 of these events ($= 2/3 \times 430$) were considered to represent incidents where, in fact, maintenance personnel utilized BIT and the maintenance action was successful.)

EXAMPLE:

CAT II False Alarm: B00 (161 occurrences in system 2 data base)

False alarm rates have been computed for systems 1 and 2 using the ground rules described in this section. These rates are presented and discussed in Section 5, Analysis Results.

4.3 SPECIAL ANALYSES

The analyses described in the preceding paragraphs have dealt with BIT performance over periods of time and have correlated BIT unit callouts with maintenance actions taken. In addition to BIT unit callouts, the BER cards related to system 1 have provided us with a wealth of detailed information as to which tests have failed, in the form of DP (decision point) numerics. We have taken advantage of this information in two ways. Firstly, we have simply identified which DPs occurred most frequently and then have performed engineering analyses of these tests, on the assumption that these tests are most likely to be associated with false alarms. Results of these analyses are presented and discussed in Section 5. In addition to this structured approach, we have sorted the DP information in every reasonable way we could think of, without having any particular objective in mind, but simply for the purpose of clarifying such issues as to whether or not particular aircraft could be singled out as having peculiar characteristics. These data are extremely informative but are considered too detailed to be included in this report.

5. ANALYSIS RESULTS

The results of analyses performed for the BIT false alarm study are summarized in this section. These results logically lead to design guidelines for avoiding the problem, presented in the next section. It might be noted that we have consciously attempted to conduct most analysis from a system point of view, with detailed analysis in a supportive role. It was hoped that such an approach would lead to system solutions (i.e., generic approaches) to the false alarm problem. The approach was adopted in recognition of the fact that attacking the problem on a "bits-and-pieces" basis, as has been done over the years, has not brought very satisfactory results. In effect, we were continually on the lookout for large-scale, generic problems that could be solved with one stroke, so to speak, by a large-scale system approach. For example, if it can be shown that the predominant cause of the problem is a general tendency of complex, military systems to exhibit momentary anomalies unrelated to the presence of faults, it would be more efficient to develop a general, unified approach for coping with this characteristic rather than trying to upgrade subordinate tests on a test-by-test basis.

5.1 RATE OF OCCURRENCE OF FALSE ALARMS

It should be recalled that system 3 is still in an early stage of development. For comparison purposes it seemed appropriate to obtain data from systems 1 and 2 when they were roughly at the same stage of development. The early experience of all three systems is strikingly similar in that the major difficulty initially encountered for each system has been with non-hardware, non-fault system anomalies, as described below.

SYSTEM 1 EARLY EXPERIENCE

The Navy customer became so alarmed at the high incidence of system anomalies during early flight testing that the customer insisted that the monthly reliability report be expanded to include a regular report on system anomalies. Initially, there were approximately 3 reported troubles per flight hour, excluding troubles which led to the removal of hardware and excluding troubles which had previously been reported. Typically half of the troubles were never "confirmed." It was observed that the number of troubles observed per flight hour was a function of delivery of new or modified software, modification of the equipment and the testing of new modes or parameters. Exhaustive effort was expended in trying to understand and correct root causes of the individual anomalous conditions. At the end of approximately a year, the rate of reported troubles per flight hour was reduced to slightly less than two, with half of these not being confirmed. This was

considered satisfactory and the extra effort pertaining to resolution of system anomalies was discontinued. The biggest single type of corrective action was in the software area. In recognition of the fact that some anomalous performance must be considered to be a system characteristic, the specification was changed by the customer to permit a certain rate of anomalous occurrence. A residual part of the system anomaly problem continues to exist today, and this is believed to be a major contributor to the false alarm problem.

SYSTEM 2 EARLY EXPERIENCE

During the initial flight tests, a "false latch" problem was causing several false latches per flight (where the term "latch" refers to setting of a failure-indicating annunciator on a unit). BIT was deemed to be too sensitive, with tolerances overly tight--as tight or tighter than factory tests or intermediate level maintenance limits--and with BIT being overly sensitive to "one time fails" or "short duration faults." After fixes were incorporated, the false latches declined to less than one per flight.

SYSTEM 3 EARLY EXPERIENCE

This system represents the most modern and sophisticated of the three systems. The BIT designers are well aware of the problem of "anomalous performance" and have taken design measures to minimize the impact of such events. In further recognition of this characteristic, during the reliability demonstration test, from which our data base is derived, certain rules were established which precluded random happenings from being classified as relevant failures. Two systems were involved, with 723 hours of BIT-monitored hours for the first and 600 hours for the second. During this time, BIT generated 2352 and 1128 fault indications, respectively. This computes to be 3.5 and 1.9 fault reports per operating hour. (The second system had some improvements not incorporated in the first one.) It should be noted that the fault messages represented a great number of duplications of the same small set of faults. For example, the most commonly occurring fault numbers occurred 570, 486, 296 and 144 times on one system and 222, 379, 116 and 65 times on the other. The vast majority of these occurrences were invalid fault messages caused by such things as interference from external RF radiation. For these messages, BIT was correctly identifying anomalous system performance, but such performance was not indicative of the presence of faults.

SUMMARY OF EARLY EXPERIENCE

Although different phraseology was used, all three systems were, in effect, faced with the same type problem: Some BIT indications were generally not indicative of faults, i.e.,

each system was faced with a false alarm problem. The predominant factor was that BIT was detecting some form of anomalous system performance but such performance was not a manifestation of a fault. Generalizing, we can say that there is a high probability that any new system will be faced with a false alarm problem, although this problem is likely to be described in some other manner ("false latches," "troubles," "system anomalies," "invalid fault messages," etc.). Some relief is achieved by desensitizing the BIT tests (i.e., broadening the test tolerances) but there are definite limits to this approach. Possibly as a matter of coincidence, the rate of occurrence of false alarms (or apparent false alarms) during this early period was approximately two per operating hour for all three systems. It is our judgment that most of these false alarm events can be categorized as Category II, that is they are failure indications when in fact the systems are fault free (or at least fault free in the functional area being faulted).

For systems 1 and 2 it may appear that the same problem--the false alarm problem--has existed from the first days of operational service until the present time (a period of more than 6 years). In some respects, this is true. However, since many of the early deficiencies have been corrected, it must be assumed that the makeup of the problem during the early period is quite different than during later stages. At the start of the operational life cycle, there are many fundamental issues of basic performance. Either the systems don't do what they are supposed to or else there is incorrect information (possibly via faulty specification) as to what the system is supposed to be able to do--and also basic problems with BIT mechanization. In time, most of these basic issues are resolved. For example, by trial and error, BIT test tolerances will be gradually optimized. Also, obvious test defects will be discovered and corrected. The false alarm problem may continue, but for a different set of reasons. This study is primarily concerned with investigating the root causes of the false alarm problem in mature systems.

SUMMARY OF CURRENT FALSE ALARM EXPERIENCE FOR SYSTEMS 1 AND 2

Using the method described in section 4, false alarm rates have been computed for systems 1 and 2 and are summarized in Table 5-1, with a more detailed breakdown presented as Tables 5-2 and 5-3. In these tables, Category II false alarm rate is defined as the percentage of the total number of fault indications which have been classified as Category II false alarms (no fault), and Category I false alarm rate is defined as the percentage of valid fault indications (valid in the sense that there is a real fault in the system) which have been classified as Category I false alarms (wrong unit called out).

	CATEGORY I FALSE ALARM RATE	CATEGORY II FALSE ALARM RATE
System 1	28%	53%
System 2	38%	22%

TABLE 5-1 SUMMARY OF FALSE ALARM RATES

5.2 CAUSES OF FALSE ALARMS

5.2.1 CATEGORY I FALSE ALARMS

In discussing Category I false alarms with BIT design engineers, it became clear that this problem is a natural fallout of severe hardware and software constraints placed on BIT designs. In both systems 1 and 2, the original BIT design included fault isolation features which were subsequently dropped or scaled back. For example, in system 1, the single most severe isolation problem occurs when the displays are disabled. (The problem can be either a fault in one of the computer units, in one of the display units or in the interconnecting wiring.) Original plans to include a fault indicator on each of the computer units were abandoned in order to cut costs. In the case of system 2, a weight saving effort resulted in elimination of considerable BIT hardware, with the avionic hardware dedicated to BIT being reduced to less than 2%. This was achieved with no loss in system test capability, since it is still possible to test system response to a test signal inserted at the front end of the system. However, considerable loss resulted in the area of fault isolation. Another important constraint was the amount of computer capability/memory that could be dedicated to BIT. The lessons have been well learned and both systems are now in the process of incorporating expanded memories. For system 2, this will more than double the software capacity for use by BIT. In the case of system 1, the expanded memory will permit an independent self test of the computer units, with two indicators mounted directly on the face of one of the computer units (one to indicate a computer failure and the other to indicate a failure of the computer power supply unit). Much of the fault isolation ambiguity will thus be eliminated.

The above discussion leads to the conclusion that the Category I false alarm problem is affected by program policy decisions as well as by technical problems. These decisions may have appeared to be appropriate at the time, but in hindsight it is possible to say that such decisions

TYPE FALSE ALARM	NUMBER OF FALSE ALARM OCCURRENCES														
	ONLY MAINTENANCE EVENTS INVOLVING SEQ. 3 BIT*						EVENTS WERE EIT WAS UTILIZED BY MAINTENANCE PERSONNEL ONLY (NO IN-FLIGHT BIT)**						ALL MAINTENANCE EVENTS		
	CLUSTERS	PAIRS	SINGLES	CLUSTERS	PAIRS	SINGLES	CLUSTERS	PAIRS	SINGLES	CLUSTERS	PAIRS	SINGLES	TOT		
CAT I FALSE ALARM (REAL FAULT-INCORRECT ISOLATION) MULT. REMOVAL	91	33	-	3	6	-	94	39	0	133					
CAT II FALSE ALARM (PURE FALSE ALARM - NO FAULT) NO REMOVAL	99	79	365	-	-	-	99	79	365	543					
VALID FAULT IND. (NON-FALSE ALARM) REMOVAL & NO MORE FAILS	66	192	-	-	-	89	66	192	89	347					
SUBTOTALS	256	304	365	3	6	89	259	310	454	1023					
TOTALS	925			98			1023								

* SEQ. 3 BIT IS THE RADAR SUBSYSTEM BIT.

NUMBER OF TRUE FAILS = 133 + 347 = 480

** This category is based on Ground Rule No. 5 which assumes that BIT is used by maintenance personnel in two-thirds of the instances for which there is no record of in-flight BIT use.

CAT I FALSE ALARM RATE = 543/1023 = 53%

CAT I FALSE ALARM RATE = 133/480 = 28%

TABLE 5-2. TABULATED DATA FOR COMPUTING FALSE ALARM RATES--SYSTEM 1

TYPE FALSE ALARM	NUMBER OF FALSE ALARM OCCURRENCES											
	ONLY MAINTENANCE EVENTS INVOLVING PILOT BIT (PBIT)				EVENTS WHERE BIT WAS UTILIZED BY MAINTENANCE PERSONNEL ONLY (NO PILOT BIT) *				ALL MAINTENANCE EVENTS			
	CLUSTERS	PAIRS	SINGLES	CLUSTERS	PAIRS	SINGLES	CLUSTERS	PAIRS	SINGLES	TOT		
CAT I FALSE ALARM (REAL FAULT-INCORRECT ISOLATION) MULT. REMOVAL	120	85	0	17	49	0	137	134	0	271		
CAT II FALSE ALARM (PURE FALSE ALARM - NO FAULT) NO REMOVAL	19	19	161	-	-	-	19	19	161	199		
VALID FAULT IND. (NON-FALSE ALARM) REMOVAL & NO MORE FAILS	12	30	116	-	-	287	12	30	403	445		
SUBTOTALS	151	134	277	17	49	287	168	183	564	915		
TOTALS	562			353			915					

* This category is based on Ground Rule No. 5 which assumes that BIT is used by maintenance personnel in two-thirds of the instances for which there is no record of in-flight BIT use.

CAT II FALSE ALARM RATE = 199/915 = 22%

NUMBER OF TRUE FAULTS = 271 + 445 = 716

CAT I FALSE ALARM RATE = 271/716 = 38%

TABLE 5-3. TABULATED DATA FOR COMPUTING FALSE ALARM RATES--SYSTEM 2

were not based on a full appreciation of the impact of the support task on combat readiness. It is finally being recognized that slighting the BIT task on complex systems is not cost effective.

5.2.2 CATEGORY II FALSE ALARMS

Fundamental to understanding the root cause of Category II false alarms is the fact that such events are characterized by being inconsistent or intermittent. They behave like random variables. But this same characteristic is also descriptive of true intermittent faults. It will be recalled that the analysis methodology was developed with full cognizance of the need to sort out the false alarms from true intermittent faults. Figure 5-1 is presented as a reinforcement of the idea that to the maintenance person symptoms of Category II false alarms and those of true intermittent faults are identical. This is another way of saying that fault-free equipment periodically exhibits short intervals of "failure-like" performance. This is not just a maintenance phenomenon, but also a very real phenomenon to the radar operator. Table 5-4 illustrates this point by listing the makeup of the type squawks generated during the period the data base for system 2 was being compiled. Note the large number of squawks that are described in such general terms as "breaks lock" or "scan abnormal." These are very real system "failures" to the operator, but the big majority are not failures in a maintenance sense, that is, they are not associated with broken or failed parts. This problem is particularly insidious because if the pilot has squawked a fault-free system, the maintenance person is obligated to run BIT on the ground and every time that BIT is run, there is some probability that a false alarm will be generated. Thus, momentary "non-fault" anomalies which have been detected via operator observation can lead to BIT false alarms and, subsequently, unnecessary maintenance action.

The makeup of the problem of momentary anomalous system performance is summarized below.

1. Variability of functional performance of fault-free equipment due to natural, external phenomena associated with radar operation, such as varying, target radar-cross-section, ground reflections, doppler effects. (Symptom generally observable by operator, but not by BIT).
 - o False Targets
 - o Detection Problems
 - o Multiple Detections
 - o Lock-on Problems, Break-lock Problems.

