

nasa-cr-170,396

NASA-CR-170396
19850006535

Analytic Redundancy Management Mechanization and Flight Data Analysis for the F-8 Digital Fly-By-Wire Aircraft Flight Control Sensors

James C. Deckert

LIBRARY COPY

MAY 7 - 1983



NF02740

LANGLEY RESEARCH CENTER
LIBRARY, NASA
HAMPTON, VIRGINIA

Contract NAS4-2675
April 1983

~~Because of its significant early commercial potential, this information, which has been developed under a U.S. Government program, is being disseminated within the United States in advance of general publication. This information may be duplicated and used by the recipient with the express limitation that it not be published. Release of this information to other domestic parties by the recipient shall be made subject to these limitations.~~

~~Foreign release may be made only with prior NASA approval and appropriate export licenses. This legend shall be marked on any reproduction of this information in whole or in part.~~

~~Review for general release~~

~~12-31-84~~

NASA

National Aeronautics and
Space Administration

Analytic Redundancy Management Mechanization and Flight Data Analysis for the F-8 Digital Fly-By-Wire Aircraft Flight Control Sensors

James C. Deckert
The Charles Stark Draper Laboratory, Inc.
Cambridge, Massachusetts 02139

Prepared for
Ames Research Center
Dryden Flight Research Facility
under Contract NAS4-2675

1983



National Aeronautics and
Space Administration

Ames Research Center
Dryden Flight Research Facility
Edwards, California 93523

X83-10192#

This Page Intentionally Left Blank

TABLE OF CONTENTS

<u>Section</u>	<u>Page</u>
SUMMARY	1
1 INTRODUCTION	3
2 FAULT-DETECTION AND ISOLATION METHODOLOGY	7
2.1 Introduction	7
2.2 Fault Detection Using Direct Redundancy	9
2.3 Fault Isolation Using Analytic Redundancy	11
2.4 False Alarm Protection	15
2.5 Summary	16
3 FAULT ISOLATION TEST DETAILS	17
3.1 Introduction	17
3.2 Rotational Kinematics	21
3.3 Translational Kinematics	26
3.4 Translational Dynamics	31
3.5 Altitude Kinematics	33
3.6 Computational Requirements	36
4 FLIGHT TEST RESULTS	37
5 CONCLUSIONS AND RECOMMENDATIONS	49
REFERENCES	55

LIST OF FIGURES

<u>Figure</u>		<u>Page</u>
2-1	Probability density for average with BFM/2 bias	10
3-1	Block diagram of turbulence estimator	28
4-1	Actual failure of alpha vane two	40
4-2	Pitch rate gyro one simulated 0.5 scale factor	40
4-3	Pitch rate gyro one simulated 0.053 rad/s bias with simulated instrument two misalignment	43
4-4	Lateral accelerometer one simulated 0.3 g bias using 0.9 CYB	43
4-5	Alpha vane one simulated 0.053 rad bias using 0.9 CL	46
4-6	Longitudinal accelerometer one simulated -0.3 g bias with simulated instrument two misalignment	46
4-7	Normal accelerometer one simulated 0.3 g bias	48

LIST OF TABLES

<u>Table</u>	<u>Page</u>
3-1 Phase II Arm signal parameters	18
3-2 Compensation parameters	20
3-3 Rate gyro rotational kinematics test parameters	22
3-4 Attitude signal rotational kinematics test parameters	25
3-5 Turbulence estimator parameters	28
3-6 Accelerometer translational kinematics test parameters	30
3-7 Mach meter translational kinematics test parameters	31
3-8 Lateral accelerometer translational dynamics test test parameters	32
3-9 Alpha vane translational dynamics test parameters	33
3-10 Functional representation of aerodynamic coefficients	34
4-1 Inserted fault average isolation times at trim (seconds)	41

ANALYTIC REDUNDANCY MANAGEMENT
MECHANIZATION AND FLIGHT DATA ANALYSIS
FOR THE F-8 DIGITAL FLY-BY-WIRE AIRCRAFT
FLIGHT CONTROL SENSORS

James C. Deckert

SUMMARY

The details are presented of an onboard digital computer algorithm designed to reliably detect and isolate the first failure in a duplex set of flight control sensors aboard the NASA F-8 Digital Fly-by-Wire aircraft. The algorithm's successful flight test program is summarized, and specific examples are presented of algorithm behavior in response to software-induced signal faults, both with and without aircraft parameter modeling errors.

This Page Intentionally Left Blank

SECTION 1

INTRODUCTION

In recent years the desire to improve aircraft performance and to employ advanced aerodynamic designs has moved operational aircraft control technology into the realm of active control systems. These systems not only reduce pilot workload but, in many cases, are essential for the flight control of the aircraft. Likewise, the evolution of the control-configured vehicle, an aircraft whose aerodynamic design alone is unable to provide satisfactory handling qualities, requires a control system that is automatic and highly reliable. However, individual controller components do not have the reliability, of the order of that of the airframe itself, required for such flight-critical applications. This dictates that a fault-tolerant control system be implemented, utilizing component replication and incorporating some technique for redundancy management.

Although control sensor redundancy management is relatively simple for triplex and higher instrument redundancy, other considerations such as weight, volume, power, and life-cycle costs suggest that the required level of sensor redundancy should be supplied by keeping direct redundancy, i.e., sensor replication, to a minimum and utilizing in its place the analytic redundancy inherent in the various physical relationships among the variables measured by unlike sensors. In June 1975 the Charles Stark Draper Laboratory (CSDL) began a study sponsored by the NASA Langley Research Center (LaRC) to investigate the feasibility of control sensor complement reduction through

the use of analytic redundancy. Specifically, the study was aimed at the problem of isolating the first failure in a duplex subset of the control sensors aboard the NASA F-8 Digital Fly-by-Wire (DFBW) aircraft. Such a technique, if feasible, would allow operational capability following the first sensor failure, in contrast to standard voting techniques for failure isolation that require at least triplex sensors measuring a scalar quantity to provide the same single-failure-operational capability.⁽¹⁾

During the 2-year study period, a preliminary fault-detection and isolation (FDI) algorithm was designed and coded in FORTRAN. The algorithm was successfully tested on the F-8 DFBW aircraft iron-bird simulation facility at LaRC, and it also performed well (following minor modification) on sensor output telemetry data from an early flight test, supplied by the NASA Dryden Flight Research Center (DFRC).^(2,3,4)

At that time, a variety of other approaches to the problem of sensor fault isolation via analytic redundancy had been proposed, including failure-sensitive filters^(5,6,7) designed to enhance failure detectability, multiple-hypothesis techniques^(8,9) involving a bank of filters for a wide class of failure modes, jump process formulations^(10,11) to detect abrupt changes in the system, and innovations-based detection systems.^(12,13,14,15,16) However, it was felt that none of these techniques gave sufficient consideration to robustness in the presences of inevitable modeling errors and allowable sensor errors, and the multiple-hypothesis techniques and several of the innovations-based techniques appeared to possess innate complexity exceeding current-generation-aircraft flight-computer constraints. Therefore, based upon the relative simplicity of the preliminary CSDL FDI algorithm and its successful initial testing, a decision was made by DFRC to sponsor a program to develop an airworthy FDI algorithm, to program the developed algorithm on the F-8 DFBW aircraft computers, and to flight test the algorithm sufficiently to prove the concept of reliable sensor fault isolation through analytic redundancy.

References 17 and 18 document the mathematical development and software specifications, respectively, for the Phase I analytic redundancy management (ARM) algorithm. As will be discussed in the following, the ARM algorithm differs from the earlier approach⁽⁴⁾ in the use of a single, comprehensive test statistic for each suspect sensor in place of the three test statistics required previously for a suspect sensor pair. Although implemented in different ways, both strategies recognize the inherent requirement for the explicit accommodation of the effects of normal sensor errors and modeling errors on analytic redundancy relationships.

The Phase I ARM algorithm was tested on seven flights between September 1979 and February 1980. In order to accommodate the analytic redundancy residual behavior observed on these flights, and also to incorporate more complete knowledge concerning the behavior of the onboard barometric altimeters, the Phase I ARM algorithm was modified to the Phase II version. In addition to parameter value changes, the only major coding changes in the new version involved the addition of first-order filters to several altimeter-related quantities and the addition of filter "traps" to accommodate hysteresis in the altimeters.

The Phase II algorithm was initially tested on three flights between October 1980 and March 1981. Following analysis of the data from these flights, it became apparent that the onboard attitude gyros, which were not representative of currently available control-grade attitude references, did not have performance capabilities consistent with ARM algorithm requirements. A particularly troublesome problem was the inherent difficulty in avoiding inaccurate directional gyro output at high roll angles, reflecting built-in gimbal-lock avoidance, due to the "free-azimuth" configuration of these instruments. Therefore, the duplex vertical gyros and duplex directional gyros were replaced with duplex all-attitude platforms, whose performance proved more than adequate.

Phase II algorithm testing was completed with three flights during June and July, 1981. The analysis of the data from these flights, as well as all previous flights, was facilitated by the use of the ground-edit program, a FORTRAN emulation of the actual flight code. This program was designed to read the sensor data stored on the downlink tape recorder during actual or simulated flights and duplicate the calculations performed by the ARM flight code. The ground-edit program proved useful not only in verifying the flight code implementation but also in validating the effects of proposed parameter value and coding changes on algorithm performance.

The Phase II ARM algorithm performance in isolating simulated sensor failures injected via software during the last three flight tests was excellent with one exception. The algorithm had difficulty isolating a positive normal accelerometer bias with the aircraft in either a positive- or negative-roll-angle high-g turn. Following extensive analysis, the problem was deduced to be a -0.05 bias in each of the two Mach meters, which use common pressure orifices. Running the ground-edit program on the downlink data with the postulated Mach meter biases removed resulted in excellent isolation performance in the previously troublesome situations.

The remainder of this report consists of four sections. Section 2 outlines the basic structure of the ARM algorithm. Section 3 presents the details of the analytic redundancy relationships and error accommodation terms used in the Phase II algorithm. Section 4 presents representative results of the Phase II algorithm performance during the flight tests. Section 5 gives the conclusions and recommendations for future work.

SECTION 2

FAULT-DETECTION AND ISOLATION METHODOLOGY

2.1 Introduction

The analytic redundancy management algorithm is dual mode, with fault detection accomplished by the comparison of like-sensor outputs and fault isolation accomplished using modified sequential probability ratio tests (MSPRT) operating on analytic redundancy residuals. By exploiting the available duplex measurements for fault detection, the algorithm has high computational speed in the normal situation in which there are no faults. Another advantage of this dual-mode structure is that it allows the MSPRT fault-isolation tests to be made quite robust, since fidelity in the analytic redundancy relationships must be maintained for only the short time between detection and isolation.

Each MSPRT resembles a simplified generalized likelihood ratio test (SGLRT),⁽¹⁹⁾ in which the failure time is known and the failure mode is assumed to be a bias of predefined magnitude whose sign is consistent with the difference between the duplex measurements at the time of failure detection. The SGLRT is a simplification of the full-blown GLRT that, in order to isolate a bias jump in a sensor, requires the calculation of the maximum likelihood estimates of both the failure size and failure time. The SGLRT simplifies these calculations by assuming a fixed failure size, and usually the estimate of failure time is restricted to lie within a window in the past. The duplex measurements available on the F-8 allow straightforward deter-

mination of the failure time as discussed below, resulting in further simplification. The major difference between the MSPRT and the SGLRT is the inclusion in the former of what amounts to a decision threshold offset to accommodate the effects of irreducible modeling errors and normal sensor error characteristics in the analytic redundancy relationships. The bias failure hypothesis is used in the absence of detailed failure mode information because bias failures are considered most likely, and also because the resulting tests are quite effective, though not optimal, in isolating other failure types, such as ramps and scale factors.

The determination of whether one or both of the duplex sensor outputs of a particular type will be used to determine an aircraft variable is based upon a hierarchy of signal status levels. In decreasing order of reliability, the four status levels used in the ARM algorithm are: unfailed, provisionally failed, conditionally failed, and unconditionally failed. The average of two signals having equal status is used, while the signal with the better status is used when there is a difference in status.

Conditional and unconditional failure status declarations are made by analytic redundancy MSPRTs as discussed in the following. Provisional failure declarations are made using sensor output selftest logic, which is included in the algorithm to minimize the effect of a hard-failed sensor on critical calculations. An unfailed signal is declared provisionally failed when it differs from its value on the previous sample (and from the present value of its companion signal if the companion's status is unfailed) by a predefined threshold magnitude. If this selftest violation disappears on either of the next two samples, the provisional failure status declaration is removed and the signal reverts to unfailed status.

2.2 Fault Detection Using Direct Redundancy

For fault detection and isolation purposes, it is convenient to define a faulty sensor as one having an output error magnitude larger than a stipulated bias failure magnitude (BFM), and in practice we would like to isolate an instrument having an output error magnitude of the order of BFM. Thus, the BFM for each signal type is chosen to be larger than the available instrument error specifications or the observed output errors in good instruments, and large enough to be isolated by the available analytic redundancy at cruise flight conditions. A signal fault is detected when the moving window average of the output of instrument one minus the output of instrument two is larger than three-quarter BFM in magnitude. This three-quarter BFM threshold results in equal probabilities of detecting a BFM/2 bias (a false alarm) and not detecting a BFM bias (a missed alarm).

Assuming that the noise in the instrument outputs is gaussian and uncorrelated from sample to sample, the number of samples in the moving window, N , may be chosen as follows. The standard deviation of the noise in the average of N samples of the difference in outputs of the two instruments is $\sqrt{2/N} \sigma$, where σ is the standard deviation of the noise on a single instrument. Figure 2-1 depicts the probability density function for the average of N samples in the presence of a BFM/2 bias. The shaded area, P_F , represents the probability that the threshold is exceeded (a false alarm). For a value of P_F of 10^{-4} , gaussian distribution tables indicate that the distance from the threshold to the mean, BFM/4 in this case, should be equal to 3.65 times the standard deviation of the noise. Thus

$$N = 426 \left(\frac{\sigma}{\text{BFM}} \right)^2 \quad (2-1)$$

Because the threshold is chosen midway between BFM and BFM/2, an identical result follows from consideration of missed-alarm probabil-




Figure 2-1. Probability density for average with BFM/2 bias

ity with a failure of BFM present. For this study, sensor output recordings were examined to determine the values of σ for the various sensor types, and integer window lengths were then chosen that most closely satisfied Eq. (2-1). The resulting values of these parameters are shown in Table 3-1.

Following the detection of the failure of one instrument of a pair, fault isolation tests are initiated using analytic redundancy. In addition to the failure indication, the sign of the moving window average of the output of instrument one minus the output of instrument two is also passed to the fault isolation tests. If this sign is positive, either instrument one has a positive error or instrument two has a negative error, and the opposite situation occurs if the sign of the average difference is negative. This failure sign information is utilized by the isolation tests to decrease the amount of processing that must be performed.

2.3 Fault Isolation Using Analytic Redundancy

The SPRT utilizes sequential observations of a process to decide which of two hypothesis concerning the probability distribution of the process is true.⁽²⁰⁾ The SPRT is independent of the a priori probabilities of the two hypotheses, and minimizes the average number of observations necessary to reach a decision while meeting prespecified misclassification probabilities.⁽²¹⁾ Because of these desirable characteristics and the simple form of the test when gaussian errors are assumed, the SPRT is an ideal candidate for use with analytic redundancy for fault isolation. In particular, assuming that the process being observed is the difference between the output of one suspect sensor and a synthesized output using analytic redundancy, appealing choices for the two hypotheses are that the process has a mean equivalent to a BFM-sized bias (i.e., the instrument has failed) or that the process has zero mean (i.e., the instrument is unfailed).

Specifically, assume that the noises on the residual process γ_k for instrument j ($j = 1$ or 2) at time t_k are independent for all k ; that either the failure hypothesis, H_1 , or the no-failure hypothesis, H_2 , is true; and that H_1 and H_2 are the following:

H_1 : γ_k^j is gaussian with variance σ^2 and mean m_k^j

H_2 : γ_k^j is gaussian with variance σ^2 and mean 0

It follows that the log likelihood ratio (LLR) of the joint probability density function for n successive observations conditioned on H_1 , divided by the joint probability density function for n successive observations conditioned on H_2 , is given by

$$\mu_n^j = \sum_{k=1}^n \frac{m_k^j}{\sigma^2} \left(\frac{m_k^j}{2} - \gamma_k^j \right) \quad (2-2)$$

(Note that σ is a general variable and the values in Eq. (2-1), Eq. (2-2), and subsequent equations are not necessarily equal.)

Defining P_m as the probability of choosing H_2 when H_1 is true, and P_f as the probability of choosing H_1 when H_2 is true, then the SPRT optimal decision rule is given by

$$\begin{aligned} \mu_n^j &\leq \delta && \text{accept } H_1 \\ \delta &< \mu_n^j < \eta && \text{take another sample} \\ \eta &\leq \mu_n^j && \text{accept } H_2 \end{aligned} \tag{2-3}$$

where

$$\begin{aligned} \delta &= -\ln[(1 - P_m)/P_f] \\ \eta &= -\ln[P_m/(1 - P_f)] \end{aligned} \tag{2-4}$$

Note that if the LLR is between the two thresholds, a choice of hypothesis that meets the specified acceptable misclassification probabilities P_m and P_f cannot be made, and another sample must be taken.

Unfortunately, such factors as allowable biases on unfailed sensors, errors in the sensor input/output models, and parameter uncertainties in analytic redundancy relationships all contribute to low-frequency errors in analytic redundancy residual processes for unfailed sensors. Direct application of the above SPRT to such a process may result in the acceptance of the failure hypothesis in spite of the fact that the sensor is operating within acceptable tolerances.

The MSPRT is a fault-isolation test that systematically accommodates those irreducible factors contributing to low-frequency analytic redundancy residual errors that cannot be explicitly removed by modeling. The ARM algorithm utilizes the MSPRT to make conditional and unconditional failure status declarations following fault detection. Underlying the technique is the assumption that either H_1 or H_2 for suspect sensor j is true. Although a straightforward approach would be to design a decision rule such as Eq. (2-3) for each sensor and then declare a failure when either hypothesis for either sensor was accepted, a more conservative approach has been chosen that avoids the tenuous situation of inferring that one sensor has failed merely because its companion sensor appears to be unfailed.

Consistent with the preceding discussion, following fault detection at time t_1 , the MSPRT test statistic, the modified log likelihood ratio (MLLR), is defined at time t_n for suspect sensor j as

$$u_n^j = \sum_{k=1}^n \left[\frac{m_k^j}{\sigma^2} \left(\frac{m_k^j}{2} - \gamma_k^j \right) + \frac{|m_k^j|}{\sigma^2} E_k \right] \quad (2-5)$$

and the following decision rule is used

$$\begin{aligned} u_n^j &\leq \delta && \text{declare instrument } j \text{ unconditionally failed,} \\ &&& \text{terminate the test} \\ \delta &< u_n^j < 0 && \text{declare instrument } j \text{ conditionally failed,} \\ &&& \text{take another sample} \\ 0 &\leq u_n^j && \text{take another sample} \end{aligned} \quad (2-6)$$

In Eq. (2-6), the mean m_k^j is computed assuming a BFM-sized bias in sensor j consistent with the direct redundancy moving window average

at the detection time. The negative threshold δ in Eq. (2-6) is the original SPRT threshold calculated using prespecified misclassification probabilities via Eq. (2-4).