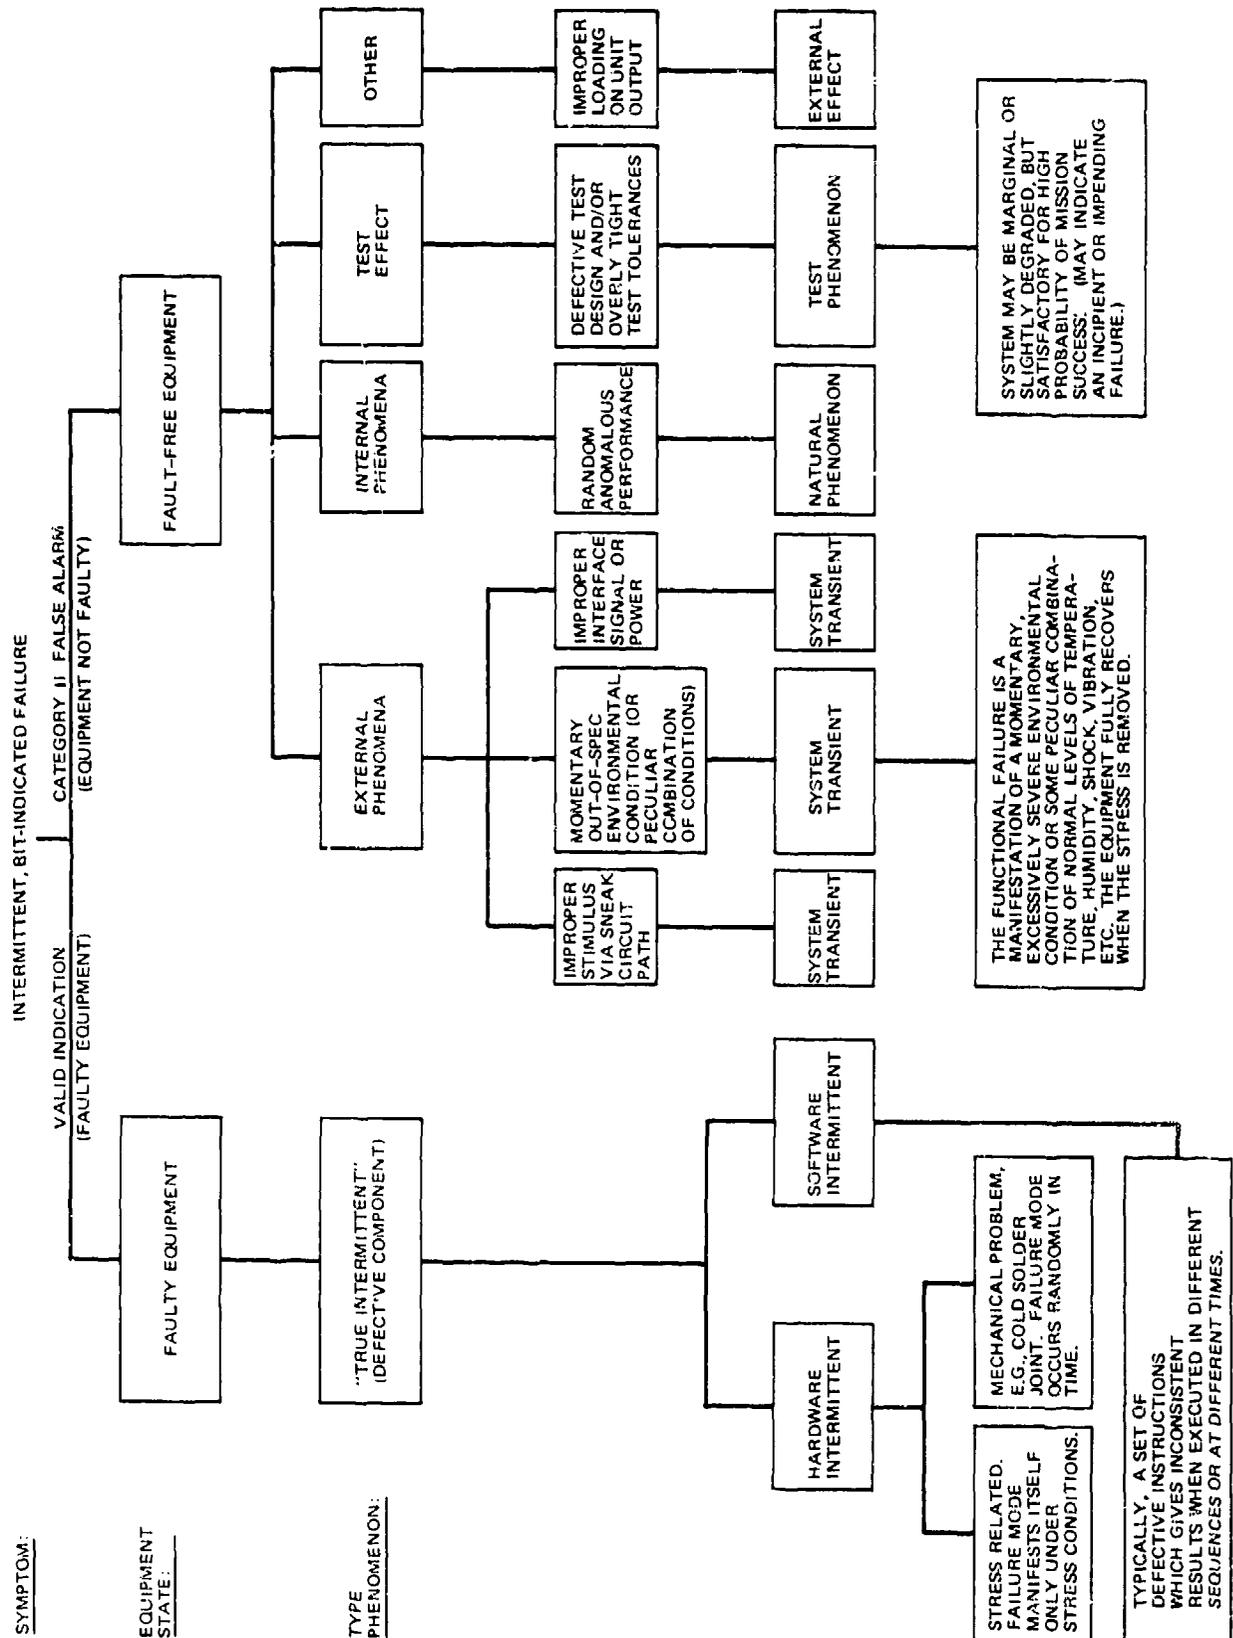


Figure 5-1. Makeup of Intermittent/Inconsistent/Non-repeatable BIT Indications.

RADAR FLIGHTLINE MAINTENANCE REPORT

AIRCRAFT/FLIGHTLINE	LOCATION	SPECIAL ACTIVITY	COMPLETION DATE		
ACN	A/C TN	FLT PILOT	RELATED ACN		
1. SQUAWK					
A.	B.	C.	D.		
E.	F.				
1. RFNG 179	POWER UP/TIME OUT 56	NO TGTS 113	BIRDS/FALSE TGTS 366	BREAKS LOCK 315	SCAN ABNORMAL 229
2. TD BOX ABNORMAL 73	POOR DETECTION SENSITIVITY 81	RANGE ERROR 30	LOCKS ON TO BIRDS/GROUND 16	POFF CIRCUIT BREAKER 0	ANT BANGS STOPS 3
3. ANT OSCILLATES IN TRACK 18	NO/SLOW LOCK ON (B'S OR SS) 65	NO/SLOW MANUAL LOCK ON (LRS) 108	HOJ AOJ JAM 112	VSD FLASHED WENT BLANK 41	NO/SLOW MANUAL LOCK ON (SRS) 106
4. LATCHED 122	ASP 2 SET 956	ASP 25 (WOP) SET 29	ASP 24 (LCP) SET 68	BIT DETECTED	NO TRACK TEST TGT 93
5. HYDRAULIC LEAK 22	INOP 28	INOP AFTER TAKEOFF 60	GRID LINES 10	GOES TO MEMORY 5	341

TYPE SQUAWK

NUMBER OF OCCURRENCES

NON-BIT DETECTED

2097, TOTAL

INOPERATIVE

185

1B. POWER UP/TIME UP	56
3E. VSD FLASHED/WENT BLANK	41
5B. INOP	28
5C. INOP AFTER TAKEOFF	60

TRACKING

675

1E. BREAKS LOCK	315
2B. POOR DETECTION SENSITIVITY	81
3B. NO/SLOW LOCK ON	65
3C. NO/SLOW MANUAL LOCK ON	108
3F. NO/SLOW MANUAL LOCK ON(SRS)	106

SCAN

323

1F. SCAN ABNORMAL	229
2A. TD BOX ABNORMAL	73
2F. ANT BANGS STOPS	3
3A. ANT OSCILLATES IN TRACK	18

BIRDS/JAM

494

1D. BIRDS/FALSE TARGETS	366
2D. LOCKS ON TO BIRDS/GROUND	16
3D. HOJ/AOJ/JAM	112

NO TARGETS

211

1C. NO TARGETS	113
4F. NO TRACK TEST TARGET	93
5E. GOES TO MEMORY	5

RF NO-GO

179

1A. RF NO-GO	179
--------------	-----

RANGE ERROR

30

2C. RANGE ERROR	30
-----------------	----

TABLE 5-4. SYSTEM 2 RADAR FAILURES SQUAWKED BY PILOTS

2. Variability of functional performance of fault-free equipment due to internal phenomena associated with system operation. (Generally observable by BIT, but not by operator). These are functional failures which are not manifestations of faults. They are maintenance false alarms!
 - o Transmitter dumps, power dumps, computer hangups
 - o Interferences from external sources. (RFI)
 - o Stress due to environmental factors. (Temperature, etc.)
 - o Interference from internal sources. (Power transients)
 - o Momentarily improper interface signals
 - o System noise
 - o Random anomalous performance
 - o Design problem, such as a sneak circuit path. (Functional failure, but not "equipment failure".)
3. Variability of functional performance of faulty equipment, caused by the variability inherent to non-catastrophic failure modes. (Generally observable by both BIT and operator).
 - o Intermittent fault, random occurrence in time
 - o Intermittent fault, occurs under certain stresses or combination of stresses
 - o Connector problems (more generally the "connection" problem)
 - o Random, "one-shot" failures
 - o Soft failures
 - o Degradation of equipment. ("incipient failure")
 - o Ground-peculiar failures.
4. Variability of functional performance of fault-free equipment due to errors in the software, i.e., due to software "unreliability." (Generally not observable by either BIT or operator)
 - o Improper logic design

- o Improper implementation
- o Timing problem

The system 1 data base was particularly useful in deriving root causes of Category II false alarms. This results from the fact that the DPs (test numbers) recorded on the BER (BIT evaluation report) card can be translated into system functional failures from DP descriptions available in both O-level and I-level publications. By observing the DP patterns on the subject BER cards it was possible to project possible root causes for 70% of the CAT II false alarms. Table 5-5 is a list of the possible causes of CAT II false alarms and Table 5-6 provides the number of occurrences of CAT II false alarms subdivided by root cause category (and also by cluster designation). The categories are somewhat arbitrary but are tailored to the types of data available for analysis. For instance category 3 (High Voltage/Transmitter Anomaly) could be divided among the transient failures, hardware failures, and environmentally induced failures, but the failures are more readily identified (and correctable) simply as high voltage/transmitter anomalies. A brief description of each of the categories follows.

Invalid Test

These are tests which have been improperly mechanized, such as incorrect logic, timing, or stimulus, or tests in which the test tolerances are excessively tight or use an incorrect nominal value. Due to the extensive time that system 1 has been in the fleet, practically all improperly mechanized tests have been detected and corrected. Only one DP (DP 216) in sequence 3 could be identified as invalid. This DP ranked number one in frequency of occurrence. There may be other DPs that occur only with certain combinations of WRAs installed because of tolerance buildup. If the system is still functional, however, the test tolerances should be widened and no DP displayed. Other than DP 216, no DPs were associated with this category during this analysis although there are some possible candidates.

Power Transient

Power transients can cause indicated BIT failures in two ways. The first is a transient that occurs during a critical measurement and causes the signal being monitored to fall outside test limits. The second is the transient that causes the related power supply to "crow-bar" or shut down. The affected unit will then be non-functional until power is recycled by the system operator. There are 19 WRAs in system 1 with power fault indicators monitored by the central computer.

TABLE 5-5 POSSIBLE CAUSES OF SYSTEM 1
CATEGORY II FALSE ALARMS

1. Invalid Test
 - a. Test mechanization incorrect
 - b. Test tolerance incorrect
2. Low Voltage Power Transient
 - a. Momentary
 - b. Unit Power Shutdown
3. High Voltage/Transmitter Anomaly
 - a. Electrical anomalies (arcs, etc)
 - b. Environmental anomalies (oil cooling, waveguide pressure, etc)
4. System Anomalies
 - a. Transitory hardware phenomena
 - b. Interface problems
5. Environmentally Induced Failure
 - a. Altitude
 - b. Temperature
 - c. Vibration
 - d. Acceleration
 - e. Humidity
 - f. RFI
6. Operator Switchology
 - a. Incompleted actions
 - b. Incorrect actions

TABLE 5-6. NUMBER OF OCCURRENCES OF SYSTEM I FALSE ALARMS, BY ROOT CAUSE AND BY CLUSTER DESIGNATION

SUSPECTED CAUSE	NUMBER OF OCCURRENCES OF CAT II FALSE ALARMS					
	CLUSTERS	PAIRS	SINGLES	TOTALS	PERCENT	
INVALID TEST	13	16	52	81	15	
POWER TRANSIENT (Power Fault)	0	6	58	64	12	
HIGH VOLTAGE/XMTR (XMTR Dump)	2	12	25	39	7	
SYSTEM ANOMALIES	64	24	51	139	26	
ENVIRONMENT	3	1	41	45	8	
OPERATOR ACTION	0	0	10	10	2	
UNDETERMINED	17	20	128	165	30	
TOTALS	99	79	365	543	100	

A failure indication for any of these units activates the same DP (DP 180). The power fault code is decoded by the central computer and the appropriate WRA is displayed but the use of a single DP for all power faults results in a high incidence of this particular DP.

High Voltage/Transmitter Anomaly

The system 1 transmitter is a high power rf source and requires two high voltage power supplies of 11 and 18 KV dc. Due to the high levels of power and voltages, the transmitter subsystem seems to be more susceptible to environmental anomalies than the low voltage units. The transmitter protection circuitry will shut down the transmitter whenever an anomaly is detected. If this occurs during the performance of the BIT procedure, DPs will be displayed. Recycling of the transmitter by the operator may restore normal operation. The BIT failures will, in this case, be scored as false alarms.

System Anomalies

This subset of CAT II false alarms includes those system anomalies which appear to be momentary, anomalous performance of fault-free equipment but, on the basis of DP analysis, may be explainable in terms of transitory hardware phenomena (e.g., incipient and intermittent faults) and interface problems.

It should be recalled that fault indications that disappear with no unit removal (i.e., do not occur in subsequent BIT runs) are classified as CAT II false alarms. This convention was adopted simply as a matter of practicality in analyzing the data. The intent is to flag out those random happenings which are "pure" false alarms, i.e., momentary, anomalous performance of fault-free equipment. Although these events must be written off as false alarms, in fact, it is reasonable to expect that each of these events has some rational (though hidden) explanation. Many of the explanations are based on assumptions of transitory hardware phenomena.

The term "transitory hardware phenomena" is intended to encompass three types of problems:

1. Intermittent faults where the repetition rate is so low that they escape detection by the filtering technique in use here (i.e., they don't occur during two missions subsequent to the original cluster of failure indications, but then they do occur in later missions).
2. Incipient faults which have deteriorated to the point of being close to the borderline between acceptable and unacceptable performance, such that any momentary system perturbation can cause the failure to exhibit

unacceptable performance for the duration of the perturbation.

3. Temporary, borderline performance not representing a deteriorating condition but simply representing normal performance for the existing equipment configuration.

As an example of transitory hardware phenomena possibly being the root cause of apparent false alarms, we might cite one particular test for which there were 15 cases of CAT II false alarms (DP 146, Continuous Wave Illumination test). Failing this test means that the Continuous Wave Transmitter has failed to turn off. A momentary failure of a particular relay can cause the observed momentary failure of the test. On the basis of historical experience with the particular type relay used, it can be said that the relay performance is often suspect. Relay experience includes all of the transitory type problems mentioned. In some cases, intermittent operation has been caused by loose particles becoming engaged in the contacts. In others, incipient failures were encountered as a result of gradual buildup of contamination on the contacts. Additionally, even fault-free relays have had a "dry current" problem when the equipment had not been used for some time (high contact resistance until a high current is passed through the contacts). These problems have now been generally eliminated but were present when our data were being collected.

All of the transitory hardware phenomena described above can exist outside of the unit under test as well as within it. Where outside performance affects the performance of the tested unit, an interface problem exists. Of particular importance are inter-unit wiring problems and connector failures. From an analysis point of view, these problems are doubly troublesome because (1) all transitory problems are difficult to analyze and (2) corrective action taken elsewhere in the system may not be associated with disappearance of a problem in the unit under test. Sometimes a maintenance action may inadvertently correct a problem. For example, dirty contacts may be cleaned in the act of removing a unit. It is inherently difficult from later analysis of collected data to know what happened.

An interesting example of an interface problem is associated with the BIT target test (DP 54). The airframe manufacturer has installed a particular type of coaxial cable with an especially low loss characteristic for carrying the BIT horn target signal from the BIT target generator. Although this was initially considered highly desirable, in fact it created problems since the target level adjustment in the unit had not been designed to accommodate such low level signals. For cases where the adjustment of the signal level is marginal, intermittent failures of the DP will occur.

Another type of interface problem is compatibility with test circuits. By experience it sometimes comes to be recognized that fault-free equipment will not always pass certain tests. For example, it was found that the synchronizer would not always lock up at zero range in the BIT mode every time the zero range lock test (DP 141) was run. Tests of this type need to be recognized as contributing a residual false alarm rate, and maintenance personnel need to be aware of the problem so that inappropriate actions are not taken.

Environmentally Induced Failure

BIT false alarms attributed to environmental factors are primarily tests involving acquisition and track of x-band BIT targets. Antenna angle track tests are performed using an x-band target located in the front of the radome. The receiver shutter is open during this test and thus the receiver is susceptible to external radiation. If the aircraft is in motion, the antenna array is susceptible to g forces. Some relay malfunctions have been attributed to landing shock but none of the DP patterns analyzed was associated with this. Although wet computer boxes were frequently reported and humidity is a contributor to intermittents, these failures are usually of a nature that prevent BIT from running (or cause random DPs), thus none of the analyzed BIT failure patterns were attributed to moisture. System overheat is indicated by a cockpit light. DPs accompanied by this condition light were included in the environmental category.

Operator Switchology

Certain BIT failure indications can be caused by incorrect or missing switch settings. Most of the switch settings required by BIT can be monitored by the central computer. If one of these switch actions is not performed by the operator, a mnemonic is placed on the BIT display to indicate the required switch action. However, certain aircraft switches that are not monitored by the computer can cause BIT to fail. For example, if the ground cooling switch is not in the radar position, a transmitter interlock is opened and BIT DPs will be displayed. The DP pattern is always the same for this switch and is recognized by experienced operators. For data analyzed from Miramar Naval Air Station, only 2% of the false alarms were attributed to switchology, so this is not a major contributor to false alarms.

Undetermined

This category encompasses those items for which there is insufficient data to speculate on root causes.

5.2.3 PROBABLE CAUSES OF OCCURRENCE OF SPECIFIC DPs

The analysis contained in the previous section has been based on BIT WRA callouts. BIT also identifies specific tests which have failed, by DP numerics. The top 20 DPs, as determined by frequency of occurrence, have been analyzed and the predominant cause of the high rate of occurrence for each DP has been determined. Table 5-7 lists the probable causes.

It is of interest to note that when this system first entered flight testing, the ratio of software-to-hardware problems was relatively high. With software maturity, this ratio has become very low. Nevertheless, every time a new software package is added, there is a risk of introducing a new software problem. An example of this is the fact that the number 1 BIT false alarm problem (BIT DP 216) was introduced with the introduction of the latest software modification.