The last term in the summation in Eq. (2-5) differentiates the MLLR from the LLR of the standard SPRT, and represents the contribution of a postulated worst-case residual error magnitude at time t_k , E_k , to LLR calculation. Eqs. (2-5) and (2-6) indicate that the MSPRT is in essence a one-sided SPRT with a threshold offset arising from the worst-case error term. It follows that so long as the threshold offset is conservative, the misclassification probabilities for the MSPRT will be no larger than those specified to determine the original SPRT threshold.⁽¹⁷⁾ The choice of the worst-case error magnitude for each analytic redundancy test requires considerable engineering judgement. An optimistic choice lowers the reliability of the test while an overly pessimistic choice may result in prohibitively long isolation times, although the inclusion of conditional failure declaration using a relaxed test criterion tends to lower the mean isolation times seen by the control system without corresponding increases in the ultimate misisolation probabilities. Specific choices for the worst-case error terms are discussed in Section 3.

In addition to the decision rule of Eq. (2-6), the ARM algorithm avoids open-ended isolation tests by utilizing an elapsed time limit (ETL) for each sensor type. If ETL is reached before an unconditional failure has been declared, the detected fault indication is removed, isolation computations cease, and the direct redundancy detection process is reinitiated. Because failure observability and worst-case error magnitude are often maneuver dependent, pilot response to notification that ETL for a sensor type has been reached could result in an enhanced fault isolation environment during the subsequent isolation period. Alternatively, reaching ETL could initiate hardware selfcheck routines.

2.4 False Alarm Protection

In order to provide increased false alarm protection, a direct redundancy SPRT is initiated for the suspect sensor type following fault detection. The process examined by the DRSPRT is the output of instrument one minus the output of instrument two, the same process whose moving window average triggered the fault detection. The DRSPRT test statistic, the DRLLR, is given at time t_n as

$$DRLLR_n = \sum_{k=1}^n \frac{m}{\sigma^2} \left[\frac{m}{2} - (o_k^1 - o_k^2) \right] \quad (2-7)$$

where o_k^1 and o_k^2 denote the outputs at time t_k of suspect instruments one and two, respectively. The mean m in Eq. (2-7) has a magnitude equal to BFM, and its sign is the sign of the moving window average at the time of fault detection. Whenever the DRLLR crosses the specified positive threshold, indicating that a difference between the two instrument outputs of BFM magnitude with the stipulated sign does not exist, the analytic redundancy fault isolation tests are terminated and the direct redundancy fault detection test is reinitiated.

In addition to providing protection against false alarms, the DRSPRT also performs a "rapid reset" function when the failure's observability decreases following fault detection, as an alternative to waiting until ETL is reached. As will be illustrated in flight test results shown in Section 4, this reset capability is particularly useful as an aid to isolating scale-factor failures when the instrument's input changes sign following fault detection, and it allows the simple bias-failure-hypothesis-based isolation tests to perform quite effectively in these situations.

2.5 Summary

To summarize the ARM FDI process, each signal type utilizes a threshold test on the moving window average of the difference in the duplex signals to detect a fault. Following fault detection, one MLLR is computed via Eq. (2-5) for each suspect sensor for each form of analytic redundancy used, and the MSPRT threshold logic of Eq. (2-6) is applied to the lowest MLLR. This process is repeated until an unconditional failure is declared or ETL is reached. Additionally, a direct redundancy LLR is computed as shown in Eq. (2-7) following fault detection to provide false alarm protection. Finally, signal selftest is continuously performed for all signals having unfailed status.

SECTION 3

FAULT ISOLATION TEST DETAILS

3.1 Introduction

The ARM algorithm monitors ten duplex instruments aboard the F-8 DFBW aircraft: longitudinal accelerometer, lateral accelerometer, normal accelerometer, roll rate gyro, pitch rate gyro, yaw rate gyro, attitude platform, barometric altimeter, Mach meter, and alpha vane. (Although the accelerometer and rate gyro complement is triplex, only a duplex subset is utilized by the ARM algorithm.) Additionally, a simplex beta vane is used in some calculations, but not monitored for failure. Each attitude platform gives outputs of Euler roll angle, ϕ , pitch angle, θ , and azimuth angle, ψ , and these outputs are considered to be three independent signal types although, in practice, the failure of one channel would probably dictate the failure of the unit.

Table 3-1 indicates relevant Phase II ARM parameters for the 12 signal types being monitored. As is discussed later in this section, the mechanization of the MSPRTs for the Mach meters, altimeters, and attitude platforms is such that there is no ETL or DRLLR associated with these signal types. The ARM algorithm sample period is 0.06 second.

From the large number of analytic redundancy relationships available, practical considerations and aircraft-specific signal-to-noise values reduce the number used in the ARM algorithm to four general types: rotational kinematics, altitude kinematics, translational

Table 3-1. Phase II ARM signal parameters.

Signal Type	Symbol	BFM	Selftest Threshold	DRLLR Variance	ETL (Samples)	Window Length
Mach	M	0.05	0.1	---	---	3
Altitude	h	76.2 m	304.8 m	---	---	5
Angle of attack	α	0.035 rad	0.1 rad	0.00005 rad^2	250	9
Longitudinal acceleration	Ax	0.2 g	0.5 g	0.002 g^2	1000	10
Lateral acceleration	Ay	0.2 g	0.5 g	0.002 g^2	500	10
Normal acceleration	Az	0.2 g	1.0 g	0.002 g^2	500	10
Roll rate	p	0.087 rad/s	1.0 rad/s	$0.00005 \text{ rad}^2/\text{s}^2$	133	3
Pitch rate	q	0.035 rad/s	0.25 rad/s	$0.00001 \text{ rad}^2/\text{s}^2$	133	2
Yaw rate	r	0.035 rad/s	0.25 rad/s	$0.00001 \text{ rad}^2/\text{s}^2$	133	2
Roll angle	ϕ	0.087 rad	0.6 rad	---	---	6
Pitch angle	θ	0.087 rad	0.2 rad	---	---	6
Yaw angle	ψ	0.087 rad	0.2 rad	---	---	6

kinematics, and translational dynamics. In the following four sections, the analytic redundancy residual equations and worst-case error terms are discussed for each sensor type, grouped by the type of analytic redundancy employed. The lateral and normal accelerometers employ two types of analytic redundancy each while the remaining signal types employ one each.

Before proceeding to the details of the analytic redundancy calculations, the compensation performed on particular sensor outputs will be discussed. The compensated instruments include the accelerometers, the alpha and beta vanes, and the Mach meters.

The accelerometers are compensated in order to transform their readings to the acceleration at the aircraft center of mass. Using a tilde (-) to denote uncompensated sensor output, the compensated reading of the j^{th} longitudinal, lateral, and normal accelerometer are computed as follows

$$A_x^j = \tilde{A}_x^j + l_x (\dot{q}^2 + r^2) \quad (3-1)$$

$$A_y^j = \tilde{A}_y^j - l_x (\dot{r} + pq) + l_y \dot{p}^2 + l_z \dot{p} \quad (3-2)$$

$$A_z^j = \tilde{A}_z^j + l_x (\dot{q} - pr) - l_y \dot{p} + l_z \dot{p}^2 \quad (3-3)$$

where l_x , l_y , and l_z denote the displacement in body axes of the sensor package from the center of mass. Eq. (3-1) through (3-3) represent the dominant terms from the exact expressions,⁽²²⁾ consisting of those terms involving the longest distance, l_x , the roll rate, or the roll acceleration. The angular acceleration terms in these compensation equations are computed by simple back-differencing in the ARM algorithm.

The alpha and beta vane readings are compensated for their distance from the center of mass and for bias, with the compensated readings of the single beta vane and the j^{th} alpha vane computed as

$$\beta = \tilde{\beta} - d \frac{r}{V} - \beta_b \quad (3-4)$$

$$\alpha^j = \tilde{\alpha}^j + d \frac{q}{V} - \alpha_b \quad (3-5)$$

where d is the longitudinal distance from the center of mass to the vanes and V is the total air-relative velocity computed using the Mach number and the velocity of sound computed as a function of altitude.⁽¹⁸⁾ The compensated readings of the j^{th} Mach meter and the rudder position transducer are computed as

$$M^j = \tilde{M}^j - M_b \quad (3-6)$$

$$R = \tilde{R} - R_b \quad (3-7)$$

The values of the various compensation parameters mentioned above are given in Table 3-2. Supersonic flight is defined as the region where the first-order-filtered, uncompensated, voted Mach signal exceeds 1.0. The time constant of this filter is 0.5 second.

Table 3-2. Compensation parameters.

l_x (m)	l_y (m)	l_z (m)	d (m)	β_b (rad)	R_b (rad)	M_b		α_b (rad)	
						Super-sonic	Other-wise	Super-sonic	Other-wise
6.42	0.31	0.17	11.0	-0.0087	0.0209	-0.01	-0.05	0.014	0

In addition to these compensations, each rate gyro is compensated for bias by subtracting a bias estimate that is updated using analytic redundancy residuals. This process is discussed in the next section.

3.2 Rotational Kinematics

Rotational kinematics (RK) is used for fault isolation in the rate gyros and attitude platforms. The roll, pitch, and yaw rate gyros provide measurements of the aircraft body rates p , q , and r about the aircraft x , y , and z axes, respectively; and these body rates are related to the rates of change of the Euler angles measured by the attitude platforms. Thus, following a rate gyro fault detection at time t_1 , the RK residual for instrument j of the suspect type is calculated at general time t_k using the appropriate equation from the following

$$\gamma_k(p^j) = \sum_{i=1}^k \{ \bar{p}_i^j T - [\phi_i - \phi_{i-1} - (\psi_i - \psi_{i-1}) \overline{\sin \phi_i}] \} \quad (3-8)$$

$$\begin{aligned} \gamma_k(q^j) = \sum_{i=1}^k \{ \bar{q}_i^j T - [(\theta_i - \theta_{i-1}) \overline{\cos \phi_i} \\ + (\psi_i - \psi_{i-1}) \overline{\cos \theta_i} \overline{\sin \phi_i}] \} \end{aligned} \quad (3-9)$$

$$\begin{aligned} \gamma_k(r^j) = \sum_{i=1}^k \{ \bar{r}_i^j T - [-(\theta_i - \theta_{i-1}) \overline{\sin \phi_i} \\ + (\psi_i - \psi_{i-1}) \overline{\cos \theta_i} \overline{\cos \phi_i}] \} \end{aligned} \quad (3-10)$$

In Eq. (3-8) - (3-10) and all subsequent equations, T is the ARM sample period of 0.06 second and an overbar indicates that the quantity represents the average of its present and previous sample values. This averaging is used to reduce computational errors during high angular rate maneuvers, and the forms of Eq. (3-8) - (3-10) avoid differentiation of the noisy attitude measurements. Variables with no "j" superscript represent voted values obtained using the status level logic discussed in Section 2.1.

At every sample time t_k , following fault detection, the residual for each suspect rate gyro is used to update its MLLR using Eq. (2-5), where the mean has magnitude equal to the rate gyro BFM times ($t_k - t_0$), the variance reflects attitude gyro noise variance, and the worst-case error is computed as the sum of the magnitudes of terms reflecting initial attitude gyro noise, roll attitude bias times pitch rate, rate gyro misalignment, and rate gyro scale factor error. Table 3-3 indicates the values of these parameters in the Phase II ARM code.

Table 3-3. Rate gyro rotational kinematics test parameters.

Rate Gyro Type	Variance (rad ²)	Initial Attitude Noise (rad)	Scale Factor Error	Roll Attitude Bias (rad)	Cross-axis Misalignment (rad)
Roll	0.0004	0.01	0.05	---	0.02
Pitch	0.0004	0.005	0.05	0.057	0.02
Yaw	0.0004	0.005	0.05	0.057	0.02

As was done for all signal types, the contributions to be used in the rate gyro RK MLLR worst-case error terms were initially chosen by examining the variational equations for the analytic redundancy residuals⁽¹⁷⁾ and applying engineering judgement. Naturally, the

final forms of the error terms and their coefficient magnitudes were strongly influenced by the flight test data.

Because of the sensitivity of the translational kinematics relationship, used for longitudinal and normal accelerometer fault isolation, to angular rate estimation errors, the rotational kinematics residuals given by Eq. (3-8) through (3-10) are used to estimate biases in the individual rate gyros. At the beginning of each cycle, these estimates are subtracted from the sensor outputs before they are used in any calculations. The general form of the bias update equation at time t_k is

$$\text{bias}_k = \text{bias}_{k-1} + 0.001 \gamma/T \quad (3-11)$$

reflecting an estimator time constant of approximately 2 minutes. These rate gyro bias update equations are bypassed whenever the roll or pitch attitude angle exceeds 0.2 rad in magnitude. During flight tests, the bias estimators have accurately compensated for biases of the order of 0.008 rad/s observed in the pitch rate gyros.

Treating the roll, pitch, and yaw attitude signals from a single platform as independent signal types, the RK residual for signal j of an attitude type (ϕ, θ, ψ) having a detected fault is calculated at time t_k as follows

$$\gamma_k(\phi^j) = \sum_{i=1}^k \{ \phi_i^j - \phi_{i-1}^j - [\bar{p}_i T + (\psi_i - \psi_{i-1}) \overline{\sin \theta_i}] \} \quad (3-12)$$

$$\gamma_k(\theta^j) = \sum_{i=1}^k \{ \theta_i^j - \theta_{i-1}^j - T [\bar{q}_i \overline{\cos \phi_i} - \bar{r}_i \overline{\sin \phi_i}] \} \quad (3-13)$$

$$\gamma_k(\psi^j) = \sum_{i=1}^k \{ \psi_i^j - \psi_{i-1}^j - [(\phi_i - \phi_{i-1} - \bar{p}_i^T) \cdot \overline{\sin \theta_i} + T(\bar{q}_i \overline{\sin \phi_i} + \bar{r}_i \overline{\cos \phi_i}) \overline{\cos \theta_i}] \}$$

(3-14)

In order to calculate the attitude gyro MLLRs, it is necessary to have stored the instantaneous RK residuals, corresponding to the terms enclosed in braces in Eq. (3-12) through (3-14), in a moving window of the same length as the detector window for each signal type. At the time of fault detection, the instantaneous residual window for each suspect signal is processed using the appropriate equation from Eq. (3-12) through (3-14) to form the MLLR residual for each window time, where t_1 corresponds to the time associated with the oldest window element. The MLLRs are computed at each intermediate window time using Eq. (2-5). The MLLR mean has BFM magnitude, the variance is of the order of the attitude gyro noise variance, and the worst-case error is the sum of the magnitudes of terms reflecting initial attitude gyro error, rate gyro bias, rate gyro misalignment, and roll rate gyro scale factor error. The values of the attitude platform RK test parameters for the Phase II algorithm are given in Table 3-4.

After processing the entire window of instantaneous residuals and calculating the MLLRs corresponding to the present time, the threshold logic of Eq. (2-6) is applied to the lower MLLR. If no unconditional failure declaration is made, the detected failure indication is removed and the detection process proceeds smoothly on the next sample.

The framework used for the attitude signals of processing a stored window of instantaneous residuals before applying the MSPRT threshold logic is also utilized for the Mach meters and altimeters, and these five signal types do not require direct redundancy LLR cal-

Table 3-4. Attitude signal rotational kinematics test parameters.

Signal Type	Variance (rad ²)	Inertial Attitude Noise (rad)	Roll Rate Bias (rad/s)	Pitch Rate Bias (rad/s)	Yaw Rate Bias (rad/s)	Misalignment Toward Roll Axis (rad)	Misalignment Toward Pitch Axis (rad)	Misalignment Toward Yaw Axis (rad)	Roll Rate Scale Factor Error
Roll	0.000025	0.01	0.031	---	---	---	0.02	---	0.02
Pitch	0.000025	0.005	---	0.02	0.02	0.02	---	0.02	---
Yaw	0.000025	0.005	0.031	0.02	0.02	0.02	0.02	---	---

culations since the isolation process does not extend beyond the time of detection. Additionally, the MLLR for any signal of these five types is reset to zero whenever it becomes positive. This is done to accommodate the uncertainty in failure time within the window, and requires a slightly higher magnitude threshold for the MSPRTs for these sensor types.⁽¹¹⁾ For the 10^{-4} misclassification probabilities used in this study, the threshold δ in Eq. (2-6) is -11.4 for these five signal types and -9.2 for the others.

In order to make the fault detection tests consistent with the fault identification tests, which are sensitive only to bias changes in the attitude, Mach, and altimeter signals, the elements whose moving window averages are used for fault detection for these five signal types are calculated as the deviations in output difference from the difference most recently discarded from the window.

3.3 Translational Kinematics

Translational kinematics (TK) refers to the redundancy between changes in aircraft air-relative velocity measured by the air data sensors and changes in aircraft velocity obtained by integrating the acceleration computed using the inertial sensor outputs. The ARM algorithm utilizes TK to isolate faults in the longitudinal accelerometers, normal accelerometers, and Mach meters.

The TK residual for longitudinal accelerometer j at time t_k , following a detected fault at time t_1 , is given by

$$\begin{aligned} \gamma_k(Ax^j) = & \sum_{i=1}^k T[Ax_i^j - \overline{\sin \theta_i} g \\ & + V_{s_i} M_i (r_i \sin \beta_i - q_i \sin \alpha_i)] \\ & - [V_{s_k} M_k \cos \alpha_k - V_{s_0} M_0 \cos \alpha_0] \end{aligned} \quad (3-15)$$

where g is the acceleration of gravity, V_s is the speed of sound, periodically recomputed as a function of altitude, and α , β , and M are the angle of attack, sideslip angle, and Mach number, respectively. (In Eq. (3-15) and all subsequent equations involving air-relative velocity, $\cos \beta$ is approximated as unity and therefore does not appear explicitly. Similarly, with α and β measured in radians, the ARM algorithm approximates $\sin \beta$, $\sin \alpha$, and $\cos \alpha$ as β , α , and $(1 - \alpha^2/2)$, respectively.) In the MLLR calculations, the mean at time t_k has magnitude equal to the longitudinal accelerometer BFM times $(t_k - t_0)$ the variance reflects air data sensor noise, and the worst-case error is the sum of the magnitudes of terms reflecting initial air data noise, a wind-shear doublet, misalignment of the suspect accelerometer, and transonic Mach meter behavior.

The TK residual for normal accelerometer j at time t_k , following fault detection at time t_1 , is given by

$$\begin{aligned} \gamma_k(Az^j) = & \sum_{i=1}^k T[Az_i^j - \overline{\cos \theta_i} \overline{\cos \phi_i} g \\ & + V_{s_i} M_i (q_i \cos \alpha_i - p_i \sin \beta_i)] \\ & - [V_{s_k} M_k \sin \alpha_k - V_{s_0} M_0 \sin \alpha_0] \end{aligned} \quad (3-16)$$

The MLLR mean has magnitude equal to the normal accelerometer BFM times $(t_k - t_0)$, the variance reflects air data sensor noise, and the worst-case error is the sum of the magnitudes of terms reflecting initial air data noise, a wind-shear doublet, suspect accelerometer scale-factor error, and pitch rate gyro scale-factor error.