TABLE 5-7. 20 TOP SYSTEM 1 DP FAILURES RANKED BY NUMBER OF OCCURENCES

<u>RANK</u>	<u>SEQ 3 DP</u>	<u>OCCUR- ENCES</u>	<u>TEST</u>	<u>PROBABLE CAUSE</u>
1	216	223	Frequency processing failed using VTPE and TDRF	TM
2	141	167	Failed ROT on transmitter in PC mode	BH
3	153	142	Transmitter failed	H
4	180	138	Power fault	H
5	176	129	CWI power fail	DM
6	149	128	Antenna not tracking horn target during PDSTT	E/H
7	62	115	Antenna scan failure in +/- 65° mode	DM
8	175	107	Transmitter not on for LPRF test	H
9	164	99	Antenna not tracking horn target during PSTT	E/H
10	177	98	Transmitter flood antenna switch not enabled	DM
11	198	92	Dummy load indicated with flood horn selected	DM

TABLE 5-7 (continued). 20 TOP SYSTEM 1 DP FAILURES RANKED BY NUMBER OF OCCURRENCES

<u>RANK</u>	<u>SEQ 3 DP</u>	<u>OCCUR- ENCES</u>	<u>TEST</u>	<u>PROBABLE CAUSE</u>
12	73	87	Low false alarm rate with low external threshold	TM
13	54	86	HPRF BIT target no. 4 level incorrect	TT
14	146	83	CWI failed to turn off	DM
15	136	71	High LJET false alarm rate	DM
16	199	70	A/G lobing failed	DM
17	170	65	Failed to generate ACM LAOT on horn target	E
18	191	65	BIT Log DC out of tolerance	TT
19	99	64	MLC notch fails to take out target	DM
20	169	64	ACM threshold calibrate fail	DM

PROBABLE CAUSE LEGEND

BH BIT HARDWARE MARGINAL
DM DEFERRED MAINTENANCE
E ENVIRONMENTAL INFLUENCE
H HIGH ANOMALY RATE HARDWARE
TM BIT TEST MECHANIZATION INCORRECT
TT BIT TEST TOLERANCE INCORRECT

5.3 FALSE ALARM RATE PREDICTION FACTORS

Considerable thought has been given to the feasibility of developing a mathematical model for predicting false alarm rates of new systems. Conceptually, to use such a model, one would simply select the appropriate system type (avionic, ground, tank, submarine, etc.) and then insert an appropriate set of coefficients. Each coefficient would be associated with a system characteristic known to be related to false alarm generation. The magnitude of each coefficient would be a measure of the likelihood of generating such false alarms for the particular system being investigated. The model would provide a predicted false alarm rate and perhaps a means for determining where resources should be expended to reduce the false alarm rate. Such a model would be based on characteristics of presently existing systems. The accuracy of the model would be dependent upon the size of the population of present systems investigated.

While the idea of a predictive model is extremely attractive, it should be recognized that there are reasons why any attempt to predict the absolute false alarm rate of an equipment planned for development is likely to prove fruitless. Two such reasons are:

- o The causes of false alarms (and apparent false alarms) are diverse and unique, and past experience is not necessarily a good indicator of future experience.
- o The measurement of false alarm rate is very complex and difficult.

The first of these two reasons is related to the fact that false alarms represent events which should not be happening. When such indications do turn up in a new system under development, no one knows why. Is it because the BIT test mechanization is wrong? Is it because the testing tolerances are too tight? Is it because the equipment does not meet its specifications? The causal factors (answers to these kinds of questions) must be determined for each type of potential false alarm--and there may be hundreds. Finding the answers and deciding what to do next may require years of engineering work to resolve. And while this work is going on, it becomes pure "busy work" to tabulate how many false alarms are occurring. Invariably, the same statistics will be repeated over and over again, until the individual causes are identified and corrected.

With respect to the statement that measurement of a BIT false alarm rate is complex and difficult, the results

of the present investigation make this seem almost a truism. Even for systems that have been operational for years--and after years of engineering work in sorting out and eliminating most spurious BIT indications--the problem of deciding whether certain of the remaining indications are actually false alarms still remains. Although the number of spurious indications has been quantitatively reduced, the general nature of the problem to be dealt with is, in many ways, qualitatively the same as when the BIT/system engineering process first began. Since it is very hard to measure the BIT false alarm rate, it makes little sense to attempt a quantitative prediction, in absolute terms, of what this rate is likely to be for a new system. It might be noted that this same line of reasoning indicates the futility of putting stringent false alarm requirements in specifications.

Considerations such as those described in the preceding paragraphs have led us to believe that, in attempting to predict BIT false alarm rates at the early conceptual stage of system development, the emphasis should be on predicting relative, rather than absolute, false alarm rates. By having a procedure for predicting relative false alarm rates, it is possible to conduct design trade studies and in this way choose design approaches to minimize BIT false alarms. The study has resulted in the development of factors for relative false alarm rate prediction. These are applicable to the early design phase of a BIT system. Although refinement will be required in applying these factors, the technique, to be described, illustrates a methodology for evaluating alternative BIT designs to determine their susceptibility to false alarms.

It is considered desirable to treat BIT false alarm prediction factors separately for Category I false alarms (real fault--incorrect isolation) and for Category II false alarms ("pure" false alarm--no real fault). The CAT II type false alarm is the more insidious of the two since it falsely indicates the need for maintenance and is more likely to cause unnecessary mission aborts. This type of false alarm will be dealt with first. Only the most basic factors will be addressed in this report.

Prediction Factors Related to CAT II False Alarms

The objective in this case is to develop a false alarm index which can be used in conducting trade studies among several possible BIT designs for new equipment to be developed. It is assumed that any signal being measured by BIT is evaluated with "reasonable" tolerances and these same tolerances are applied to all of the BIT design approaches being considered. Four factors of most importance are listed below.

1. Number of signals/parameters BIT evaluates.
2. Number of times BIT is performed.
3. Operating environment factor.
4. Filtering effectiveness.

1. Number of Signals/Parameters BIT Evaluates

Given prime equipment of any complexity, the number of BIT false alarms should be positively correlated with, if not directly proportional to, the number of signals/parameters that BIT checks. Thus, if the number of parameters being checked is doubled, it seems reasonable to expect that the false alarm rate might also be doubled. (NOTE: Specifications demanding a high degree of BIT "thoroughness" tend to cause more tests to be designed into the system and, as a direct result, tend to result in more false alarms.)

2. Number of Times BIT Is Performed

By definition, each time BIT yields a fault indication for a fault-free system, it is a false alarm. Clearly, the more times BIT is performed, the more opportunity there is for a false alarm to occur. A basic consideration here is that a false alarm can only occur if the operator or maintenance person sees the erroneous indication. This is clarified in the discussion of factor 4, below.

3. Operating Environment Factor

In part, false alarms are generated as a function of the environment in which the equipment operates. For example, a severe operational environment can cause momentarily severe temperature excursions, which in turn can cause fault-free equipment to momentarily malfunction (which can be misinterpreted by BIT as an indication of a catastrophic fault). The operating environment factor can be thought of as a kind of index which will have the same value for all of the BIT design alternatives being compared. A nominal value for the index might be 1.0, which would represent an average set of operating conditions.

4. Filtering Effectiveness Factor

The number of potential Category II false alarms is a function of the number of different signals which BIT evaluates and the number of times each of these signals is tested in a given time period. If the system which BIT is checking is complex and if BIT is mechanized to provide a high degree of thoroughness in checking system functions,

there will be a substantial probability of triggering a false alarm each time the signals are tested. In addition, when continuous monitoring is applied to all of the tested signals, each signal may be tested many times per hour of system operation and the chances for false alarms are even greater. It is therefore essential that there be some method of "filtering" the potential false alarms so that the NO GO indications which BIT displays to the system operator have a reasonable chance of being valid.

This objective can be achieved by taking the following two actions in the way BIT is mechanized:

1. Mechanize the continuous monitoring function so that the individual BIT checks of a given signal are far enough apart in time to be uncorrelated (i.e., so that if the results of the "i"th check of a signal are within the range of normal signal excursions, this is not indicative of where within that range the signal may be on the "j"th check).
2. Establish criteria, using the results of successive BIT checks, for deciding when failed BIT results for a given signal are sufficiently consistent to merit displaying a NO GO to the system operator (and/or to the maintenance person).

The intent of these actions is, of course, to minimize the frequency with which spurious NO GOs are displayed to the operator. This will ensure that when NO GOs are displayed the air crew can be confident that a valid malfunction phenomenon is present and act accordingly. Similarly, maintenance personnel can know that when a BIT NO GO is present, maintenance is actually required.

The purpose of applying filtering to BIT testing is to censor out BIT test failures which probably represent normal performance of the equipment and do not require maintenance actions to be taken. If BIT testing limits are established symmetrically about the range of values which a tested analog signal usually assumes when performing normally, it is to be expected that the usual test result for that signal will be a PASS. However, if the BIT testing limits are relatively "tight" in relation to the actual behavior of the signal, some fraction of sampled signal values will fall outside the BIT testing limits and register a FAIL when the BIT procedure is performed.

One measure of the tightness of BIT testing limits is the standard deviation, designated by σ , of the values which the testing signal assumes during normal operation. Calculation of σ assumes a normal distribution and a large value of σ indicates more dispersion of the signal values than when the value of σ is small. But if all BIT testing limits are placed at the same

number of σ units above and below the average value of the normally performing signal, it is possible to estimate the fractions of the time that the signals exceed the BIT testing limits on either the high side or on the low side. For example, if BIT testing limits were placed at $\pm 2 \sigma$, it is to be expected that 2.27% of the signal values will fall outside the BIT testing limits on the high side and another 2.27% will fall outside the BIT testing limits on the low side. Thus with BIT testing limits set at $\pm 2 \sigma$, a normally performing signal can be expected to fail $2.27 + 2.27 = 4.5\%$ of the time that the signal is tested. If the BIT tolerances are loosened to $\pm 3 \sigma$, the FAIL rate for normally performing signals drops to 0.26% of the signals tested. And with the BIT tolerances loosened still further to $\pm 4 \sigma$, the FAIL rate drops to 0.006%. These values all assume that BIT measurement error is zero.

There are advantages and disadvantages to having BIT tolerances tight or loose. The advantage of tight tolerances is the high precision provided by a PASS result, i.e., the values of the signal are known within relatively narrow limits. The disadvantage is the relatively high false alarm rate; frequently a normally performing signal will be measured to be outside of the testing limits. These characteristics are reversed for loose tolerances; information about the signal provided by a PASS result is less precise (a disadvantage) but the false alarm rate will be lower (an advantage).

The actual conditions under which engineers establish BIT tolerances is somewhat less refined than the above discussion of σ limits might imply. For a complex electronic system, the actual behavior of certain tested signals under operating conditions may not be precisely known. In this case, it may be necessary to initially establish the BIT testing tolerances in accordance with what the behavior of the tested signal is supposed to be--the "specification values." This procedure will sometimes lead to unexpected test failures even though the testing tolerances are seemingly loose (e.g., $\pm 4 \sigma$). These failures occur because the actual mean and the actual σ of the signal under operational conditions are not accurately known.

In order to illustrate the filtering of BIT test results, a hypothetical BIT design will be assumed. This BIT design is one in which 100 different signals are tested. The hypothetical BIT system uses some form of continuous monitoring such that each of the 100 signals is tested 100 times per hour. For all 100 of the tested signals, BIT testing limits are set at $\pm 2 \sigma$. Assuming the use of $\pm 2 \sigma$ limits is entirely arbitrary; we could choose $\pm 3 \sigma$ or $\pm 4 \sigma$. But the use of $\pm 2 \sigma$ in this example helps to show that filtering can be effective in eliminating false alarms even when the BIT testing tolerances are quite tight.

Filtering of test results for the hypothetical BIT system is illustrated in Table 5-8. The method of filtering used in

TABLE 5-8. TECHNIQUE FOR ANALYZING FILTER EFFECTIVENESS FACTOR

The following table is based on a filtering concept of at least "m" test fails out of "n" opportunities before a NO GO indication is displayed. The table provides in rows (a) the probability of generating a NO GO indication and in rows (b) the average number of such indications per hour of system operation, for various combinations of m and n.

- Assumptions: (1) 100 different signals are each tested with + 2 σ test limits 100 times per hour.
 (2) The system under test is performing normally.
 (3) Successive individual tests are independent and criterion test groupings are mutually exclusive.

NO GO Display Criteria: Number of Fails, m		n, Number of Successive Tests on which GO/NO GO Indication is Based				
		1	2	3	4	5
1	(a)	4.54×10^{-2}	8.87×10^{-2}	1.30×10^{-1}	1.69×10^{-1}	2.07×10^{-1}
	(b)	454	444	433	422	414
2	(a)		2.06×10^{-3}	6.00×10^{-3}	1.16×10^{-2}	1.88×10^{-2}
	(b)		10	20	29	38
3	(a)			9.36×10^{-5}	3.61×10^{-4}	8.73×10^{-4}
	(b)			0.31	0.90	1.75
4	(a)				4.25×10^{-6}	2.05×10^{-5}
	(b)				0.01	0.04
5	(a)					1.93×10^{-7}
	(b)					0.0004

EXAMPLE: If the NO GO display criterion of at least 3 fails out of 4 tests is selected, the probability of a NO GO indication in a sequence of 4 tests is 3.61×10^{-4} . The average number of NO GO indications displayed to the operator per hour of system operation would be 0.90. Without filtering (m=n=1), 454 NO GOs would be generated. The filtering effectiveness factor is thus $0.90/454 = 0.00198$.

this example is one in which filtering is accomplished by examining successive test results for consistency. Test-to-test consistency of failed test results is less likely to be present for a system that is performing normally than for a system containing an actual malfunction. The data presented in Table 5-8 assumes that the system being tested is performing normally.

The (a) rows of Table 5-8 present the probabilities that a NO GO display will be generated, for varying filtering criteria. The first (a) row presents probabilities when the criterion for displaying a NO GO indication is at least 1 failed test ($m=1$). From left to right, the numbers present the probability of a test failure every time one of the 100 signals is tested by the hypothetical BIT system, the probability of at least one failure in two tests of the same signal, the probability of at least one failure in three tests, etc. The second (a) row presents probabilities when the criterion for displaying a NO GO indication is at least 2 failures ($m=2$). From left to right, the numbers present the probability of two failures in two tests, the probability of at least two failures in three tests, the probability of at least two failures in four tests, etc. Similarly, the third (a) row presents probabilities when the criterion is at least 3 failures ($m=3$), the fourth (a) row presents probabilities when the criterion is at least 4 failures ($m=4$) and the fifth (a) row presents probabilities when the criterion is at least 5 failures ($m=5$).

The probabilities described above are computed by straightforward application of probability statistics. Each time one of the signals is tested by the hypothetical BIT system, there is a probability of $p=0.9546$ that the test will pass (under the $+2\sigma$ testing limit assumption) and a probability of $q=0.0454$ that the test will fail--as long as the tested equipment continues to perform normally as we assume. Thus, the first entry in the first (a) row is 0.0454, or 4.54×10^{-2} . Computation of the other probabilities is facilitated by use of the binomial expression $(p+q)^n$, where n is the number of successive tests of the signal. For $n = 2$, the binomial expression expands to $p^2 + 2pq + q^2$ and the values of the successive terms are, respectively, the probability that both tests pass ($p^2 = 0.91126$), the probability that one test passes and one test fails ($2pq = 0.08668$), and the probability that both tests fail ($q^2 = 0.00206$). For $n = 2$, the probability of at least 1 failed test is 0.0887 ($= 0.08668 + 0.00206$), i.e., the second entry in the first (a) row.

For $n = 3$, the binomial expression expands to $p^3 + 3p^2q + 3pq^2 + q^3$ and the values of the successive terms are, respectively, the probability that all three tests pass (0.87), that two tests pass and one test fails (.124114), that one test passes and two tests fail (0.005903), and that all three tests fail (0.0000936). The (a) row values in the column labeled

3 in Table 5-8 are computed in the following manner: $0.124114 + 0.005903 + 0.0000936 = 0.130$, $0.005903 + 0.0000936 = 0.006$ and, directly, 0.0000936 . The binomial expression was applied similarly for values of $n=4$ and $n=5$.

All of the probabilities in the (a) rows of Table 5-8 pertain to individual "criterion events." If the criterion for displaying a NO GO indication to the system operator is that at least three out of four tests for a given signal have failed, the criterion event in that case is that four successive tests of the signal are made by BIT. As seen from Table 5-8, the applicable probability of a NO GO indication for the three-out-of-four criterion is 3.61×10^{-4} . The other probabilities given in the table have a similar significance.

The (b) rows of Table 5-8 list the average number of NO GO indications which an operator could expect to see during one hour of system operation. These average numbers are tied to both the false alarm probabilities given in the (a) rows and to the number of criterion events which will occur in one hour of system operation. This can be illustrated by using again the example in which the criterion for displaying a NO GO indication is that three out of four tests of the signal are failed. In this case, the size of a criterion event is four tests. Since the number of BIT tests performed is 100 per hour for each of the 100 signals evaluated, a total of $(100)(100) = 10,000$ tests is performed during each hour of operation. But for filtered test data, each criterion event involves two or more tests. This means that there are fewer than 10,000 criterion events per hour. With the example of three out of four tests as the NO GO display criterion, the number of criterion events is $10,000/4 = 2,500$. Since NO GO indications are being generated at the rate of 3.61×10^{-4} per criterion event, the average number of NO GO indications per hour for the three-out-of-four criterion is $(3.61 \times 10^{-4})(2,500) = 0.90$. Other values in the (b) rows are similarly computed. If the total number of tests performed is 10,000 and the number of successive tests on which a GO/NO GO indication is based is 2, the number of criterion events per hour is 5,000; when the number of successive tests is 3, there are 3,333 criterion events; for 4 successive tests, 2,500; and for 5 successive tests, 2,000. It would, of course, be possible to mechanize the NO GO display arrangement using overlapping criterion groups (e.g., first criterion event consists of tests, 1, 2, 3, and 4; second criterion event consists of tests 2, 3, 4, and 5, etc.) but the average number of false alarms per hour would then be greater than given in the (b) rows of Table 5-8.

With this much background on the quantitative aspects of Category II false alarm rate prediction, let us now consider how a BIT design tradeoff study might be conducted. Assume that two basic BIT designs are being compared. Design A is a conventional, operator-initiated BIT performed two times per hour,

checking 100 signals each of these times. BIT Design B uses continuous monitoring and checks 100 signals 100 times per hour. Design A uses no filtering; that is, each BIT test failed is displayed to the operator as a NO GO. With Design B, on the other hand, there is the possibility of mechanizing the design with any of several degrees of filtering. These assumptions allow the following false alarm evaluations to be made for the two designs:

	<u>No. of BIT Checks Per Hour</u>	<u>No. of Signals Checked</u>	<u>Average No. of Tests Failed($\pm 2\sigma$ Testing Limits)</u>	<u>Predicted No. Filtering of Operator NO GO Indications</u>	<u>Filtering Factor = NO GOs/ Test Fails</u>
<u>Design A:</u>	2	100	9.08	9.08	1.00
<u>Design B:</u>	100	100	454	(a) 1.75 (b) 0.90 (c) 0.04	3.85×10^{-3} 1.98×10^{-3} 8.81×10^{-5}

- (a) NO GO display criterion of at least 3 fails out of 5 times tested
- (b) NO GO display criterion of at least 3 fails out of 4 times tested
- (c) NO GO display criterion of at least 4 fails out of 5 times tested

Notice that the comparisons have ignored the "environmental factor" mentioned earlier; it is assumed that the environmental factor affects both designs in the same way and can be omitted from the comparison. The average number of BIT tests failed is the product of the probability that a single test is failed when performed and the number of times the test is performed. For example, we know from Table 5-8 that the probability of failing a single test when $\pm 2\sigma$ limits are used for a normally performing system is 4.54×10^{-2} . For Design A, 100 signals are checked two times each hour (200 total tests); for Design B, 100 signals are checked 100 times each hour (10,000 total tests). These test repetitions result in 9.08 test failures for System A and 454 test failures for System B. But since System B utilizes any of three degrees of filtering, the number of NO GO indications displayed is in all three instances less for System B than for System A. A filtering factor can be calculated by dividing the predicted number of operator NO GO indications by the number of failed BIT tests (e.g., 454 for system B).

For the degrees of filtering illustrated here, BIT Design B is clearly superior to BIT Design A from the Category II false alarm point of view. However, BIT Design B is not without penalties because it will require added cost and complexity to mechanize the functions of continuous monitoring, BIT data recording and filtering. Tradeoffs may even be

involved in the three degrees of filtering for BIT Design B since a high degree of filtering may possibly entail a longer lag time in making BIT results available to the operator than would be the case for a lower degree of filtering and may also require additional memory space.

Prediction Factors Related to Category I False Alarms

Prediction of Category I false alarms requires use of different factors than were applied in connection with prediction of Category II false alarms. The applicable factors for prediction of Category I false alarms are these:

1. Number of system elements to which faults are to be isolated by use of BIT.
2. Category II BIT false alarm rate index.
3. "Federated" BIT design factor.

These three factors and their use in connection with BIT design tradeoff studies are described in the paragraphs which follow.

1. Number of System Elements to Which Faults Are to be Isolated

With one LRU in a system, and neglecting cabling/connector problems, the probability of isolating a known fault to that LRU is 1.00. If fault isolation is a completely random process and if all boxes have an equal fault likelihood, the probability of correct fault isolation is 0.5 with two LRU's, 0.33 with three LRU's, etc. Corresponding probabilities of incorrect isolation are 0, 0.5, 0.67, etc.

2. CAT II False Alarm Rate Factor

While CAT I and CAT II false alarms can be regarded as being basically independent of each other, it is important to be cognizant of real world problems associated with measuring the two types of false alarms. From this viewpoint, the CAT II false alarm rate can have a subtle but significant impact on the apparent CAT I false alarm rate. For example, assume that BIT has detected two independent system anomalies and that one is a pure false alarm (CAT II) and the other is the result of a hard failure. Also assume that BIT displays two units, one for each of the apparent failure events. The first unit is actually fault free and the second one contains a fault. If the maintenance person removes the fault-free unit first, it will be concluded that a CAT I false alarm has been experienced, since there is a very real fault in the system but apparently BIT has isolated this fault to the

wrong unit. Thus, some CAT II false alarms will, in effect, be translated into apparent CAT I false alarms. The CAT II false alarm rate factor, to be used in generating the CAT I false alarm rate index, is an attempt to deal with this phenomenon.

3. "Federated" BIT Design Factor

BIT designs are sometimes keyed to the results of end-to-end (system-level) tests. The system-level parameters represent functions which the over-all system must perform to meet its requirements. When the result of such testing is a NO GO, isolation to the system element considered to be faulty is accomplished by using logic in conjunction with BIT testing results for parameters other than the particular one found to be failed. Under field conditions--particularly when multiple indications of possible failures within a system are present--the elements of the system to which the indicated faults are isolated in this way are sometimes not faulty. When this is the case, a Category I false alarm is present.

The cure for Category I false alarms is to use a "federated" BIT design---one in which most BIT tests, when failed, isolate directly to the failed element of the system. If it is desired to isolate faults to an LRU, BIT is designed to individually test each LRU of the system. When one of these LRU tests is found to be failed, it is known with a probability approaching 1.00 that the detected fault is within the tested LRU.

In practice, a federated BIT design cannot lead to perfect fault isolation (i.e., cannot lead to elimination of Category I false alarms). Possible reasons for imperfections in fault isolation include the following:

1. Even though all the LRUs of a system may be individually tested by BIT, it is still necessary to verify that certain system-level functions are performing as required. The adequate performance of these functions depends upon correct operation of more than one LRU, creating a possible fault isolation problem when the system-level parameters are found to be NO GO.
2. Some functions tested at the unit level may also be tested at the system level, creating the possibility of contradictions between unit-level and system-level testing.
3. It may not be feasible to test and continuously monitor certain aspects of LRU operation by use of either unit-level or system-level BIT checks.

These possible imperfections in federated BIT design must be taken into account quantitatively in conducting tradeoff studies to select an optimal BIT design for a new system. This can be done by allocating the total failure rate of each system LRU into the following categories:

1. $\lambda_{\bar{s}}$ = LRU failure rate which is BIT tested by unit-level testing only.
2. λ_s = LRU failure rate which is BIT tested by system-level testing only (more than a single LRU involved in each test).
3. $\lambda_{\bar{s},s}$ = LRU failure rate which is BIT tested by both unit-level and system-level testing.
4. $\lambda_{\bar{t}}$ = LRU failure rate which is untested by either unit-level or system-level testing.

For that fraction of the system failure rate which falls into the first of the above categories, BIT fault isolation to the LRU level will practically always be correct. For that fraction of the system failure rate which falls into the second category, BIT fault isolation effectiveness can be estimated on the basis of experienced fault isolation effectiveness for systems already operational. This procedure can also be applied in the case of the third category, although this category is not expected to be much of a problem. Since testing tolerances for unit-level tests are usually tighter than testing tolerances for system-level tests, a unit-level test can be expected to be generally more effective in fault identification than a system-level test. That is, when a parameter is borderline bad, the unit-level test is more likely than a system-level test to detect the condition. The system-level test may say GO while the unit-level test says NO GO, but it will seldom be the other way around. The fourth category will not enter into BIT fault isolation effectiveness although it may contribute to possible anomalies between operator-reported evaluations of system performance and BIT results.

Tradeoff studies for selection of an appropriate BIT design can be conducted for Category I false alarms in a fashion similar to that which has been shown for Category II false alarms. This is illustrated in Table 5-9 for two hypothetical BIT designs C and D. Explanations of the entries for Table 5-9 are given in the paragraphs which follow.

The first column of Table 5-9 is entitled, "Number-of-System-Elements Factor". It is postulated that ease of fault isolation tends to be inversely proportional to the number of system elements to which faults must be isolated. Systems C and D are each assumed to be comprised of five LRUs to which BIT

TABLE 5-9. Technique for Comparing Systems' Category I False Alarm Characteristics.

Number-of-System-Elements Factor *	Federated BIT Factor				Product	Cat. II False Alarm Factor (Valid Fraction of Displayed NO GOs)	Fault Isol. Effect. Factor	Normalized Cat. I False Alarm Avoidance Index
	Type of Fault Isolation Testing	Failure Rate Fraction Tested	Fault Isolation Effectiveness					
0.2	$\lambda \bar{S}$	0.5	1.0	0.50	0.83	0.7	0.1162	1.00
	λS	0.1	0.5	0.05				
	\bar{S}, S	0.4	0.7	0.28				
		1.0	System Factor:					
0.2	$\lambda \bar{S}$	0.7	1.0	0.70	0.89	0.8	0.1424	1.22
	λS	0.1	0.5	0.05				
	\bar{S}, S	0.2	0.7	0.14				
		1.0	System Factor:					

* As used here the number-of-system-elements factor is merely the reciprocal of the number of elements to which BIT must fault isolate. Hypothetical systems C and D are each assumed to have five such elements.

must be capable of isolating any faults identified. Since the number of LRUs is the same for both systems, each is assigned a Number-of-System-Elements Factor of 0.2, the reciprocal of 5.

The second column of the table, "Type of Fault Isolation Testing", divides the total failure rate of the systems into the previously defined categories of $\lambda_{\bar{g}}$, λ_s and $\lambda_{\bar{g},s}$. The fact that the $\lambda_{\bar{f}}$ symbol is not included in the table signifies either an assumption that the system's failure rate is 100% tested by BIT or an assumption that the portion of the system's failure rate not subject to testing by BIT is not included in the comparison between Systems C and D.

The "Failure Rate Fraction Tested" column allocates the failure rate of the system to the indicated BIT testing categories. In order to arrive at these allocations, it is necessary to know or be able to predict the failure rate of the system down to the level at which BIT testing is performed. It is also necessary to know or be able to predict how the various BIT checks will be mechanized and the fraction of the circuitry each BIT test will evaluate. For System C, the assumption is that half of the system failure rate is evaluated by unit-level BIT testing. For the BIT tests of this type, fault isolation will be excellent; for the other two categories, fault isolation will not be as effective. System D shows an advantage over System C in that 70% of the system's failure rate is subject to unit-level type BIT testing. System D achieves this advantage by having a lower fraction of its tests subject to both unit-level and system-level testing.

The "Fault Isolation Effectiveness" column provides figures of merit for each of the three categories of BIT testing. These figures of merit are the same for the two systems. A figure-of-merit value of 1.0 for unit-level testing implies an assumption that that method of BIT testing leads to perfect fault isolation effectiveness. BIT testing which involves a combination of unit-level and system-level testing is assumed to be next most effective with a figure of merit of 0.7. This assessment in relation to unit-level (only) testing assumes that when there is overlapping unit-level and system-level testing, BIT fault isolation information will sometimes be ambiguous. BIT fault isolation is assumed to be least effective in those instances in which there is system-level BIT testing with no unit-level testing as a backup (figure of merit of 0.5).

The column labeled "Product" includes values which are merely the product of the two preceding columns. This product is a way of summarizing in a single number the effect of failure rate fraction tested and fault isolation effectiveness. The three values are then added together for each system to obtain a system federated BIT factor. The maximum possible system factor value is 1.0 whereas System C has a value of 0.83 and System D a value of 0.89.

The results of this BIT false alarm investigation indicate that a substantial proportion of BIT NO GO indications can be classified as Category II false alarms. When BIT indications representing Category II false alarms are present along with valid indications of system faults, the problem of correct BIT fault isolation is made more complicated. This relationship between Category I and Category II false alarms is acknowledged in the column entitled "Cat. II False Alarm Factor". It is assumed that in BIT Design C, 70% of the displayed BIT NO GO indications are valid (i.e., do not represent Category II BIT false alarms). For BIT Design D, 80% of the displayed BIT NO GO indications are assumed to be valid.

The column of Table 5-9 entitled "Fault Isolation Effectiveness Factor" combines the three principal numerics given in the table. The two values given in this column are the products for Systems C and D respectively of the "Number-of-System-Elements Factor", the product value under the "Federated BIT Factor", and the "Category II False Alarm Factor." Since the value in this column is higher for System D than for System C, it is implied that the BIT system for System D is more effective in avoiding Category I false alarms than System C. The right hand column of the table provides normalized values for the values given in the preceding column. This is accomplished by assigning a value of 1.00 for the "Fault Isolation Effectiveness Factor" of 0.1162 for System C and a value of $0.1424/0.1162 = 1.22$ for System D.

The approach outlined here for Category I false alarm prediction is but a starting point. The method will allow rough tradeoff studies to be conducted among different BIT designs. There is, however, a need for further research to specify the steps in more detail, to provide better quantification of the factors and to refine the overall method. These actions are beyond the scope of this investigation.

Prediction of BIT False Alarms During System Development

The preceding discussions of false alarm prediction have pertained to the selection of a BIT design from two or more candidate designs. Once this selection is accomplished, there is a need to monitor BIT false alarm status during system development. This monitoring will pertain primarily to the status of Category II false alarms. The status of Category I false alarms cannot be readily evaluated until system development is well down stream and not until the various elements to which faults must be isolated are operating together as a complete system. But there is a need to monitor Category II false alarm status throughout development. In effect, the monitoring amounts to successive updatings of the original

Category II false alarm prediction made during design selection.

The original Category II false alarm prediction will ordinarily have been made before any of the individual tests of which BIT will be comprised have been planned or mechanized. In making a false alarm prediction at such an early stage, it is necessary to estimate the total number of individual tests to be used and to assume that all of these tests fit a standard pattern of implementation (e.g., use of testing limits set at $+2\sigma$). But in actual fact, each individual BIT test is unique as to its potential for generating Category II false alarms. This uniqueness is determined by such considerations as the following:

1. Is the signal to be tested analog or digital?
2. What are the requirements which the tested signal must satisfy?
3. How well does the signal perform in relation to its stated requirements?
4. How are values of the signal affected by the modes or conditions of system operation?
5. To what extent do values of the signal vary from system to system?
6. How is signal functioning affected by the evolving design of the system?

Answers to questions such as these are needed for every signal which BIT checks. Such answers can be obtained using either of two approaches:

1. Trial and error.
2. Direct collection of data pertaining to signal performance.

The trial and error approach consists of trial implementation of a given BIT design and given testing limits based on a priori assumptions as to the performance of the tested signal. If the collection of BIT data indicates that the initial BIT implementation is yielding too many NO GOs, the testing tolerances can be loosened accordingly. The other approach---direct collection of data pertaining to the signal---is concerned with actual values of the tested signal rather than with BIT results as such. Preferably, the data as to signal performance should be collected by use of an automatic recording procedure so that the amount of data assembled is large enough to be representative. With such data in hand, rational

decisions can be made as to how the particular BIT test should be mechanized and testing limits established. Further information about establishing BIT testing limits is given in Appendix B.

As the development of the BIT design proceeds and data relating to BIT performance, or expected performance, are accumulated, it will be possible to make a succession of predictions as to the Category II false alarm rate. These predictions will be similar to those described earlier for use in tradeoff studies for BIT design selection. The principal difference will be that, as the BIT design develops, information and data will become available as to individual BIT checks. Such information and data will increase the accuracy of prediction for subsequent BIT performance in military operational environments.

In order to make false alarm predictions on a signal-by-signal basis, it will be necessary to know, or estimate, the false alarm rate of each tested signal or class of tested signals. Instead of an across-the-board basis for establishing BIT testing limits---such as $+ 2\sigma$ ---individual BIT checks will have their own limits and these may differ from any average system-wide standard. Also, individual BIT checks may utilize individual methods of filtering out spurious NO GO indications. That is, one BIT check may utilize a criterion of two fails out of three tests performed, while another check utilizes a criterion of three fails out of four tests performed. In this way, actual performance of the overall BIT can be optimized in a way that would not be possible with a standard, across-the-board method of filtering.

There will be a need to combine the data from individual BIT checks for use in a system-wide prediction of Category II false alarms. This can be accomplished by assembling the data in accordance with the following format:

(A)	(B)	(C)	(D)	(E)	(F)
Test Type	No. of Tests of This Type	Unfiltered False Alarm Probability	Filtering Standard	Filtered False Alarm Probability	Weighted False Alarm Rate (B) x (E)

A brief explanation of how these headings might be utilized will make them more understandable:

- (A) Test Type. This heading can be used to group together tests which have common characteristics--same method of establishing pass-fail limits (e.g., $+ 3\sigma$), same degree of test filtering, and perhaps other features.

- (B) Number of Tests of This Type. (Self explanatory)
- (C) Unfiltered False Alarm Probability. This is the probability that a given BIT test will yield a false alarm when the test is performed one time for a system which is performing normally (is not malfunctioning). It may prove desirable to subdivide this heading by using two subheadings--"Expected Probability" and "Observed Probability". For example, if the testing limits for the test were established at $+3\sigma$, the Expected Probability would have a value of 0.0026.
- (D) Filtering Standard. This heading is addressed to the issue of the criterion used for deciding to display a NO GO indication to the equipment operator. For example, the criterion might be three tests failed out of four tests performed.
- (E) Filtered False Alarm Probability. This is the probability that a false alarm will occur even when the Filtering Standard, as given in (D) above, is applied.
- (F) Weighted False Alarm Rate. In order to obtain the Weighted False Alarm Rate required for an entry under this heading, the number of tests of this type, obtained from Col. (B), is multiplied by the Filtered False Alarm Probability, obtained from Col. (E). The individual weighted values could later be added up for the system and then divided by the total number of tests to obtain an average filtered false alarm probability for the system as a whole.

5.4 DISTRIBUTION OF BIT CALLOUTS AMONG AIRCRAFT

BIT data shows that certain types of fault callouts occur much more frequently than others. The question arises as to whether these frequently occurring callouts tend to be generated by one or two aircraft or are occurring among many different aircraft in the data sample. This issue was examined in the context of System 1 in which individual fault callouts are known as DPs (Decision Points).