The wind-shear doublet magnitude and MLLR variance for the longitudinal and normal accelerometers each assume one of two values de-

pending upon the binary output of a wind turbulence filter operating on lateral channel TK residuals, analogous to Eq. (3-15) and (3-16). These residuals are high-pass filtered, squared, low-pass filtered, and then passed through a two-level turbulence flag logic with hysteresis. Figure 3-1 indicates a block diagram of the calculation of the wind turbulence estimate, VAREST, and the turbulence flag logic. Table 3-5 shows the values of the associated parameters, with the values of K_1 and K_2 reflecting equal 5 second time constants for the high-pass and low-pass filters.⁽⁴⁾

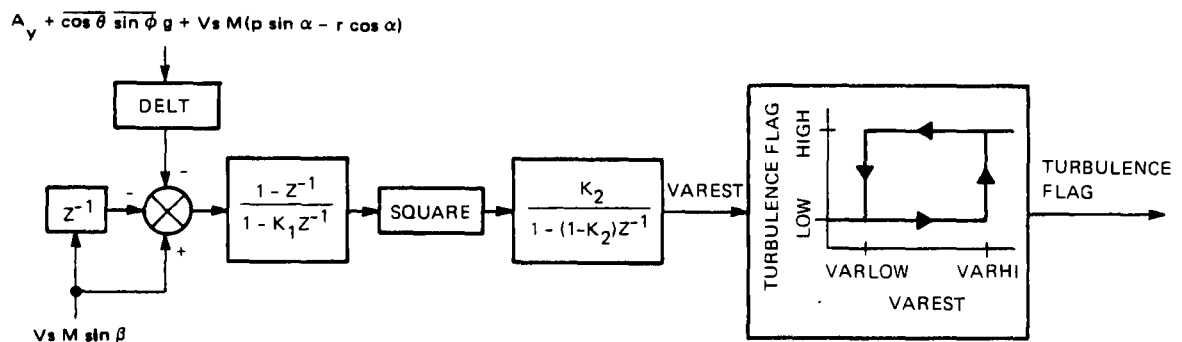


Figure 3-1. Block diagram of turbulence estimator

Table 3-5. Turbulence estimator parameters

K_1	K_2	VARLOW (m^2/s^2)	VARHI (m^2/s^2)
0.9881	0.0119	0.464	0.743

Table 3-6 lists the values of the longitudinal and normal accelerometer TK test parameters used in the Phase II ARM algorithm. A scale-factor error term for pitch rate is employed instead of a rate bias term in the normal accelerometer test for three reasons. First, it reflects the excellent performance of the rate gyro bias estimators. Second, it reflects the existence of an observed scale-factor error in the number two pitch rate gyro of approximately -0.05. Third, the form of this term also allows accommodation of uncompensated Mach bias. The transonic region is defined as values of the filtered, uncompensated, voted Mach signal between 0.92 and 1.04.

The TK residual for Mach meter j is calculated at time t_k as

$$\begin{aligned} \gamma_k(M^j) = & \sum_{i=1}^k \{ V s_i M_i^j \cos \alpha_i - V s_{i-1} M_{i-1}^j \cos \alpha_{i-1} \\ & - T [A x_i - \overline{\sin \theta_i} g + V s_i M_i^j (r_i \sin \beta_i - q_i \sin \alpha_i)] \} \end{aligned} \quad (3-17)$$

As for the attitude gyros discussed earlier, the instantaneous residuals for each Mach meter, the terms in braces in Eq. (3-17), are stored in a moving window. At the time a Mach meter fault is detected, each sensor's instantaneous residual window is processed using Eq. (3-17) to compute its TK residual at the intermediate window times, and simultaneously Eq. (2-5) is used to compute each Mach meter's MLLR, with the threshold logic of Eq. (2-6) applied to the lower MLLR following complete window processing. The MLLR mean has magnitude equal to the Mach meter BFM times $V s \cos \alpha$ at the time of detection, the variance reflects the effect of Mach meter noise, and the worst-case error is the sum of terms arising from initial Mach meter error and acceleration uncertainty. Table 3-7 indicates the values of these parameters used in the Phase II code.

Table 3-6. Accelerometer translational kinematics test parameters.

Type	Variance (m ² /s ²)		Initial Noise (m/s)	Scale Factor Error	Alignment Error (rad)	Transonic Error Rate (g)	Supersonic Error Jump (m/s)	Pitch Rate Scale Factor Error	Doublet		
	Low Turb.	High Turb.							Duration (s)	Magnitude (g)	
										Low Turb.	High Turb.
Ax	9.29	33.4	1.52	---	0.03	0.031	-30.5	---	12	0.155	0.31
Az	9.29	33.4	1.52	0.02	---	---	---	0.035	12	0.155	0.31

Table 3-7. Mach meter translational kinematics test parameters.

Variance (m ² /s ²)	Acceleration Error (m/s ²)	Initial Noise (m/s)	
		Transonic	Otherwise
27.9	4.0	30.5	1.52

3.4 Translational Dynamics

Translational dynamics (TD) refers to the redundancy between the acceleration of the aircraft measured by the accelerometers and the acceleration predicted by stored aerodynamic coefficient functions using air data sensor measurements. TD residuals are used by the ARM algorithm to isolate failures in the lateral accelerometers and alpha vanes.

The TD residual for lateral accelerometer j at time t_k is given by

$$\gamma_k(Ay^j) = Ay_k^j - (CYB_k \beta_k + CYDR_k R_k) Q_k S/m_k \quad (3-18)$$

where CYB and $CYDR$ are stored lateral coefficient functions of Mach and alpha, Q is computed dynamic pressure, S is the surface area of the wing, and m is the estimated aircraft mass. The absence of fuel tank level measurement inputs to the onboard computers precludes automatic mass estimate update. As a compromise between accuracy and operational complexity, the current technique requires pilot selection of which of three mass estimates is used on the basis of cockpit (or telemetry) indications of fuel remaining.

On every sample following fault detection, the residuals given by Eq. (3-18) are used in Eq. (2-5) to compute the lateral accelerometer TD MLLRs. The TD MLLR mean has lateral accelerometer BFM magni-

tude, the variance reflects air data noise, and the worst-case error is the sum of the magnitudes of terms reflecting the effects of beta vane bias, lateral accelerometer misalignment, neglected lateral coefficients, and scale-factor error in the computed aerodynamic side-force. The values of these parameters used in the Phase II ARM code are shown in Table 3-8.

Table 3-8. Lateral accelerometer translational dynamics test parameters.

Variance (m ² /s ⁴)		Beta Bias (rad)	Misalignment (rad)	Scale Factor Error	Neglected Coefficients	
Low Turbulence	High Turbulence				p/V (m)	Aileron
0.93	3.72	0.002	0.02	0.2	1.63	0.052

The TD residual for alpha vane j at time t_k is given by

$$\gamma_k(\alpha^j) = -(L_k^j \cos \alpha_k^j + D_k^j \sin \alpha_k^j)/m_k - Az_k \quad (3-19)$$

where the lift, L, and drag, D, are computed using each alpha vane output individually in stored functions of Mach, elevator position, and angle of attack. The alpha vane TD MLLR mean has magnitude equal to the alpha vane BFM times the magnitude of the computed TD residual gradient, the variance assumes one of two values depending upon indicated turbulence level, and the worst-case error is the sum of the magnitudes of terms reflecting the effects of normal accelerometer scale factor error and aerodynamic coefficient error. Table 3-9 indicates the values of these parameters used in the Phase II code. The "good fit" region for the aerodynamic coefficient error contribution for alpha vane j is currently defined by the following two inequalities

$$-0.175 \leq \delta_e \leq 0.0872$$

$$-0.07 \leq \alpha^j \leq (0.42 - 0.232 M) \quad (3-20)$$

where δ_e is the measured elevator position in radians. The total contribution of aerodynamic coefficient error is currently calculated as 0.6 times the sum of the two individual vane contributions.

Table 3-9. Alpha vane translational dynamics test parameters.

Variance (m ² /s ⁴)		Accelerometer Scale Factor Error	Aerodynamic Coefficient Contribution (m/s ²)		
Low Turbulence	High Turbulence		Good Fit	Otherwise	
				M <u>≤</u> 0.8	M > 0.8
18.6	37.2	0.1	M ² +0.305M	0.3048 + (5.52M-1.66) ²	19.5

Table 3-10 lists the lift, drag, and sideforce aerodynamic coefficient functional representations. The lift coefficient is formed by linear interpolation between the values calculated at the two Mach break points bracketing the Mach estimate; there are ten Mach breakpoints between 0.18 and 1.9. The complete drag coefficient is formed as the sum of CD and CDDE.

3.5 Altitude Kinematics

Altitude kinematics (AK) refers to the redundancy between the changes in altitude measured by the altimeters and changes in altitude computed from the vertical acceleration measured by the accelerometers. This vertical acceleration at time t_1 is computed as

Table 3-10. Functional representation of aerodynamic coefficients.