Before describing the analysis conducted to illustrate the distribution of DPs among aircraft, it is appropriate to speculate why some DPs occur more frequently than others. Here are three possible reasons for high DP occurrence rates:

1. A particular element of the system may be failing at a high rate.
2. Something may be wrong with the way the test corresponding to that DP is mechanized.

3. The high DP occurrence rate may indicate the presence of Category II false alarms.

Since System 1 has been in operational use for a number of years, the first two of these reasons can be rejected almost out of hand. Over these years, modification programs have been conducted to improve the reliability of any high failure rate items. Similarly, gross errors in BIT mechanization have been discovered and corrected. Elimination of 1 and 2 as predominant reasons for high DP rates puts the focus on the third reason.

When military maintenance men discover that their efforts to eliminate given BIT NO GO indications are not rewarded---when they take the implied kind of maintenance action and the phenomenon still continues---they are inclined to decide that the particular indications are not too significant. The indications get ignored and continue to occur. Under these conditions, it seems reasonable to expect that high-frequency DPs would be found occurring in a number of aircraft rather than only one or two. In fact, a finding for a mature system like System 1 that high-frequency DPs are occurring randomly among various aircraft can almost be taken as prima-facie evidence that these DPs are false alarms. This correlation is illustrated in the scatter diagram of Figure 5-2.

In order to test the hypothesis that frequency of DP occurrence is correlated with the number of different aircraft in which the DP occurs, the product moment correlation between these two variables was calculated. The data sample used in this analysis included 183 different DPs and 31 different aircraft. These data were collected over a period of 18 months of operational usage. The findings in the form of distribution statistics on the two variables and the correlation coefficient were as follows (with supporting data in Appendix E):

	<u>Number of Times DP Occurred</u>	<u>Number of Different Aircraft On Which DP Occurred</u>
Mean	24.9	8.6
Standard Deviation:	33.7	6.6
Product Moment Correlation:		0.80
Number of Different DPs encountered:		183
Total number of aircraft involved:		31

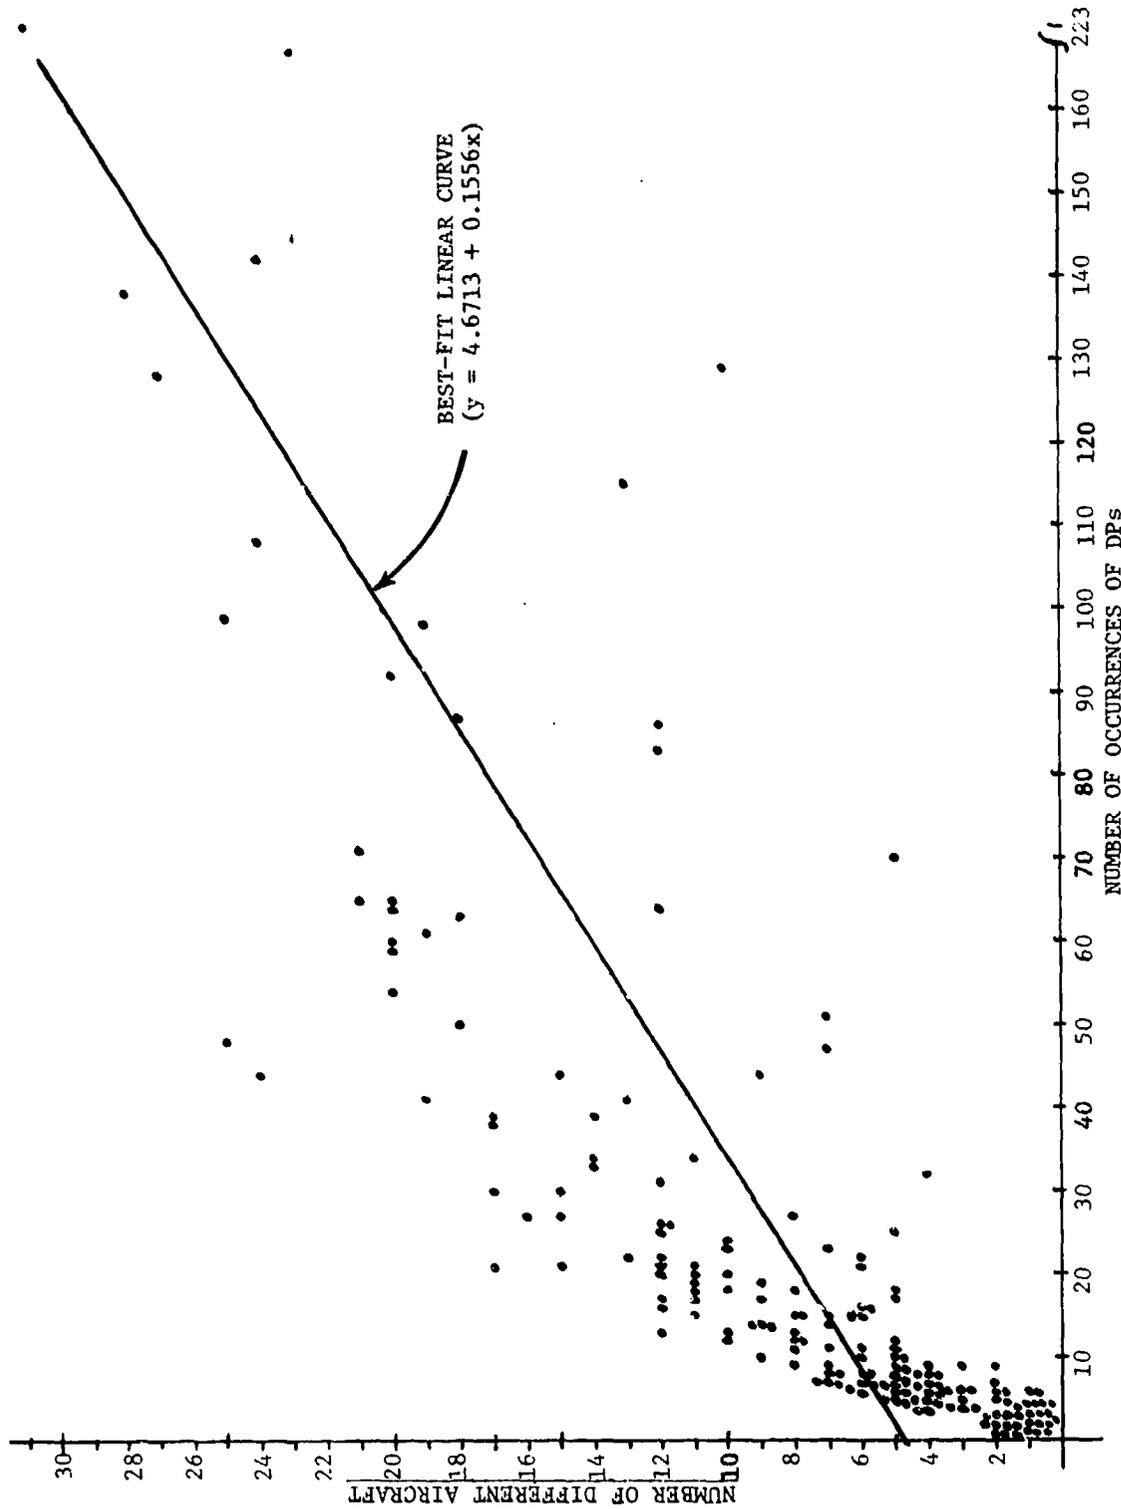


Figure 5-2. Scatter Diagram, Number of Times DP Occurred versus Number of Different Aircraft on Which DP Occurred

The probability that a correlation coefficient of this magnitude could occur by chance is statistically very unlikely--less than one chance in a thousand. It can therefore be stated that there is a strong tendency--at least for this data sample--for high DP occurrence rates to be associated with the occurrence of DPs on a number of different aircraft instead of only on one or a few aircraft.

Although these findings are necessarily limited to the one weapon system from which the data was obtained, it does seem reasonable to expect that the findings would be similar for other data samples of a like nature. In fact, the occurrence of the same BIT fault callout on a number of different equipments could be taken as a symptom indicative of the possible presence of a Category II false alarm condition. Such a finding could be used as a starting point for investigation of this possibility.

6. DESIGN GUIDELINES

The guidelines presented in this section address only the need for avoiding BIT false alarms. In order to arrive at a truly optimal BIT design, other criteria such as those in RADC-TR-80-111 and NAVMATINST 3960.9 must be considered.*

Our approach to specifying BIT (deemphasis on false alarm numerics, more emphasis on techniques) represents a sharp demarcation from the traditional role played by specifications but we need only point out that past specifications have not only not solved the problem, but to some extent they have created the problem. Clearly there is a need for a radical departure from the past.

These guidelines have been based primarily on Hughes experience with the performance of real systems working in an operational environment. In addition, some guidelines are presented which are described in the literature.

6.1 GUIDING PRINCIPLES

1. The problem of BIT false alarms must be treated with the same level of respect that is now accorded the field of reliability.

Although tremendous strides have been made in making systems more reliable, relatively little has been accomplished in reducing the required maintenance effort. Partly this can be explained by saying that reliability gains have been offset by continually increasing system complexity. But many units brought into the I-level shop do not contain part failures.

-
- * (1) RADC-TR-80-111, Design Guidelines and Optimization Procedures for Test Subsystem Design, D. N. Lord, G. A. Walz, S. Green, April 1980, Rome Air Development Center. (Appendix D provides an interesting discussion of the interrelationships between intermittent malfunctions and BIT false alarms.)
 - (2) NAVMATINST 3960.9, Built-in-Test (BIT) Design Guide, Test and Monitoring Systems Office (MAT 04T), Naval Material Command, 1 July 1976. (An upgraded version was released on 19 September 1979, as NAVMATINST 3960.9A.)

The explanation appears to be that maintenance is not based on failures but on test indications of functional failures. Thus, if a high proportion of such indications are false, it can be expected that a large proportion of the resulting maintenance actions will be false (i.e., unnecessary). This study concludes that this is indeed the explanation of why reliability gains have not resulted in commensurate lower support costs. The problem is not primarily a reliability problem; instead it is one of false alarms and can be identified with design of built-in-test. Because of false alarms, missions can be aborted, systems can be grounded, and units can be removed unnecessarily. The consequences of the false alarm problem are thus indistinguishable from the consequences of a serious reliability deficiency. Clearly, both problems must be treated with the same level of respect. As long as systems continue to reflect increasing complexity, it can be expected that both problems will require continuing attention.

2. In the past, BIT designers have concentrated on detection/isolation tasks. In future generations of BIT, designers must expand their viewpoint to include heavy emphasis on interpretation of detected system anomalies.

Current BIT designs are superb at detecting system anomalies. The assumption is generally made that such anomalies represent system failures and, by definition, failures are functional manifestations of faults. Our experience has been that the true situation is much more complex than this and the solution lies in a broader view of system performance. In many ways, the performance of complex systems is comparable to the performance of the human being.* The interpretation task consists of deciding when such system anomalies are truly indicative of defects justifying maintenance action. The task is extremely difficult (more so as systems become more complex) because false symptoms can be very convincing. In the human, attacks of indigestion can almost perfectly mimic heart attacks. In the machine, there can be periods when the system has clearly died, only to be resurrected

* A recent magazine article indicates many physicians estimate that as many as half of their patients either need no medical help or have temporary conditions for which they believe the value of medical aid is extremely marginal. Here, too, there is a large amount of unnecessary maintenance!

by the simple act of setting the system. Too often in the past the interpretation task has been left to the maintenance person. This is totally at odds with the fundamental concept of smart machine/low-skilled operator. The maintenance person has enough trouble wrestling with the type failure that he can detect by straightforward observation. In the future, any BIT that requires operator interpretation must be considered as being not very "smart." In fact, this is precisely the situation that generally exists today.

3. Adequate resources must be allocated to the task of eliminating false alarms.

Current BIT designers are frequently faced with an unreasonable allocation of resources. Typically, they must contend with severe constraints and limitations, such as limited computer memory and constraints on the amount of hardware that can be allocated to BIT. Solution to the false alarm problem requires that adequate resources be made available to the designers. BIT performance must be accorded the same attention as, say, power output and radar detection range.

4. The design of future generations of BIT must be optimized to meet real world conditions, including consideration of:

- o Real world system performance.
- o Real world environmental factors.
- o Real world operations environment and needs.
- o Real world skill levels of maintenance personnel.
- o Real world support constraints and limitations.

Both the customer and contractor must recognize that it is virtually impossible to anticipate all real world factors early in the development of any system.

Early in the program, specification descriptions of system performance represent absolutely the best source of information and BIT must be tailored to be compatible with such performance. On the other hand, when the systems enter field trials and flight testing, inevitably it will be found that the system does not perform precisely as expected and system/BIT incompatibilities will be discovered (for example, test tolerances will need adjustment). BIT must be adjusted to match real world performance. This process is usually repeated when the system moves out to the field and is declared operational. New factors will be discovered which will mandate additional BIT changes. (The painful consequence

of not doing so is unnecessary and expensive maintenance, i.e., the current field problem.)

5. BIT must be designed with a great deal of flexibility, encompassing:

- o Ease of incorporating BIT software changes without affecting tactical software.
- o Ease of changing test limits. (Ideally, several selectable limits should be built into the system, beforehand.)
- o Ease of operator control. (Perhaps the operator should be able to change the makeup of different test sequences.)

Design of BIT must be recognized as an evolutionary process, never ending as long as changes are permitted in the prime equipment.

6. BIT and I-level test equipment designs must be compatible: they are both parts of a single test system.

In a practical sense, judgment whether or not BIT has generated a false alarm depends upon whether or not a failure is found at I-level. If I-level test mechanizations are put together independently of BIT test mechanizations, the opportunities for I-level test results to imply the presence of BIT false alarms are maximized. For example, if a unit with a BIT-detected fault checks No Fault at I-level, the BIT failure indication--although correct--will be categorized as a false alarm. If the I-level mechanizations mirror the BIT mechanizations, the problem is minimized. This is, of course, the underlying philosophy of the concept of federated BIT. Having the test circuit built into the unit permits the unit to be tested at I-level in precisely the same manner as at O-level, with precisely the same test stimuli and precisely the same assessment standards.

Generalizing the preceding discussion, it should be recognized that I-level can either call bad equipment good or good equipment bad (in addition to correctly categorizing the equipment). If O-level and I-level were congruent, I-level could correctly filter out BIT CAT I and CAT II false alarms. If the two type tests are not congruent there will be a higher incidence at I-level of calling bad equipment good and good equipment bad.

There are many aspects to the problem. For example, there is the dilemma of "hidden defects," i.e., faults which only exhibit themselves under operational conditions and are therefore invisible in shop environments (assuming no stress testing). A unit with such a defect may cycle

back and forth between O- and I-levels, swelling the total number of maintenance actions to many times what it should be. One solution to the "hidden defect" problem is to perform limited stress testing at I-level. (In a 6 month experiment at Holloman Air Field, Hughes supplied a three-axis, quasi-random vibrator mounted inside a portable thermal chamber to be used as diagnostic tool. Although very limited testing was actually performed, in the case of a central computer, failures were detected/isolated in 3 different units--out of 7--which checked no fault when tested in the normal manner.)

Another approach to the problem of hidden defects is simply to return suspect units to the contractor for detailed troubleshooting. In system 1, over a 3 year period, 200 such problem units were returned to the contractor site. Although it sometimes took extraordinary effort, in practically every case a hidden defect was eventually detected.

6.2 SYSTEM DESIGN GUIDELINES

As used here, the term "design guidelines" pertains to design characteristics which an end-item of electronic equipment must have if the kinds of BIT false alarm problems currently extant are to be avoided. These design guidelines should be identified to the contractor/designer of the equipment as functional features which must be present in order for the end-item to meet its specifications.

Until now, requirements for BIT have been stated quantitatively in such terms as "testing thoroughness", fraction of detected faults which must be isolated to the correct faulty unit or module, and allowable rate at which BIT indications can be false alarms. This approach to stating BIT requirements has not worked in the sense that it has not led to desired improvements in maintenance effectiveness. There are two principal reasons why these quantitative approaches to the establishment of BIT requirements haven't been effective:

1. Compliance with the requirements is difficult to verify.
2. The "testing thoroughness" requirement may in some instances have been counterproductive in that a stringent requirement can lead to unnecessary proliferation of tests and, hence, to an unnecessarily high rate of BIT false alarms.

The BIT design guidelines which follow are a direct outgrowth of the findings from this investigation and findings

of earlier Hughes investigations into the problems of BIT design (e.g., see Appendix A, "Analysis of Anomalous Maintenance Incidents").

1. Federated BIT

There is a need to decentralize the tests of which BIT is comprised so that a NO GO on a given test directly localizes the implied fault to an element of the system normally replaced at organizational level.

BIT for an electronic system can be thought of as consisting of "system-level" and "unit-level" tests. When NO GO indications are present with current BIT designs, a common basis for localizing faults to the unit which should be replaced is to use logical relationships among the system-level NO GOs. This approach is frequently ineffective, particularly in instances in which BIT indications imply the existence of multiple faults. Fault isolation logic which does not apply to the conditions which exist at the time the test is performed results in Category I BIT false alarms (fault exists but is incorrectly isolated).

The cure for Category I BIT false alarms is to put the BIT tests "in the box". Under these conditions, when a NO GO is present, the location of the fault is unmistakable. The federated BIT concept is to have BIT comprised primarily of "unit-level" tests (some "system-level" tests still needed to verify certain functions). Under these conditions, most faults indicated to be present will be correctly localized to the unit which is faulty.