Aerodynamic Coefficient	Polynomial	Range
CL	For Mach = M_j , $j = 1, 10$	
	$CL_\ell = (a_0 + a_1\alpha) + \delta_e(a_2 + a_3\alpha)$	$\alpha \leq 0.1745$
	$CL_m = (a_0 + a_1\alpha) + \delta_e(a_2 + a_3\alpha)$	$0.1745 < \alpha < 0.2094$
	$CL_h = (a_0 + a_1\alpha + a_2\alpha^2) + \delta_e(a_3 + a_4\alpha)$	$0.2094 \leq \alpha$
CD	$CD_\ell = (a_0 + a_1\alpha + a_2\alpha^2) + M(a_3 + a_4\alpha + a_5\alpha^2)$	$0.18 \leq M \leq 0.9$
	$CD_m = (a_0 + a_1\alpha + a_2\alpha^2) + M(a_3 + a_4\alpha + a_5\alpha^2)$	$0.9 < M < 1.2$
	$CD_h = (a_0 + a_1\alpha + a_2\alpha^2) + M(a_3 + a_4\alpha)$	$1.2 \leq M \leq 1.9$
	$+ M^2(a_5 + a_6\alpha) + M^3a_7$	
CDDE	$CDDE = (a_0 + a_1\alpha + a_2\alpha^2 + a_3\alpha^3)$ $+ \delta_e(a_4 + a_5\alpha + a_6\alpha^2 + a_7\alpha^3)$ $+ \delta_e^2(a_8 + a_9\alpha + a_{10}\alpha^2 + a_{11}\alpha^3)$ $+ \delta_e^3(a_{12} + a_{13}\alpha + a_{14}\alpha^2)$	
CYB	$CYB_\ell = (a_0 + a_1\alpha + a_2\alpha^2 + a_3\alpha^3)$	$0.18 \leq M \leq 0.9$
	$+ M(a_4 + a_5\alpha + a_6\alpha^2 + a_7\alpha^3)$	
	$CYB_m = (a_0 + a_1\alpha + a_2\alpha^2 + a_3\alpha^3)$	$0.9 < M < 1.2$
	$+ M(a_4 + a_5\alpha + a_6\alpha^2 + a_7\alpha^3)$	
CYB	$CYB_h = (a_0 + a_1\alpha + a_2\alpha^2 + a_3\alpha^3)$	$1.2 \leq M \leq 1.9$
	$+ M(a_4 + a_5\alpha + a_6\alpha^2 + a_7\alpha^3)$	
CYDR	$CYDR = a_0 + a_1M + a_2M^2 + a_3M^3$	

$$Av_i = Ax_i \sin \theta_i - (Ay_i \sin \phi_i + Az_i \cos \phi_i) \cos \theta_i - g$$

(3-21)

Examination of Eq. (3-21) indicates that, depending upon the orientation of the aircraft, an AK test could be used to isolate a failure in an altimeter, any type of accelerometer, and the pitch and roll output channels of the attitude platform; the ARM algorithm contains AK tests for isolation of failures in the lateral and normal accelerometers and the altimeters.

During initial ARM algorithm design, it was anticipated that the AK test would be the most powerful test for isolating normal accelerometer failures at shallow bank angles. At that time, the major source of error in the altimeters, except during transonic flight, was thought to be the 3.4 m output quantization.⁽⁴⁾ However, during the flight testing of the Phase I ARM code, significant unexpected errors were found in the AK residuals during simulated accelerometer failures. Discussion with DFRC personnel revealed that earlier attempts to model the observed behavior of these altimeters had resulted in a model containing a first-order lag with a time constant of 0.5 second and hysteresis of uncertain magnitude between 10 m and 30 m.

Because of the uncertainty associated with the altimeter hysteresis, it was decided that no attempt would be made to model this effect in the Phase II algorithm. Instead, "traps" were added to the AK filters integrating Eq. (3-21) that bypass the incorporation of altitude measurement residuals smaller than a stipulated magnitude, and the worst-case error terms corresponding to initial altimeter error were increased significantly to accommodate the observed hysteresis effects. In addition, the vertical acceleration given by Eq. (3-21), computed using inertial instrument data, is passed through a first-order lag, with a 0.5-second time constant, before being double integrated and compared with the altimeter outputs.

The relatively large initial error term required to accommodate barometric altimeter hysteresis error in the AK tests renders the effectiveness of the Phase II AK tests inferior to that of the TK and TD tests in isolating failed accelerometers. Additionally, the altimeter fault-isolation performance achieved with the AK test, which is sensitive only to change in altimeter bias and not absolute bias level, could be obtained just as reliably via signal selftest. For these reasons, the details of the AK test implementation are not discussed here. The interested reader is referred to References 4 and 24 for additional information on the formation of AK test residuals.

3.6 Computational Requirements

The Phase II ARM algorithm occupies approximately 8000 16-bit words of computer memory. Bench tests have indicated that its timing requirements are approximately 5.6 ms per 60 ms with no detected faults and approximately 12 ms per 60 ms with 12 detected faults. The longest timing requirement for a single detected failure is approximately 7.0 ms per 60 ms. The bulk of the timing requirement with no detected faults represents the overhead associated with sensor read and scaling; moving window fault detection processing; aircraft state determination; sensor output selftest; analytic redundancy residual calculations for the Mach meters, altimeters, rate gyros, and attitude platforms; and the turbulence estimator. If the AK tests for normal and lateral accelerometers were removed, as suggested by the flight test results, together with all research-specific operations, the timing and memory figures could be lowered by approximately ten percent.

SECTION 4

FLIGHT TEST RESULTS

The ARM algorithm has been implemented in the F-8 DFBW aircraft computers in a parallel mode in which it is able to obtain all required sensor and effector position readings, but its sensor signal status information does not affect the choice of signals used in control law calculations; these choices continue to be made by the baseline FDI programs. The ARM software includes extensive error simulation and signal fault insertion capability, controlled by the pilot through the computer input panel (CIP), to allow inflight evaluation of ARM algorithm performance. Simulated signal faults that can be inserted (on the number one instruments only) include bias, drift, scale-factor error, hardover, transient pulse, and loss of signal. Errors that can be simulated include scale-factor errors in the normal and lateral aerodynamic coefficients CL and CYB, error in the assumed center of mass location, and misalignment of the number two rate gyro and accelerometer triads. In addition, the choice of which of three stored aircraft mass values is used in the ARM TD calculations is updated, on the basis of fuel remaining, by the pilot via CIP entry.

In addition to the excellent performance of the Phase II algorithm in identifying inserted sensor failures during flight testing, as will be discussed, the Phase I ARM algorithm performed well when confronted with two actual inflight failure situations. In one instance, an opening developed in the potentiometer for alpha vane two. The ARM algorithm declared alpha vane two unconditionally failed 24

seconds before the baseline code, which by itself cannot isolate the failed sensor of a duplex set, detected the alpha vane fault. Figure 4-1 shows the behavior of the Phase I ARM algorithm during this in-flight failure, with the aircraft at an altitude of 12 km and Mach 1.05. The first frame shows the voted, compensated normal accelerometer output of approximately -0.9 g. The second frame shows the two alpha vane readings, the third frame shows the TD residuals for each vane and the TD worst-case error magnitude, and the fourth frame shows the TD MLLR for each vane. Note that following the declaration of alpha vane two as conditionally failed at the time of detection, the fault isolation process is delayed by the output spike on alpha vane two on sample 14. This spike drives alpha vane two out of the aerodynamic coefficient error "good fit" region, analogous to Eq. (3-20), that was used in the Phase I algorithm. This, in turn, increases the magnitude of the TD worst-case error magnitude on that sample, and each TD MLLR is incremented by a large positive amount. With the removal of the alpha vane 2 output spike, the isolation process proceeds steadily, ending with the declaration of alpha vane two as unconditionally failed on sample 28. It should be noted that this and similar experiences with the alpha vane TD tests during the flight test program indicated that the "good fit" region in the Phase I algorithm was conservatively narrow. This resulted in the larger region, defined by Eq. (3-20), being used in the Phase II algorithm.

The second inflight malfunction occurred when the baseline system detected the memory parity failure of one of the triplex computers and declared it failed. Because the failed computer was dedicated to reading the number one sensor outputs and because of the mechanization of the buffer refreshing operation, the computer loss manifested itself to the ARM algorithm as all number one signals jumping to their negative maximum values. ARM sensor output selftest immediately declared all number one signals provisionally failed, and the analytic redundancy tests declared all number one signals unconditionally failed within 1.3 seconds of the computer failure.

Table 4-1 indicates the performance of the Phase II ARM algorithm in isolating inserted bias, drift, and hardover failures, with the aircraft near trim at Mach 0.6 at 6100 m altitude. These isolation times are in excellent agreement with theoretical values calculated for the MLLR residuals and worst-case error terms in a 1 g non-maneuvering environment. In all cases, sensor output selftest immediately declared a hardover sensor provisionally failed. It is important to note that between the times of failure insertion and fault isolation, the average of the two sensor outputs would be used in control calculations, thus halving the effective error. The fault isolation times shown in Table 4-1 are also representative of the injected failure isolation times observed with the aircraft undergoing moderate maneuvers.

Because of the variety of worst-case error terms in the different analytic redundancy tests, a summary of inserted fault isolation performance during extreme maneuvers similar to Table 4-1 is not possible. However, Figures 4-2 through 4-7 indicate representative results, both with and without simulated errors in the knowledge of important aircraft parameters.

Each of these figures contains four different frames, and three of these frames present analogous information in the different figures. One frame shows the two suspect signals, including the simulated failure. One frame contains the residual for each suspect signal and the worst-case error term for the particular type of analytic redundancy used: RK for the rate gyros and attitude platforms; TD for the alpha vanes and lateral accelerometers; and TK for the longitudinal accelerometers, normal accelerometers, and Mach meters. Finally, one frame contains the MLLR for each suspect signal.

Figure 4-2 indicates isolation of a simulated 0.5 scale factor in pitch rate gyro one. During the early negative pitch rate maneuver, the failure is detected but it cannot be isolated before the




Figure 4-1. Actual failure of alpha vane two.




Figure 4-2. Pitch rate gyro one simulated 0.5 scale factor.

Table 4-1. Inserted fault average isolation times at trim (seconds).

Signal Type	1.5 BFM Bias	BFM/s Drift	Hardover
Mach	0.06	NT	0
Altimeter	0.12	NT	0
Angle of attack	0.54	2.1	0.12
Long. accel.	7.26	3.6	0.78
Lat. accel.	0.3	1.38	0
Normal accel.	6.51	3.84	0.36
Roll rate	0.48	1.44	0.06
Pitch rate	0.9	1.8	0.24
Yaw rate	0.9	1.8	0.24
Roll attitude	0.18	NT	0
Pitch attitude	0.12	NT	0
Yaw attitude	0.18	NT	0

NT = no test

magnitude of the pitch rate decreases to the extent that the difference between the two pitch rate signals becomes insignificant. This forces the DRLLR to exceed the +9.2 threshold, removing the detected failure flag and reinitializing the detection process. (With the detected failure flag removed, all of the variables associated with fault isolation remain unchanged.) During the subsequent positive pitch rate maneuver, the fault is again detected, and instrument one is declared conditionally failed in six samples and unconditionally failed 12 samples later.