2. Continuous Monitoring

There is a need to have BIT results based on an integration of successive measurements of a signal over some span of time instead of having the results based on a one-shot check of the signal.

For electronic systems currently operational, BIT usually requires an operator-initiated procedure. Typically, the operator accomplishes the test by placing the system in a test mode and allowing programmed checks to be made. Under these conditions, each signal checked is usually looked at only one time. When the BIT procedure includes the checking of hundreds of different signals, a few of them may be "out-of-bounds" due to chance phenomena present at the particular instant the check is performed. This can result in a Category II false alarm (BIT indicates that a fault is present but the equipment is operative and does not require maintenance).

By having BIT outputs based on results of continuous monitoring, the requirement for the operator to place the equipment in a test mode (and thus interrupt the accomplishment of normal tactical functions) is eliminated. Actually, there is no requirement to have the monitoring of a signal be literally continuous; the only requirement is that the signal be sampled over a time span. One could use the term "continual monitoring" to describe the required capability, but the term "continuous monitoring" (CM) is already part of the technical jargon and is not likely to be supplanted.

3. BIT DATA RECORDING

Under the assumption that BIT is implemented by use of continuous monitoring, there is a need for a capability to record the successive results of this monitoring for later evaluation.

The use of continuous monitoring by itself does not lead directly to a cure for the difficulties faced by operators and maintenance personnel in utilizing the results from current operator-initiated BITs. The monitored data must be accumulated and summarized in some way if the problem of spurious BIT malfunction indications is to be solved.

The quantity of test result data from continuous monitoring is potentially enormous. However, the amount of data which gets recorded can be kept to manageable size by:

- a. Limiting the number of signals that are monitored.
- b. Limiting the maximum sampling rate.
- c. Reducing the time span over which data are accumulated.
- d. Restricting the type of data accumulated.
- e. Use computational techniques not requiring storage of old input data. (For example, mean values and standard deviations can be based on the results obtained at the last sample time and the current input only.)

4. BIT DATA FILTERING

There is a need to summarize and evaluate recorded BIT data so that the results can be used by equipment operators (to decide how the equipment can best be used to accomplish mission functions) and maintenance personnel (to decide what maintenance, if any, is required).

Recorded BIT data must be summarized and evaluated in such a way that the results serve the needs of both operators and maintenance personnel. How this is accomplished depends upon the specific characteristics of the BIT data which is stored. Stored data may be in one of the following forms:

1. Raw values obtained each time a BIT-monitored signal is sampled.
2. Number of times that a signal is sampled and number of times the signal fails to satisfy required tolerances.
3. Data from which mean and standard deviation values can be calculated for each sampled signal.

To meet the needs of a system operator, the stored BIT data must be retrieved and summarized so as to provide the operator with real-time information as to the status of the equipment. If the equipment has a malfunction, the operator must be told which equipment modes, if any, are still operative. Equipment status information must be continually updated for the operator from the start of a mission until its completion.

To meet the needs of maintenance personnel, the stored BIT data must be retrieved and summarized on the ground after a given mission and, if desired, stored for subsequent use. The process of BIT data filtering must answer the following questions for the maintenance man:

1. Does the system require maintenance?
2. If so, which unit of the system is faulty?

SUMMARY REMARKS

The findings of this investigation suggest that there should be a major shift in the way that BIT requirements are stated for contractor/designer compliance: Instead of stating these requirements entirely or primarily in quantitative terms (e.g., probability of fault detection, probability of fault isolation to the correct LRU), there is also a need to state the requirements in terms of functional capabilities (e.g., continuous monitoring, BIT data recording, BIT data filtering) known to be needed to bring about improvements in the operation and maintenance of electronic systems.

The question arises: Is it feasible to design electronic systems so that they are provided with the types of functional capabilities which have been stated to be required? All of the information available to us here at Hughes suggests that the capabilities are achievable. It has not been possible for us, under the auspices of this investigation, to evaluate the relative cost and effectiveness of different design approaches which might be used to achieve

a federated BIT design, to do continuous monitoring, or to accomplish data storage and retrieval. We have, however, concluded that implementation of these capabilities must be treated on a system-by-system basis.

6.3 PROCEDURAL GUIDELINES

Possibly the most significant procedural guideline is that development of BIT should be carried out in a very similar manner to reliability programs. Such programs include Test, Analyze and Fix (TAAF) tasks and have well defined objectives. Achieved performance is continually assessed and compared to objectives. The same kind of dedication needs to be given to the development of BIT.

Another important procedure is institution of an adequate data system for collecting BIT data and correlating such data with maintenance data. Hughes has demonstrated the feasibility of implementing such a system in a "MORPEP" study (excerpts in Appendix D). Such a data system should be considered mandatory in developing future systems.

6.4 OTHER CONSIDERATIONS

Additional techniques for preventing or suppressing false alarms due to momentary anomalies are discussed below.

(a) Environmental Sensor Input

A very common type of false alarm occurs when the real problem is some momentary environmental stress, for example, momentary high temperatures. Some authors feel that this is the major source of false alarms and "that false alarm prevention amounts to accurately distinguishing between failures in a module and failures in its environment (i.e., fault isolation up to the module)." * A possible solution to this type of problem is to provide an additional sensor so that the momentary stress can be detected. By correlating the presence of this stress with a test failure, logic can be used to avoid incorrectly isolating the problem to the equipment.

Electronic hardware that is sensitive to excessive temperature or excessive aircraft motion can fail BIT if these excesses are present during test. One of the avionics systems analyzed had a light in the cockpit that indicated excessive coolant temperature. This condition could be associated with

* "A Preliminary Study of Built-In-Test for the Military Computer Family (MCF)," report number CORADCOM-76-0100-F, by J. Clary, A. Jai, S. Weikel, R. Saeks and D. Siewiorek, and dated March 1979.

a BIT display of certain faulty units. The output of the temperature sensor should be made available to the central computer so that displayed and/or stored test results would indicate an anomalous environmental condition at the time of test. Similarly, rate and attitude sensor data should be used to modify BIT results affected by these influences. The environmental sensors should also include a humidity or moisture sensor if applicable. (One of the airborne systems analyzed exhibited intermittent computer anomalies because of wet computer boxes.) Ideally, sensor level and tolerances should be programmable or adjustable so appropriate values can be inserted after field evaluation.

(b) Test Tolerances

The subject of setting test limits and adjusting test tolerances is a crucial one. Suffice it to say here that limits and tolerances must be matched to real world performance. This process should be extended into flight testing and early field tests. This entails measurement of all critical signals via instrumentation and requires careful analysis of the data. The statistical approach is essential, i.e., such statistical parameters as means and standard deviations need to be computed. Computer processing is necessary because of the large mass of data involved. After the instrumentation package is removed and after the systems have been moved into a true operational environment, continued vigilance must be exerted. This subject is discussed further in Appendix B, "Considerations in Setting Electronic System Testing Tolerances."

(c) Power Transient Monitoring

Electronic systems using a central computer generally have a transient monitoring function to protect the computer operation during a power transient condition. If the transient is a primary power transient, it can also affect other units of the system. If the computer detects a transient during a BIT test, and the test fails, the test should be repeated or the fail display inhibited. For equipments that don't operate through a central computer or operate from other power forms, additional transient detectors can be incorporated in the tested hardware.

(d) Motion Sensor By-pass

Electronic systems having rate and attitude sensors should have provisions for providing a fixed output during BIT to prevent motion inputs from biasing system test results, or the rates could be provided to the central computer to enable automatic test limit compensation.

(e) Faulty Wire Test

In a system with multiple units, usually only one point is monitored at each unit, the output. If both input and output signals are monitored, faulty interconnections can be detected. This usually has not been done in the past because adding a wire to check a wire was not cost effective. With integrated circuits, test signals can be multiplexed back to the central computer with no additional interconnecting copper. In the case of digital commands to a unit, a simple "or" gate can monitor test input status as the computer sets the commands high one at a time.

(f) Airborne vs Ground or Operator vs Maintenance

Since the airborne environment is more severe than the ground environment, tests can fail airborne and the failure be not repeatable on the ground. By monitoring the weight-on-wheels indication, the test software could use one set of tolerances for ground test and a wider set for airborne test. Another approach would be to have loose tolerances for the system operator who is interested in mission essential performance only and a tighter tolerance set of tests for maintenance use.

(g) Multiple Run Entry

This is a test mode that automatically repeats a selected test for a given number of cycles or until the operator terminates the test. This mode does not directly prevent false alarms but does aid in confirming intermittent failures, thus separating them from false alarms. This mode is recommended for high speed electronic devices but not necessarily for low speed mechanical devices.

(h) Failure Recording and Weighting

This type of failure filtering may be applied over a period of a few seconds to a few weeks. For instance, if a failure is sensed, the test control program could repeat the test before displaying the test results, thus filtering out short term transient effects. In the case of continuous monitoring BIT, a delay should be invoked before displaying the fault, and if the fault clears, the operator is not alerted. There are in most systems some critical functions for which a short term intermittent failure could have a catastrophic effect on the mission, such as a momentary computer hangup during a missile launch. Thus failures such as an intermittent computer check sum failure should be displayed at the time of failure and the consequences assessed by the operator based on the mission profile.

Another type of filtering can be applied if the system has some type of permanent memory or recording device. One-time occurrences that were inhibited from display because of their momentary nature would still be recorded and either the on-board computer or a ground based computer (if the data are recoverable) used to make a determination between false alarms or incipient hardware failures by observing the long term failure patterns.

The primary reason for collecting empirical data is to distinguish between actual and expected performance.

(i) Prioritizing and Structuring BIT Tests

BIT false alarms can be reduced by carefully structuring BIT tests into different groups of tests, with the top level consisting of the fewest number of different tests but with these tests encompassing overall system functions. Carried to its logical extreme, the top test might include no BIT tests, per se. For example, one experienced operator reported that when he was stationed in Japan, the standard procedure was to attempt a radar lock onto Mt. Fuji immediately after takeoff. Operators became highly skilled at assessing system performance via this procedure. If the system were functioning properly, all BIT tests were dispensed with. One advantage of this approach is that the system is being checked in a closed-loop configuration. But operators were possibly coerced into using this type of makeshift testing by the presence of BIT false alarms.

The second tier of testing, generally known as confidence testing, should provide the kind of information needed to enhance the probability of mission success. Such information permits the operator to make decisions pertinent to mode selection, weapon selection, frequency channel selection, etc. BIT tests should focus on mission-relevant characteristics. Indications of faults which have negligible impact on mission success should be considered false alarms. Such indications should be recorded for subsequent assessment by maintenance personnel but should not be displayed to the operator during the mission.

The third and subsequent tiers of testing should be designed to optimize the process of returning systems to a full-up state of readiness. They should include:

- o System or "end-to-end" functional testing.
- o Supporting fault isolate tests, to be run if a functional failure has been encountered.

Where a given function is spread over a number of different boxes, a "federated BIT" concept (each unit with its own

built-in-test) accomplishes isolation of a fault to the proper box. BIT can be expanded to include isolation to a shop replaceable assembly (SRA), in which case some I-level tests can possibly be dispensed with. The testing problem is greatly simplified when all circuits associated with a single tactical function are packaged in the same box.

The above philosophy is based on the observation that, with current systems, frequently fault isolation tests will fail when the system is working perfectly according to (1) the operator and (2) the system-level BIT confidence test. Although there is no question that the fault isolation test results are valid-- for example, test limits are being exceeded (marginally)--there is also no question that such results are false alarms, in an operational sense. In general, it is recommended that results of detailed tests should be ignored in the absence of operator squawks and if the BIT confidence test passes. Additionally, it is recommended that when a system failure is encountered, only those specific fault isolation tests that are related to the system failure should be performed.

(j) Maintenance Structuring

In view of the complexity of new systems, it must be concluded that the false alarm problem is a phenomenon that can be reduced but not eliminated. Therefore, it is essential that this subject be considered in training maintenance personnel. As part of this training, maintenance personnel must be encouraged to make careful observations of any peculiar anomalous system performance that they encounter, especially when there appears to be a pattern of such performance. Such experience needs to be carefully documented and analyzed by engineering personnel. Conclusions must be given widespread distribution to all field people and engineering corrective action must be expedited. In short, there must be a systematic approach to solving the false alarm problem.

The false alarm problem is a highly complex one involving many factors. There is no one answer. The only satisfactory approach is to work off each factor as it is identified. Maintenance personnel can play a key role in this identification process. Face-to-face contact and communication between engineering and maintenance personnel need to be encouraged.

7. CONCLUSIONS

BIT false alarms should be considered a top contributor to the problem of excessive support costs for fielded military electronic systems. The attack on the BIT false alarm problem must therefore be continuous and unrelenting. How this attack is to be implemented depends upon one's point of view. The conclusion of this study is that current BIT designs have not been optimally matched to system performance, especially under field conditions. This should not be viewed as a reflection on BIT designs, per se. In effect, BIT designers have been directed by specification to detect system anomalies with high precision and this is precisely what BIT systems do (generally, "BIT does not lie"). The tacit assumption is made that such anomalies can be equated to the need for maintenance. This is a mistake. Many system anomalies do not indicate failure events requiring maintenance action. Thus, many BIT-indicated anomalies are maintenance false alarms. Although this study has taken the first step in generating guidelines for resolving this "failure without a fault" paradox, there is a need for more research. Clearly, there is a lack of understanding and appreciation of (1) the severity of the stresses encountered under operational conditions, and of (2) many of the subtleties of how complex systems perform, whether in stressful or benign environments. Evidence for the second point is the fact that many anomaly mysteries are encountered in a laboratory environment and when airborne systems are tested on the ground, as well as in flight.

We conclude that the first step in solving the BIT false alarm problem is a better understanding of how the prime equipment operates, initially under laboratory conditions and eventually under field conditions. Every instance of anomalous performance must be treated with the same respect that is now accorded hardware failures. It is well understood by all design engineers that reliability must be "designed into" systems and the only way to achieve this goal is (1) to understand the root cause of each and every failure event and (2) to take corrective action to avoid recurrence. Precisely this same attitude must be adopted in the area of false alarms. Every incident of anomalous performance not related to "reliability failures" must be analyzed and root causes established. In some cases, a design action should be taken. For example, if the root cause is a power transient, a design fix might be either a better power supply or improved filtering in the receiving unit. If the anomalous characteristic is deemed to be unavoidable and inherent to the design, BIT must be properly matched to this characteristic, i.e., BIT must be designed to identify and accept "normal" anomalies without generating a failure indication (false alarm). Unless such system behavior is well understood by both prime equipment and

BIT designers, it cannot be said that the designers know how the system works and without this knowledge it is inconceivable that a BIT system can be designed without having a false alarm problem.

Thus, the first step in solving the BIT false alarm problem is assumed to be a better understanding of why fault-free systems intermittently perform in an anomalous manner. It is assumed that such research will lead to improved designs but it is doubtful if the anomaly characteristic can ever be totally eliminated. The second step in solving the BIT false alarm problem lies in the area of designing BIT to cope with the residual false alarm problem. It should no longer be a goal simply to detect system anomalous performance. The next generation of BIT must have the "smarts" built into the design to distinguish between anomalies which are manifestations of faults and anomalies which must be tolerated as characteristic of fault-free equipment. It must be considered totally unacceptable to burden the maintenance person with this interpretation task.

Although detailed analysis of intermittent faults was explicitly excluded from consideration in this study, no discussion of false alarms would be complete without mention of this topic. There is extraordinary similarity between the symptoms of false alarms and those of intermittent faults. Both are inherently intermittent in nature and both are extremely difficult to isolate. Because of these similarities, intermittent faults are frequently written off as false alarms and false alarms are frequently misinterpreted as evidence of hidden faults. A prime requirement is that solutions to the false alarm problem not be allowed to mask the problem of detecting and isolating intermittent faults.

From a management point of view, the key to solving the BIT false alarm problem is a better BIT specification. Past specifications have dealt with the subject on a very theoretical basis, and have totally ignored the problem of intermittent faults. Future specifications must address the real world, including both false alarms and intermittent faults. Unfortunately, there are many aspects to the real world that simply can't be anticipated during the R&D phase of system development. It is suggested that future specifications put less emphasis on specific numerics and more emphasis on demanding the inclusion of techniques for making BIT systems less susceptible to false alarms. Typically, current numerics (e.g., one percent false alarm rates) are so totally unrealistic as to be meaningless (but "provable" if an appropriate definition of false alarm is adopted). One possible solution is to transfer the problem to a separate handbook, which would include design guidelines for guarding against false alarms. The specification could then reference this document, or appropriate parts of the document. Many

excellent studies have been made in defining BIT figures-of-merit, but generally these have addressed the problem on a theoretical basis. The handbook could supplement these studies with considerations based on real world experience. The specification would tell designers what characteristics are to be present and how compliance with the specification is to be verified. The handbook would provide practical ways for implementing the desired characteristics.

The preceding discussion can be summarized in the following manner:

- (a) The first step in solving the BIT false alarm problem is a better understanding of how the prime equipment works.
- (b) Because of the complexity of current and future systems, and because of the complexity of the operational environment, and especially because of the complexity of failure modes, it is too idealistic to expect a full understanding of how the equipment works under all of the conditions it will ever encounter in its life cycle. System anomalies will always exist and so the Category II false alarm problem will continue.
- (c) To cope with the problem of imperfect understanding, it is necessary to design into each electronic system those "tools" needed by both engineering and maintenance personnel in evaluating system performance:
 - o Continuous monitoring.
 - o BIT data recording.
 - o BIT data filtering.
- (d) For maintenance personnel, these tools provide a basis for making maintenance decisions and establishing maintenance policies in coping with day-to-day anomalies which the engineers either don't yet know about, don't yet understand, or don't yet have a satisfactory way of handling in BIT design.
- (e) For engineering purposes, these tools yield information not otherwise available as to how the system performs under the various conditions encountered in its operation. Lack of these tools results in some problems remaining in systems literally for years.