Figure 4-3 shows the successful isolation of a simulated 0.053 rad/s bias in pitch rate gyro one during a positive roll rate maneuver, with a simulated misalignment of the number two pitch rate gyro 0.02 rad about the aircraft yaw axis. The effect of the roll maneuver acting through the misalignment can be seen as a negative ramp in the RK residual for pitch rate gyro two in the third frame of the figure. However, it is important to note that the magnitude of this residual is always less than the postulated worst-case error magnitude, also shown in the third frame. Therefore, although the sign of the residual in pitch rate gyro two is consistent with the sign of the detected failure (positive for instrument one or negative for instrument two), the RK MLLR for pitch rate gyro two remains positive throughout the failure simulation. The plot of the MLLR for pitch rate gyro one indicates that it is declared conditionally failed 14 samples after failure injection and declared unconditionally failed 8 samples later.

Figure 4-4 shows the successful isolation of a simulated 0.3 g bias in lateral accelerometer one, with the TD calculations utilizing a value for the lateral aerodynamic coefficient CYB 0.9 times the computed value. During this test, the aircraft has a sideslip angle of 0.05 rad and a dynamic pressure of 22.5 kPa (470 lb/ft²). The use of a low CYB results in TD residuals for the two lateral accelerometers that are lower than they would normally be, and this is clearly evident in the figure as a negative bias in the residual for instrument two. Although the sign of this bias is consistent with the sign of the detected failure, it is smaller in magnitude than the worst-case error term, and the MLLR for instrument two remains positive. As indicated by its MLLR, lateral accelerometer one is declared conditionally failed three samples after failure injection and is declared unconditionally failed five samples later.

Figure 4-5 shows the successful isolation of a simulated 0.053 rad bias in alpha vane one, with the TD calculations utilizing a value for the lift coefficient 0.9 times the computed value. The aircraft




Figure 4-3. Pitch rate gyro one simulated 0.053 rad/s bias with simulated instrument two misalignment.




Figure 4-4. Lateral accelerometer one simulated 0.3g bias using 0.9 CYB.

is executing a 3.4 g turn at Mach 0.83 at an altitude of 5.4 km. The effect of the erroneous lift coefficient is reflected in residuals for the two alpha vanes that are more positive than they would normally be, and this can be seen as a positive bias in the residual for unfailed alpha vane two. Although this residual has a sign consistent with the sign of the detected failure, it is smaller than the postulated worst-case error magnitude. Thus, the MLLR for alpha vane two remains positive, while alpha vane one is declared conditionally failed five samples after failure insertion and is declared unconditionally failed five samples later.

Although this example shows correct alpha vane fault isolation by the ARM algorithm in spite of a ten percent error in calculated lift coefficient, this is not always the case throughout the aircraft's flight envelope. However, although some situations have been observed during 1.5 BFM bias insertions in which the ARM algorithm is unable to decide which alpha vane has failed, no instances of misisolation of the unfailed vane have been encountered with ten percent lift coefficient error. The observed instances of fault isolation indecision occur at high angles of attack, where the polynomial lift coefficient functions are known to be less accurate than at the lower angles of attack. This knowledge is reflected in the "good fit" region defined by Eq. (3-20). Improved alpha vane fault isolation performance could be obtained by using more complex lift coefficient models than those shown in Table 3-9, with corresponding decreases in the contribution of coefficient error to the worst-case error term.

Additional improvement in TD fault isolation test accuracy could be obtained through a more refined mass estimate update procedure. The current technique of using one of three stored mass estimates on the basis of pilot selection can, even when executed properly, result in TD residual error for a lateral accelerometer or alpha vane as large as five percent of the true lateral acceleration or normal acceleration, respectively.

The high-frequency oscillation in the compensated normal acceleration shown in the first frame of Figure 4-5, with a magnitude of approximately 0.3 g, is a manifestation of actual airframe flexure arising from what is termed "high alpha buffet." Unfortunately, the currently programmed ARM method of computing angular acceleration by simple back-differencing of the voted angular rates, when used in compensation Eq. (3-3), results in a term that is in phase with the accelerometer oscillations. In fact, the magnitude of the oscillations in the compensation term and the raw accelerometer output are approximately equal. Because of the high frequency of the oscillation in compensated normal acceleration, its overly large magnitude has not created any problems in the TD isolation tests, which are much more sensitive to low-frequency errors. If desired, the oscillation in the compensation term could be effectively eliminated by computing the angular acceleration from the back-difference over two ARM sample periods, without the need for more elaborate filtering of the rates.

Figure 4-6 shows the successful isolation of a -0.3 g bias in longitudinal accelerometer one, with the simulated misalignment of longitudinal accelerometer two 0.02 rad about the aircraft pitch axis. The failure is inserted just before the aircraft begins a 2.5 g windup turn. The effect of the simulated misalignment and the windup turn can be seen as an increasing positive TK residual for instrument two. Because this residual is significantly smaller than the postulated worst-case error, the MLLR for instrument two remains positive. Longitudinal accelerometer one is declared unconditionally failed 129 samples after failure injection.

As shown in frame three of Figure 4-6, the postulated worst-case error is extremely conservative relative to the observed residual of the unfailed instrument, and this conservatism results in a relatively long isolation time of nearly 8 seconds. As discussed earlier, the fact that the acceleration of the air mass is an unmodeled error source in the TK residual equation motivates the inclusion of a




Figure 4-5. Alpha vane one simulated
0.053 rad bias using 0.9 CL.




Figure 4-6. Longitudinal accelerometer one
simulated -0.3g bias with simulated
instrument two misalignment.

wind-shear doublet in the worst-case error term, with the magnitude chosen as one of two values based upon the output of a binary turbulence estimator. Significantly faster isolation performance could be obtained by lowering the magnitude (perhaps to zero) of the doublet used during those times, such as for Figure 4-6, when the turbulence estimate is low.

Figure 4-7 shows the successful isolation of a 0.3 g bias in normal accelerometer one during a 3 g turn. As discussed in Section 3, the TK test for the normal accelerometers is very sensitive to errors in the knowledge of the pitch rate. Extensive analysis of flight data indicates that pitch rate gyro two has a scale factor of approximately 0.95. The effect of this scale factor error on the estimated pitch rate during the windup turn having a pitch rate of 0.09 rad/s results in the negative ramp for the TK residual for normal accelerometer two seen in Figure 4-7. Since the residuals for both sensors are identically affected by the pitch rate error, the slope of the positive ramp failure signature for normal accelerometer one is smaller than it would be in nonrotating flight, accounting for the larger number of samples (268) required for isolating the failure. As for the longitudinal accelerometers, the isolation time for the normal accelerometers could be lowered by decreasing the magnitude of the low-turbulence wind-shear doublet used in the worst-case error term.




Figure 4-7. Normal accelerometer one simulated 0.3g bias.

SECTION 5

CONCLUSIONS AND RECOMMENDATIONS

The Phase I and II flight test programs have demonstrated the validity and capability of the ARM concept to achieve fail-operational performance with duplex sensors. Of particular importance has been the adaptability to unforeseen sensor behavior provided by the worst-case error terms in the analytic redundancy MLLRs. The ARM algorithm design and flight test experience have shown that the identification and accommodation of basic, irreducible, low-frequency error sources in the analytic redundancy residuals are mandatory for the design of a reliable fault isolation system using analytic redundancy. The proper interpretation of a low-frequency component in an analytic redundancy residual, i.e., sensor failure or error term, has proven a significantly more important design issue than the optimality of the decision rule that determines whether the component is present.

An alternate approach to the ARM technique for achieving fail-operational capability with duplex sensors involves discarding both sensors of a particular type when a disparity in their outputs is detected, and subsequently using a synthesized value of that variable in the control laws. The relationship used to synthesize the missing variable would be analogous to the relationships used to generate the analytic redundancy residuals in the ARM algorithm. However, it should be noted that the ARM approach has two significant advantages over this alternative. First, fidelity of the residual relationships in the ARM approach is required only over the short time period between fault detection and isolation, compared to the remainder of

the mission time for the alternate approach. Second, the ARM approach incorporates a consistent, self-contained mechanism to accommodate errors in the analytic redundancy relationships. Such accommodation via the control laws in the alternate approach is much more difficult to achieve.

Reasonability checks, in the form of sensor output selftest, have proven to be powerful complements to the analytic redundancy tests in the ARM technique. This is particularly true for kinematic-type analytic redundancy tests, some of which require significant times to isolate even a hard-failed sensor. The provision for conditional failure declaration has resulted in significant decreases in the effective isolation times for BFM-sized failures of the rate gyros, lateral accelerometers, and alpha vanes.

The use of a bias failure hypothesis in the MLLRs does not significantly restrict the ability to isolate nonbias failures in duplex systems, as indicated by the roll rate gyro scale-factor isolation shown in Figure 4-2. However, it should be noted that if more than one simplex sensor is an input to an analytic redundancy residual calculation, more elaborate failure mode modeling is required for reliable fault isolation.

Although the ARM software is more extensive than that associated with conventional redundancy management algorithms, it is similar in complexity to guidance and navigation calculations and posed no unusual implementation problems. Furthermore, its modularized structure lends itself to distributed computation systems.