8. RECOMMENDATIONS

It is recommended that investigation of BIT false alarms be continued and expanded to include the problem of intermittent faults. There is little doubt that these twin problems are the root cause of the enormous cost of maintaining systems in the field. These problems have never been resolved because there has never been a systematic, industry-wide attack on them. The payoff for such effort is incalculable, in terms of improved system readiness as well as reduced support costs.

It is recommended that research be carried out in the following areas:

- (1) Continuing research on existing systems for the purpose of further understanding root causes of both false alarms and intermittent faults (with the research to encompass both category I and category II false alarms).
- (2) Development of design guidelines for coping with the two problems, to include both BIT design and the design of I-level test equipment and encompassing the optimal way to use BIT results in the I-level shop.
- (3) Development of the technology required for implementing the required functions, including continuous monitoring, BIT data recording and BIT data filtering.
- (4) Research into all aspects of the BIT data problem, including definition of the types of data that BIT should generate (for both real-time assessment of system condition and for shop assessment), optimal ways to transmit this data (for maintenance crews and for shop personnel) and optimal ways to insert pertinent data into maintenance data systems (with minimal paper work).
- (5) Research into optimal data analysis techniques, with heavy emphasis on utilization of computerized BIT data processing.
- (6) Research into management techniques for alleviating BIT-related field support problems, emphasizing maintenance policies and procedures.
- (7) Research into creative and effective ways to handle the problem of specifying BIT.
- (8) Research into the human factors considerations associated with interpretation of BIT results by maintenance personnel.

APPENDIX A
ANALYSIS OF ANOMALOUS
MAINTENANCE INCIDENTS

APPENDIX A. ANALYSIS OF ANOMALOUS MAINTENANCE INCIDENTS*

BIT false alarms will naturally result in atypical or abnormal (anomalous) maintenance incidents, such as removals of fault-free units. Therefore, any study of anomalous maintenance incidents can be expected to provide evidence pertinent to the subject of BIT false alarms. One such study was performed during 1976, as part of an IR&D (Independent Research and Development) project entitled "Design to Support Cost Methodology." The purpose of analyzing anomalous maintenance incidents was to gain a better understanding of the maintenance phenomena which occur in military operational environments.

Data for the study was collected at Miramar Naval Air Station (near San Diego) and pertained to the maintenance of the AWG-9 weapon control system used in the F-14 aircraft. Two Hughes engineers were assigned to Miramar for a period of five months to collect the required data. Data collection took place during the period April-August, 1976 and pertained to two tactical Navy squadrons, VF-24 and VF-211, stationed at Miramar at that time.

The statements which follow represent a brief summary of the study's findings:

1. Of the 103 downing AWG-9 faults worked on by organizational-level maintenance personnel, 32 resulted in a decision that the reported fault could not be duplicated; 11 resulted in performing at-aircraft maintenance procedures which did not include sending any Weapon Replaceable Assemblies (WRAs) to the intermediate-level shop; and 60 resulted in a decision to send one or more WRAs to the intermediate-level shop for maintenance.
2. A total of 77 WRAs were sent to the shop in all; of these, 58 WRAs (75%) were found to be defective and 19 (25%) were found to be nondefective.
3. During missions, the AWG-9 system is operated by a Naval Flight Officer (NFO). It was found that when the NFO uses the BIT fault isolation sequences during the mission, there is a significantly better chance that the WRAs later removed by maintenance personnel will check out bad when tested in the intermediate-level shop.

* The investigation, "Analysis of Anomalous Maintenance Incidents" was conducted under the supervision of Dr. Richard W. Highland. The content of this appendix was extracted from the complete report published in the Hughes proprietary document, Designers' Support Cost Prediction Handbook.

4. In initiating AWG-9 maintenance (i.e., in performing fault verification and isolation), organizational-level maintenance personnel seldom used the BIT fault isolation sequences unless the NFO had previously used these sequences during the mission. If there was no NFO use of the BIT fault isolation sequences during the mission, maintenance personnel were found to use these sequences only 10.5 percent of the time in performing fault verification and isolation.
5. Use of BIT by the NFO is predominantly on the deck. For example, use of the BIT confidence test sequences during the in-flight portion of missions was found to occur on fewer than 25 percent of the missions.
6. In-flight use of BIT by the NFO has a relationship to whether or not maintenance personnel are able to duplicate a reported fault on the ground. The probability of being able to duplicate a reported fault during maintenance is significantly higher when the NFO has used BIT during the in-flight portion of a mission.
7. For those maintenance cases terminated with the decision that the reported symptom could not be duplicated, maintenance personnel always used BIT in their attempts to duplicate the symptom. However, for cases in which reported symptoms were duplicated, maintenance personnel used BIT for fault verification and isolation only 65 percent of the time.
8. Reasons for not using BIT during maintenance were approximately equally divided among:
 - a. The type of malfunction precluded running BIT.
 - b. It was feasible to confirm and isolate the fault by operating the AWG-9 system in non-BIT modes.
 - c. Maintenance personnel relied on prior experience rather than BIT in isolating the fault.
9. No significant difference was found in the tendency of WRAs to check out good in the intermediate-level shop according to whether BIT was used in accomplishing organizational-level fault isolation.
10. The mean number of different BIT sequences used by maintenance men for fault verification and isolation in cases where the reported fault was duplicated was 2.9; for cases where the reported fault could not be duplicated, the mean number of BIT sequences used was 4.7. The difference between these two means is statistically

significant (i.e., the difference is greater than would be expected under the assumption that the two samples were drawn from the same population).

11. Although the presence of a large number of BIT Decision Points (DPs) and BIT callouts of potentially-faulty WRAs can be thought of as a possible source of confusion for maintenance men in their attempts to isolate faults, the data offers no evidence that this is the case. There was no significant difference between the mean number of BIT DPs and BIT WRA callouts for cases where the WRA's sent to the intermediate-level shop checked good as compared with cases where the WRAs checked bad.
12. About one-third of the WRAs found faulty in the intermediate-level shop require some special type of maintenance action to be performed (i.e., an action other than removing and replacing a Shop Replaceable Assembly or performing adjustments). Cleaning connector pins and repairing shorted or broken wires are the most frequent of these activities.
13. No evidence was found that shop maintenance personnel use shortcut approaches in checking out WRAs. Apparently the complete checkout procedure is used each time a WRA is processed through the shop.
14. AWG-9 WRAs with high shop-check-OK rates tend to be less complex than the average AWG-9 unit, to have relatively long Mean Times Between Maintenance Actions, to have relatively short shop repair times and to have fewer Shop Replaceable Assemblies which must be sent to a depot for repair. Shop adjustment is required less frequently for these than for other WRAs.
15. AWG-9 WRAs with low shop-check-OK rates tend to have characteristics opposite of those cited above for high shop-check-OK WRAs. These WRAs tend to be more complex than average, to have short Mean Times Between Maintenance Actions, to have lengthy shop repair times, to have a relatively large proportion of Shop Replaceable Assemblies which must be sent to a depot for repair, and frequently require adjustment in the course of shop maintenance.
16. WRAs with low shop-check-OK rates tend to contain a large number of individual adjustment controls. There is a correlation between the number of individual adjustment controls that WRAs contain and the percentage of time that the WRAs get adjusted in the course of being processed through the intermediate-level shop.

17. Each WRA has an "expected" shop-check-OK rate based on its physical characteristics. Shop-check-OK rates which are higher than expected tend to occur when the WRA is not thoroughly tested by BIT or if it has inherent fault isolation ambiguity with other WRAs in the same testing loop. Shop-check-OK rates which are lower than expected tend to occur if a WRA is thoroughly tested by BIT, if the BIT approach for testing the WRA does not allow inherent fault isolation ambiguity with other WRAs to occur, or if the WRA has special features which allow faults to be isolated to it without use of BIT.
18. WRAs removed from AWG-9 systems which are judged to have multiple problems (i.e., several distinct and presumably independent AWG-9 fault symptoms) have lower shop-check-OK rates than do AWG-9 systems which exhibit only single problems.
19. When several fault symptoms are present in an AWG-9 system, selective processes operate in determining which of these symptoms get worked on by organizational-level maintenance personnel after a given flight.
20. WRAs removed from AWG-9 systems under the condition that the reported discrepancy was "BIT Only" or "Writeup Only" have higher shop-check-OK rates than is the case where the reported discrepancy is "Writeup Confirmed by BIT".
21. In instances where a reported AWG-9 discrepancy could not be duplicated, there was no reported recurrence of the symptom on the next flight by the same aircraft in 83 percent of the cases (i.e., recurrences were reported in 17 percent of the cases).
22. In instances where a reported AWG-9 discrepancy was duplicated and worked on by organizational-level maintenance personnel, there was no reported recurrence of the symptom on the next flight by the same aircraft in 86 percent of the cases (i.e., recurrences were reported in 14 percent of the cases).
23. In 14 instances where single WRAs checked good in the intermediate-level shop, reported recurrences of the fault on the next flight by the same aircraft (with a different WRA of the same type installed) were present in 21 percent of the cases.
24. The factors which tend to depress the percentages of next flights on which discrepancies are reported to recur, include the following:
 - a. There may have been no valid discrepancy during the flight on which the discrepancy was originally reported.

- b. Failure symptoms vary in their degree of observability by the NFO.
- c. There are mission-to-mission variations in the observability of a given fault symptom.
- d. There are NFO-to-NFO differences in tendencies to report discrepancies, given that the discrepancies are observable.

APPENDIX B
CONSIDERATIONS IN SETTING
ELECTRONIC SYSTEM
TEST TOLERANCES

APPENDIX B. CONSIDERATIONS IN SETTING ELECTRONIC SYSTEM
TEST TOLERANCES

One factor which affects the rate at which BIT false alarms occur is the way in which limits are established for the tested signals. The purpose of this appendix is to cite certain variables related to the occurrence of BIT false alarms and to briefly describe how these variables operate to influence observed false alarm rates.

The following definitions are pertinent to a discussion of this topic:

1. Required Limits. For a given signal or parameter, Required Limits define a range of performance needed to satisfy formally-stated objectives.
2. Decision Limits. For a given signal or parameter, Decision Limits define a range of performance treated as acceptable for testing purposes. Values indicated to be inside the Decision Limits are treated as GO and values indicated to be outside the Decision Limits are treated as NO GO.
3. Measurement Error. Measurement Error is that characteristic of a testing device or measurement procedure which causes signals to be evaluated at values which differ from the true values. The standard deviation of the differences between individual measured values and corresponding true values can be used as an index of the extent of measurement error.
4. Consumer's and Producer's Loss Probabilities. These probabilities are mentioned here primarily because the terms are sometimes used in connection with formalized statements of testing requirements. The Consumer's Loss Probability (CLP) is the probability that a measurement results in a GO indication when, in fact, the true value of the signal or parameter evaluated is outside the Required Limits. The Producer's Loss Probability (PLP) is the probability that a measurement results in a NO GO indication when, in fact, the true value of the signal or parameter evaluated is inside the Required Limits. In instances in which CLP and PLP take on values greater than zero, it is because measurement error is present or because Decision Limits have been placed at some point other than

the Required Limits. The concepts of CLP and PLP are of limited explanatory value in the understanding of Category II BIT false alarms.

5. Signal Distribution Statistics. These statistics include the mean and standard deviation of the tested signal as these quantities manifest themselves under the conditions in which BIT is performed.

The establishment of BIT tolerances is sometimes an iterative process. The usual starting point in setting the Decision Limits for a particular BIT check is to set these limits to match the equipment specification. But are the limits given in the equipment specification actually the Required Limits? The values given in a specification are often internally developed by the contractor/designer and may not represent required performance.

When a BIT check is provided with Decision Limits which match the equipment specification and it is found that the check is indicating frequent failures, the BIT designer is faced with a problem. There are the following possible explanations of the frequent failures:

1. The tested parameter may actually be failing when BIT is indicating it to be failing (i.e., failed parts may be causing the test failure).
2. The frequent failures may be attributable to Measurement Error.
3. The specification from which the Decision Limits were obtained may be in error. That is, the equipment may perform satisfactorily even though the specification values are not met.
4. Normal performance of the tested equipment may not match the requirements which the equipment must be able to satisfy.

The BIT designer must somehow choose from among these possibilities. If failed parts are causing the test to fail, this will be rather quickly discovered and eliminated as a cause of the BIT fails. Measurement Error is usually not a problem since BIT is usually highly accurate in relation to the tolerance limits in question. Having the specification values to which the Decision Limits are matched be in error is a definite possibility. The approach to discovering that this is the case is one of verifying that the equipment is operating properly at the time the BIT NO GO is observed to occur.

The fourth of the possibilities represents a serious problem. But it is not a problem that BIT is going to solve.

That is, if the equipment does not meet its requirements, having BIT indicate NO GO each time this testing procedure is performed isn't going to correct the condition. The remedy is that the equipment design or its functioning must somehow be altered so that the requirements are satisfied. If this remedy is not going to be applied and the equipment is going to become operational anyway, it could make sense to change the BIT Decision Limits to conform to the way the particular tested signal is currently behaving. BIT should not tell the maintenance man that something is wrong if the maintenance man can do nothing to correct the condition.

Most Category II BIT false alarms are probably attributable to interaction between BIT Decision Limits and actual behavior of the tested signal as reflected by Signal Distribution Statistics. There are many possibilities as to how this interaction takes place: The average signal value may be higher or lower than expected. The distribution of signal values may be skewed to the high side or to the low side. Or the signal values may be more variable--or more variable under certain conditions--than was believed to be the case. Any incompatibility between the BIT Decision Limits and the Signal Distribution Statistics may range from gross incompatibilities to very subtle mismatches. BIT designers must be able to recognize such discrepancies and, having found them, decide what course of action can be taken.

APPENDIX C

MISSILE-ON-AIRCRAFT-TEST

(MOAT) FALSE ALARM STUDY

APPENDIX C. MISSILE-ON-AIRCRAFT-TEST (MOAT) FALSE ALARM STUDY

The most severe BIT false alarm problem that we know of existed in connection with testing the missile on board system 1, during the time period 1976-1977. The problem was particularly severe because of the relatively large size of the missile (many hundreds of pounds). The BIT results had such little credibility that when BIT faulted the missile, the standard procedure was to download the missile (with great difficulty, compared to a small avionic unit), and transfer the missile to another aircraft station and then to retest it. If the missile retested faulty, the results were still suspect. Standard procedure was then to transfer the missile to another aircraft and retest. If the same fail results were achieved, the problem was assumed to be confirmed and the missile declared to be faulty. More often than not, the missile checked OK during the final check, i.e., the preceding BIT indications were false alarms.

Extensive review of field experience with MOAT was conducted to identify the problem areas. Major revisions were recommended to improve MOAT, including increased test thoroughness and fault isolation, more effective displays and reduced sensitivity to failures caused by spurious influences such as r-f interference and aircraft maneuvers (when the test was performed airborne). As a part of that program considerable attention was given to the problem of intermittencies and non-repeatable test results. These efforts are reflected in the design of MOAT as mechanized in the current software program. This program includes significant improvements specifically designed to reduce intermittent failure indications, especially those due to external test interference caused by aircraft motion and by signal corruption due to vibration or electromagnetic interference. Adaptive test limits for the autopilot analog report line are implemented. (The adaptive test limits compensate for motion.) Another improvement is the computation of accelerations along the missile pitch and yaw axes for invalidating test results when system design limits are exceeded. A third improvement is the sampling of analog report lines and the implementation of statistical computations to deal with the signal corruption problem. A statistical best linear fit is used to determine the proper analog values while the variance of the samples is used to determine when noisy analog values should be discarded. A fourth improvement is the accumulation of valid test results to deal with system degradation. Accumulation of test results obtained by consecutively repeating MOAT permit the indication of incipient as opposed to hard failures in the test result displays. Incipient failures occur when a system function begins to fail intermittently.

Our investigation indicates that the MOAT false alarm fixes are working well, including the readiness tests and the adaptive threshold. There is some difficulty being experienced by operators in making use of the intermittent failure storage and display for multiple runs of the MOAT test. Many operators do not fully understand the concept. This perhaps indicates the need for better training and education in the field of detecting/isolating intermittent faults.

APPENDIX D
MAINTENANCE, OPERATION, RELIABILITY AND
PERFORMANCE EVALUATION PROGRAM (MORPEP) STUDY

APPENDIX D. MAINTENANCE, OPERATION, RELIABILITY AND
PERFORMANCE EVALUATION PROGRAM (MORPEP) STUDY

The MORPEP study* consisted of determining the requirements of a data system for processing both BIT data and maintenance data. To prove the feasibility of such a concept, software programs were developed for processing a sample of BIT and maintenance data and a computer system was used to generate sample reports. The key feature of the system was its ability to merge separate files of BIT and maintenance data into a single computer printout. This unique capability provides the analyst with the ability to compare BIT fault indications with the success or failure of subsequent maintenance actions. This process leads to isolation of problem areas so that corrective action can be taken.

To illustrate the concept, assume that a high rate of occurrence of some BIT decision point (DP) can be linked to a high no-fault rate at I-level, i.e., units being removed per the BIT DP usually check OK at I-level. It would be easy to establish whether or not most of the removals are associated with a particular unit (by serial number). If so, this unit should be withdrawn from the inventory and given special testing, since there is a high probability that the unit contains some peculiar or intermittent fault. This illustrates the point that great care must be utilized in categorizing problems. Frequently, real, valid faults can masquerade as BIT false alarms.