The process of developing the reliable Phase II ARM algorithm from the preliminary Phase I version involved either the compensation for or accommodation of several error sources in the analytic redundancy relationships. This involved painstaking and tedious ground analysis of downlink data recorded during flight testing. Any technique requiring a similar level of effort for aircraft-specific

code initialization or to respond to day-to-day variations in the sensors or replacement of faulty units would clearly be impractical for fleet-wide application. Therefore, it is recommended that future effort be directed toward developing methods whereby FDI algorithms utilizing analytic redundancy can be made robust relative to variations in basic sensor characteristics. The ARM rate gyro bias estimators, utilizing rotational kinematics residuals, suggest one direction to pursue. Not only do these bias estimates allow certain MLLR worst-case error terms to be lowered, but the levels of the estimates themselves could be used for fault isolation or trend analysis. A recently developed related approach, in which Kalman filters are used to estimate the mathematically observable biases and scale-factor errors in direct redundant sensors, has given encouraging results in simulator studies.⁽²³⁾ While this approach requires real-time covariance matrix calculation to achieve optimal estimates, it appears to be a promising starting point for the development of (possibly suboptimal) estimators of the important sensor biases and scale factor errors appearing in analytic redundancy residuals. It must be noted, however, that the fact that each sensor error term may appear in more than one residual could pose a significant computational challenge to any sensor calibration error estimation technique.

Although the integration of the ARM technique into a redundant digital flight control system remains an applications problem, the F-8 DFBW aircraft flight test experience suggests that this is certainly within the state-of-the-art. This integrated system could either provide fail-operational capability following one more than the original design number of identical sensor failures or allow the removal of one sensor of each type, representing a significant savings in acquisition, spares, and maintenance costs over the life of the aircraft.

Such an integrated system should be capable of monitoring the health of any single remaining sensor of a particular type. Although

detail design of the required algorithms has not been performed, a promising candidate approach exists that is a straightforward adaptation of the ARM algorithm.⁽²⁴⁾ Basically this approach uses two MSPRTs for the single remaining sensor, one postulating a positive failure and the other a negative failure. A time interval is chosen, significantly shorter than the ETL for the sensor, after which the BFM failure signature dominates the noise in the analytic redundancy process. If after that interval has passed the LLR portion of the MLLR is positive, the MLLR is reset to zero and the process is repeated. Otherwise, the MLLR is updated until either: 1) the LLR becomes positive, at which time the MLLR is reset to zero, or 2) the MLLR crosses the failure threshold and the sensor is identified as failed.

The implementation of such an integrated system on the F-8 DFBW aircraft is strongly recommended. Although the flight test program reported here, incorporating approximately nine hours of flight time, indicated proper functioning of the ARM algorithm, such a test period is insufficient to exhaustively exercise a major software package or to compile accurate performance statistics such as false alarm and missed alarm rates. While an integrated system for the entire set of control sensors would be desirable in order to demonstrate the full range of benefits and to accumulate the performance characteristics of the integrated approach, such a full-scope demonstration is not mandatory. In particular, meaningful results could be obtained from an integrated system for a subset of the sensors monitored by the ARM algorithm, e.g., the attitude and attitude rate sensors.

Finally, it is well-known that significant savings in the replication of sensors measuring components of a vector quantity (such as angular velocity or linear acceleration) required to isolate a given number of successive sensor failures can be achieved by geometrically skewing the sensors. Thus, instead of the three aligned

orthogonal triads of rate gyros and accelerometers employed in the baseline F-8 design to isolate a single failed instrument of either type, skewed arrays of five instruments of each type would suffice,⁽¹⁾ a savings of eight instruments in all. It is important to note that for such a skewed array four instruments remain after the first identified failure, allowing only detection of a second failure, and that analytic redundancy relationships similar to those in the ARM algorithm could be used to isolate the second failure.

This Page Intentionally Left Blank

REFERENCES

1. Desai, M.N., and A. Ray, "A Fault Detection and Isolation Methodology," Proc. 20th IEEE Conf. on Decision and Control, San Diego, California, December 1981, pp. 1363-1369.
2. Desai, M.N., et al., "Dual Redundant Sensor FDI Techniques Applied to the NASA F8C-DFBW Aircraft," Proc. AIAA Guidance and Control Conf., San Diego, California, August 1976, pp. 502-513.
3. Deckert, J.C., et al., "F8-DFBW Aircraft Sensor Failure Identification Using Analytic Redundancy," IEEE Trans. Auto. Control, Vol. AC-22, No. 5, October 1977, pp. 795-803.
4. Deckert, J.C., et al., Reliable Dual-Redundant Sensor Failure Detection and Identification for the NASA F-8 DFBW Aircraft, CR-2944, NASA, Washington, D.C., February 1978.
5. Beard, R.V., Failure Accommodation in Linear Systems through Self-Reorganization, CR-118314, NASA, Washington, D.C., 1971.
6. Jones, H.L., Failure Detection in Linear Systems, Ph.D. Dissertation, Dept. Aeronautics and Astronautics, M.I.T., Cambridge, Mass., September 1973.
7. Kerr, T.H., "A Two Ellipsoid Overlap Test for Real Time Failure Detection and Isolation by Confidence Regions," presented at the Pittsburgh Conference on Modeling and Simulation, April 1974.
8. Montgomery, R.C., and A.K. Caglayan, "Failure Accommodation in Digital Flight Control Systems by Bayesian Decision Theory," Journal of Aircraft, Vol. 13, Feb. 1976, pp. 69-75.
9. Willsky, A.S., J.J. Deyst, and B.S. Crawford, "Two Self-Test Methods Applied to an Inertial System Problem," Journal of Spacecraft and Rockets, Vol. 12, July 1975, pp. 434-437.
10. Davis, M.H.A., "The Application of Nonlinear Filtering to Fault Detection in Linear Systems," IEEE Transactions on Automatic Control, Vol. AC-20, April 1975, pp. 257-259.

11. Chien, T.T., An Adaptive Technique for a Redundant Sensor Navigation System, CR-140313, NASA, Washington, D.C., 1972.
12. Mehra, R.K., and J. Peschon, "An Innovations Approach to Fault Detection and Diagnosis in Dynamic Systems," Automatica, Vol. 7, September 1971, pp. 637-640.
13. Willsky, A.S., and H.L. Jones, "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems," IEEE Transactions on Automatic Control, Vol. AC-21, February 1976, pp. 108-112.
14. Deyst, J.J., and J.C. Deckert, "Maximum Likelihood Failure Detection Techniques Applied to the Shuttle RCS Jets," Journal of Spacecraft and Rockets, Vol. 13, February 1976, pp. 65-74.
15. Maybeck, P.S., Failure Detection Through Functional Redundancy, AFFDL-TR-76-93, September 1976.
16. Cunningham, T.B., et al., Fault Tolerant Digital Flight Control with Analytical Redundancy, AFFDL-TR-77-25, May 1977.
17. Deckert, J.C., Definition of the F-8 DFBW Aircraft Control Sensor Analytic Redundancy Management Algorithm, Report R-1178, The Charles Stark Draper Laboratory, Inc., Cambridge, Mass., August 1978.
18. Deckert, J.C., and R.R. Bairnsfather, Software Specification for F-8 DFBW Aircraft Duplex Control Sensor FDI, Report R-1133 Revised, The Charles Stark Draper Laboratory, Inc., Cambridge, Mass., June 1978.
19. Willsky, A.S., "A Survey of Design Methods for Failure Detection in Dynamic Systems," Automatica, Vol. 12, 1976, pp. 601-611.
20. Wald, A., Sequential Analysis, Dover, New York, 1973, Chap. 3.
21. Wald, A., and J. Wolfowitz, "Optimum Character of the Sequential Probability Ratio Test," Annals of Math. Stat., Vol. 19, 1948, pp. 326-339.
22. Etkin, B., Dynamics of Flight - Stability and Control, Second Edition, John Wiley & Sons, Inc., New York, 1982, pp. 299-300.
23. Hall, S.R., Parity Vector Compensation for FDI, Report T-763, The Charles Stark Draper Laboratory, Inc., Cambridge, Mass., February 1982.

24. Desai, M.N., J.C. Deckert, and J.J. Deyst, "Dual-Sensor Failure Identification Using Analytic Redundancy," AIAA J. Guidance and Control, Vol. 2, No. 3, May-June 1979, pp. 213-220.

1. Report No. NASA CR-170396		2. Government Accession No.		3. Recipient's Catalog No.	
4. Title and Subtitle ANALYTIC REDUNDANCY MANAGEMENT MECHANIZATION AND FLIGHT DATA ANALYSIS FOR THE F-8 DIGITAL FLY-BY-WIRE AIRCRAFT FLIGHT CONTROL SENSORS				5. Report Date April 1983	
				6. Performing Organization Code	
7. Author(s) James C. Deckert				8. Performing Organization Report No. CSDL-R-1520	
9. Performing Organization Name and Address The Charles Stark Draper Laboratory, Inc. 555 Technology Square Cambridge, Massachusetts 02139				10. Work Unit No.	
				11. Contract or Grant No. NAS4-2675	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, D.C. 20546				13. Type of Report and Period Covered Contractor Report - Final July 1978 - Sept. 1981	
				14. Sponsoring Agency Code RTOP 533-02-61	
15. Supplementary Notes NASA Technical Monitor: Kenneth J. Szalai, Ames Research Center, Dryden Flight Research Facility, Edwards, CA 93523.					
16. Abstract The details are presented of an onboard digital computer algorithm designed to reliably detect and isolate the first failure in a duplex set of flight control sensors aboard the NASA F-8 digital fly-by-wire aircraft. The algorithm's successful flight test program is summarized, and specific examples are presented of algorithm behavior in response to software-induced signal faults, both with and without aircraft parameter modeling errors.					
17. Key Words (Suggested by Author(s)) Analytic redundancy Digital fly-by-wire Failure detection and identification Redundancy management Reliability			18. Distribution Statement Unclassified STAR category OR		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 63	
				22. Price	

Available through NASA's Industrial Applications Centers.

End of Document