Although the MORPEP concept has never been implemented, it is considered mandatory for future programs to give serious consideration to the need for implementing data systems which can perform the type analysis described above.

The attached figures illustrate the MORPEP concept.

* Sponsored by the Naval Air Systems Command Technical Representative at Hughes.

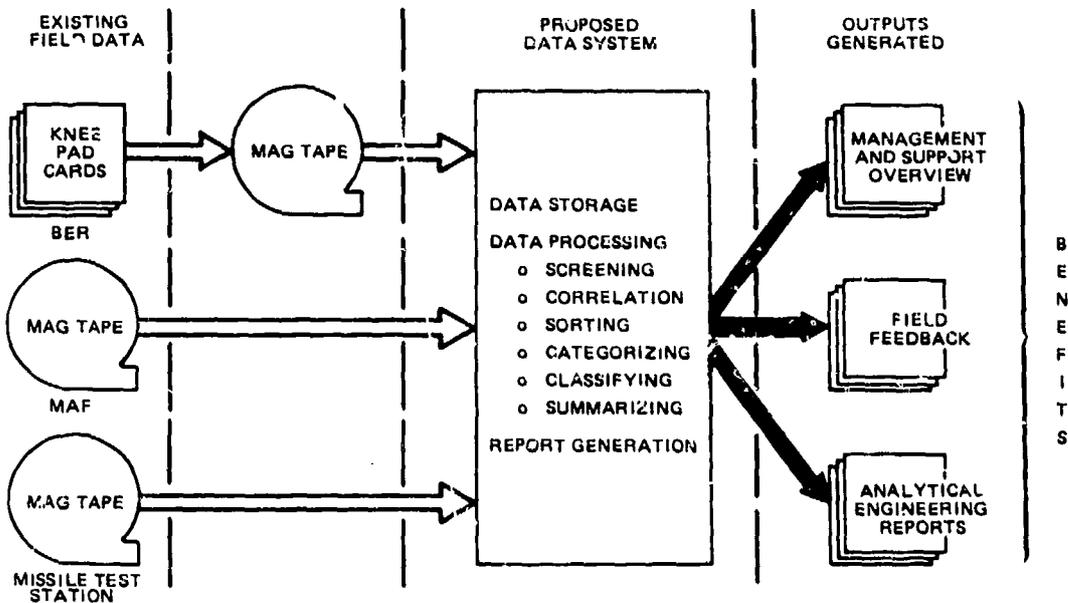


Figure D-1. Simplified Diagram Illustrating MORPEP Inputs and Outputs.

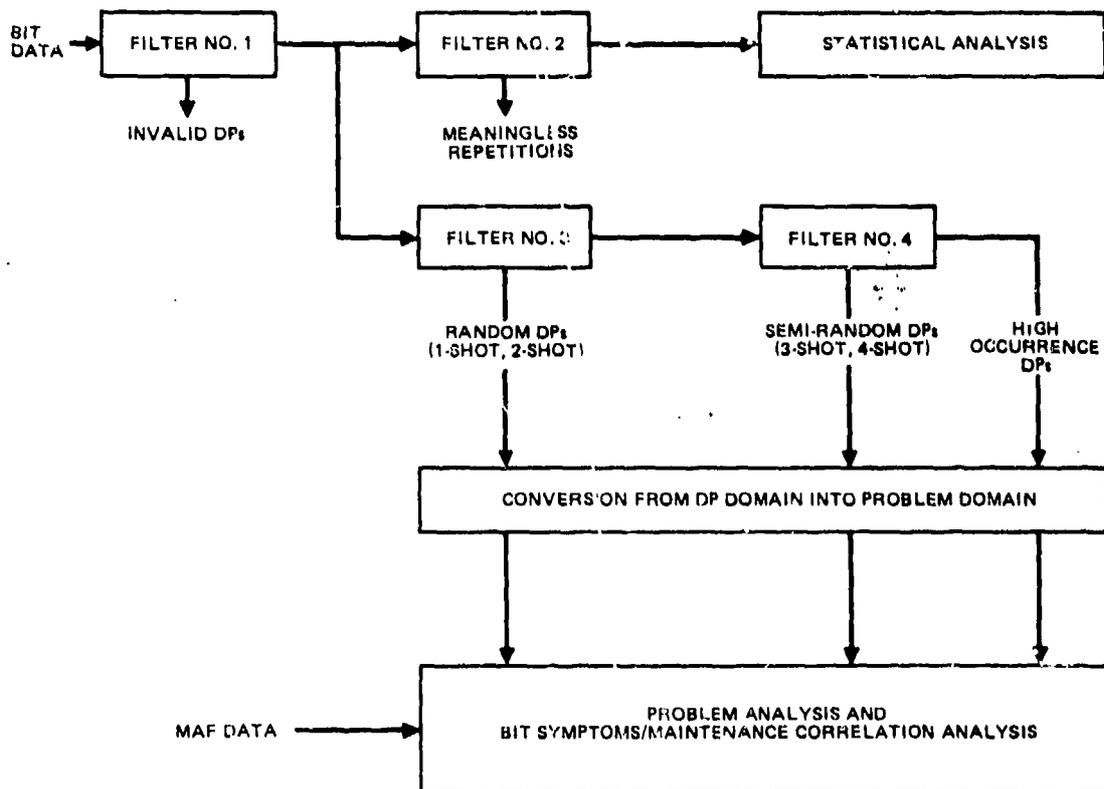
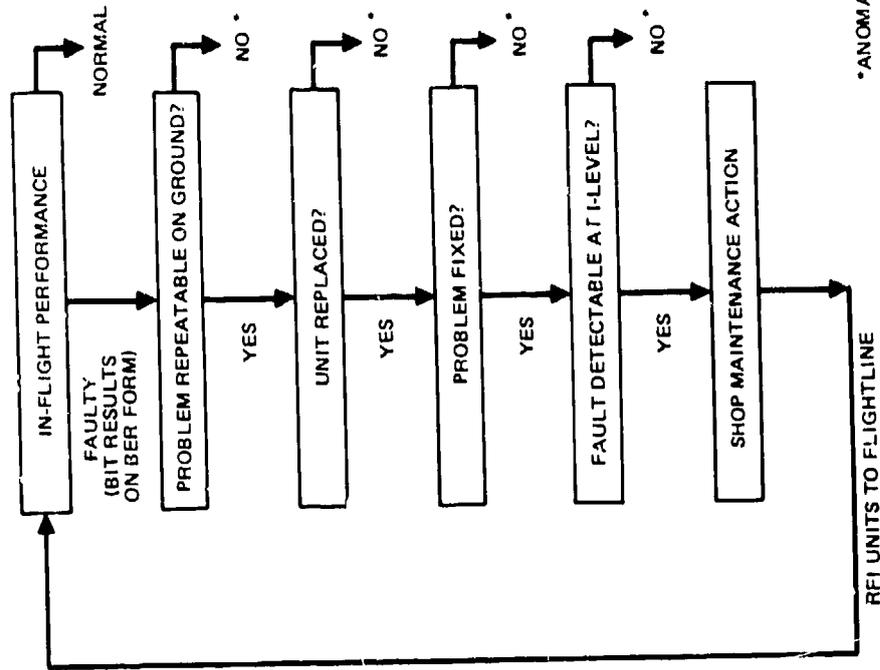


Figure D-2. Analysis Road Map

CORRELATION TOPICS/POTENTIAL BIT FIGURES OF MERIT

- BIT MODE ASSESSMENT VALIDITY
 - FALSE ALARMS (FAIL INDICATIONS VERSUS MISSION SUCCESS)
 - OVERSIGHTS (PASS INDICATIONS VERSUS MISSION FAILS)
- IN-FLIGHT TEST-TO-TEST REPEATABILITY
 - (BER/BER CORRELATION)
- CORRELATION BETWEEN FLIGHT AND GROUND RESULTS
 - EFFICIENCY OF ISOLATING "MODE-X" FAULTS
- COMPARISON BETWEEN BIT-INDICATED AND NON-BIT-INDICATED MAINTENANCE ACTIONS
- CORRELATION BETWEEN BIT INDICATIONS AND O-LEVEL MAINTENANCE SUCCESSES/FAILS
 - (BER/O-LEVEL-MAF CORRELATION)
- CORRELATION BETWEEN O-LEVEL AND I-LEVEL MAINTENANCE SUCCESSES/FAILS
 - (BER/MAF CORRELATION)
- CORRELATION BETWEEN INITIAL IN-FLIGHT SQUAWK AND FAULT FOUND AT I-LEVEL
 - NEXT-FLIGHT VERIFICATION OF MAINTENANCE PROCESS



*ANOMALOUS MAINTENANCE INCIDENTS REQUIRING ANALYSIS

Figure D-3. System Viewpoint of the Maintenance Process

APPENDIX E

SYSTEM 1 SUPPORTING DATA --
SEQUENCE 3 (RADAR CONFIDENCE TEST)
DP STATISTICS AND RANKINGS

TABLE E-1. OCCURRENCE OF BIT DPs AMONG AIRCRAFT*

<u>NUMBER OF DIFFERENT AIRCRAFT EXHIBITING DP</u>	<u>NUMBER OF DIFFERENT DPs OCCURRING ON THE INDICATED NUMBER OF A/C</u>	<u>PROPORTION OF AIRCRAFT EXHIBITING DPs</u>
1	18	0.03
2	16	0.06
3	7	0.10
4	16	0.13
5	19	0.16
6	14	0.19
7	12	0.23
8	9	0.26
9	7	0.29
10	7	0.32
11	7	0.35
12	12	0.39
13	4	0.42
14	3	0.45
15	4	0.48
16	1	0.52
17	4	0.55
18	3	0.58
19	3	0.61
20	6	0.65
21	2	0.68
22	0	0.71
23	1	0.74
24	3	0.77
25	2	0.81
26	0	0.84
27	1	0.87
28	1	0.90
29	0	0.94
30	0	0.97
31	1	1.00
<u>31 DIFFERENT A/C</u>	<u>183 DIFFERENT DPs</u>	

* A DP (Decision Point) is a BIT indication that uniquely identifies a particular BIT check which has been failed. The data in this table illustrate that certain DPs were quite common among all the aircraft in the sample (e.g., one DP occurred on all 31 aircraft) while other DPs were displayed on only a few aircraft. The occurrence of a DP on many different aircraft suggests the hypothesis that these instances represent CAT II false alarms since it seems unlikely that a particular type of hardware failure would have been so widespread.

TABLE E-2. DP LISTING, RANKED BY NUMBER OF OCCURRENCES

<u>Rank by No. of Occurrences</u>	<u>SEQ 3 DP</u>	<u>No. of Occur- rences</u>	<u>No. of A/C</u>	<u>No. of A/C, in %</u>	<u>Rank by A/C %</u>
1	216	223	31	100	1
2	141	167	23	74	9
3	153	142	24	78	6
4	180	138	28	90	2
5	176	129	10	32	-
6	149	128	27	87	3
7	62	115	13	42	36
8	175	107	24	78	7
9	164	99	25	81	5
10	177	98	19	61	20
11	198	92	20	64	17
12	73	87	18	58	22
13	54	86	12	39	38
14	146	83	12	39	47
15	136	71	21	68	10
16	199	70	5	16	-
17	170	65	20	64	16
18	191	65	21	68	11
19	99	64	12	39	42
20	169	64	20	64	15
21	130	63	18	58	23
22	173	61	19	61	19
23	71	60	20	64	12
24	147	59	20	64	13
25	163	54	20	64	14
26	162	51	7	23	-
27	6	50	18	58	21
28	3	48	25	81	4
29	182	47	7	23	-
30	85	44	15	49	29
31	151	44	9	29	-
32	210	44	24	78	8
33	5	41	19	61	18
34	142	41	13	42	-
35	122	39	17	55	25
36	126	39	14	45	34
37	69	38	17	55	24
38	82	34	11	36	52
39	179	34	14	45	35
40	63	33	14	45	33
41	178	32	4	13	-
42	220	31	12	39	50
43	116	30	15	49	30

TABLE E-2 (CONT). DP LISTING, RANKED BY NUMBER OF OCCURRENCES

<u>Rank by No. of Occurrences</u>	<u>SEQ 3 DP</u>	<u>No. of Occur- rences</u>	<u>No. of A/C</u>	<u>No. of A/C, in %</u>	<u>Rank by A/C %</u>
44	212	30	17	55	27
45	8	27	16	52	28
46	78	27	8	26	-
47	214	27	15	49	32
48	88	26	12	39	40
49	160	26	12	39	48
50	56	25	5	16	-
51	65	25	12	39	39
52	59	24	10	32	-
53	60	23	7	23	-
54	70	23	10	32	-
55	83	22	13	42	37
56	134	22	12	39	46
57	184	22	6	19	-
58	93	21	12	39	41
59	96	21	6	19	-
60	138	21	17	55	26
61	152	21	15	49	31
62	209	21	11	36	56
63	9	20	11	36	51
64	133	20	12	39	45
65	211	20	10	32	-
66	89	19	11	36	53
67	102	19	9	29	-
68	76	18	8	26	-
69	131	18	11	36	55
70	183	18	10	32	-
71	190	18	5	16	-
72	95	17	5	16	-
73	101	17	12	39	43
74	115	17	9	29	-
75	215	17	11	36	57
76	111	16	6	19	-
77	123	16	12	39	44
78	124	16	6	19	-
79	21	15	8	26	-
80	103	15	6	19	-
81	105	15	11	36	54
82	127	15	8	26	-
83	128	15	6	19	-
84	158	15	7	23	-
85	86	14	9	29	-
86	90	14	7	23	-

TABLE E-2 (CONT). DP LISTING, RANKED BY NUMBER OF OCCURRENCES

<u>Rank by No. of Occurrences</u>	<u>SEQ 3 DP</u>	<u>No. of Occur- rences</u>	<u>No. of A/C</u>	<u>No. of A/C, in %</u>	<u>Rank by A/C %</u>
87	155	14	9	29	-
88	159	14	9	29	-
89	74	13	10	32	-
90	156	13	8	26	-
91	181	13	12	39	49
92	106	12	8	26	-
93	161	12	5	16	-
94	165	12	8	26	-
95	167	12	10	32	-
96	10	11	6	19	-
97	77	11	8	26	-
98	148	11	7	23	-
99	168	11	5	16	-
100	19	10	9	29	-
101	81	10	5	16	-
102	171	10	6	19	-
103	185	10	5	16	-
104	13	9	8	26	-
105	15	9	7	23	-
106	48	9	2	6	-
107	75	9	4	13	-
108	91	9	3	10	-
109	97	9	5	16	-
110	186	9	5	16	-
111	4	8	4	13	-
112	17	8	7	23	-
113	18	8	7	23	-
114	22	8	6	19	-
115	41	8	4	13	-
116	66	8	6	19	-
117	139	8	4	13	-
118	143	8	5	16	-
119	7	7	6	19	-
120	14	7	7	23	-
121	20	7	6	19	-
122	64	7	4	13	-
123	79	7	4	13	-
124	80	7	4	13	-
125	84	7	7	23	-
126	87	7	5	16	-
127	92	7	4	13	-
128	94	7	2	6	-
129	140	7	5	16	-

TABLE E-2 (CONT). DP LISTING, RANKED BY NUMBER OF OCCURRENCES

<u>Rank by No. of Occurrences</u>	<u>SEQ 3 DP</u>	<u>No. of Occur- rences</u>	<u>No. of A/C</u>	<u>No. of A/C, in %</u>	<u>Rank by A/C %</u>
130	197	7	7	23	-
131	213	7	5	16	-
132	0	6	5	16	-
133	16	6	6	19	-
134	28	6	1	3	-
135	29	6	2	6	-
136	33	6	4	13	-
137	61	6	3	10	-
138	98	6	4	13	-
139	110	6	3	10	-
140	112	6	4	13	-
141	113	6	5	16	-
142	125	6	1	3	-
143	187	6		19	-
144	193	6	2	6	-
145	11	5	3	10	-
146	25	5	2	6	-
147	51	5	4	13	-
148	72	5	5	16	-
149	137	5	4	13	-
150	154	5	5	16	-
151	196	5	5	16	-
152	217	5	2	6	-
153	12	4	3	10	-
154	189	4	3	10	-
155	194	4	4	13	-
156	195	4	4	13	-
157	219	4	3	10	-
158	2	3	2	6	-
159	30	3	2	6	-
160	1	2	2	6	-
161	40	2	1	3	-
162	50	2	1	3	-
163	55	2	2	6	-
164	57	2	2	6	-
165	58	2	2	6	-
166	132	2	2	6	-
167	150	2	1	3	-
168	174	2	2	6	-
169	192	2	2	6	-
170	206	2	2	6	-
171	39	1	1	3	-
172	49	1	1	3	-

TABLE E-2 (CONT). DP LISTING, RANKED BY NUMBER OF OCCURRENCES

<u>Rank by No. of Occurrences</u>	<u>SEQ 3 DP</u>	<u>No. of Occur- rences</u>	<u>No. of A/C</u>	<u>No. of A/C, in %</u>	<u>Rank by A/C %</u>
173	52	1	1	3	-
174	100	1	1	3	-
175	104	1	1	3	-
176	109	1	1	3	-
177	118	1	1	3	-
178	121	1	1	3	-
179	129	1	1	3	-
180	135	1	1	3	-
181	172	1	1	3	-
182	188	1	1	3	-
183	204	1	1	3	-

MISSION
of
Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control Communications and Intelligence (C³I) activities. Technical and engineering support within areas of technical competence is provided to ESD Program Offices (POs) and other ESD elements. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.