

PERFORMANCE AND SENSITIVITY ANALYSIS OF THE GENERALIZED LIKELIHOOD
RATIO METHOD FOR FAILURE DETECTION

by

RAMON A. BUENO

B.S., Massachusetts Institute of Technology
(1974)

SUBMITTED IN PARTIAL FULFILMENT

OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

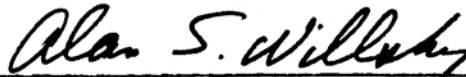
February, 1977

Signature of Author



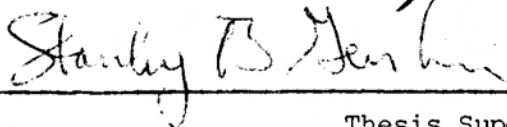
Department of Aeronautics and Astronautics
September 27, 1976

Certified by



Thesis Supervisor

Certified by



Thesis Supervisor

Accepted by

Chairman, Department Committee on Graduate Students

PERFORMANCE AND SENSITIVITY ANALYSIS OF THE GENERALIZED LIKELIHOOD
RATIO METHOD FOR FAILURE DETECTION

by

RAMON A. BUENO

Submitted to the Department of Aeronautics and Astronautics on
September 27, 1976, in partial fulfillment of the requirements
for the degree of Master of Science.

ABSTRACT

This thesis reports some results from studies of the Generalized Likelihood Ratio (GLR) technique to the detection of failures in an aircraft application. The GLR method can be used for detecting and identifying abrupt changes in linear, dynamic systems. It is designed to determine simultaneously whether a change has taken place, the time that the change occurred and an estimate of the extent of the change.

The technique is presented and its relationship to the properties of the Kalman-Bucy filter is examined. Under the assumption that the system is perfectly modeled, the detectability and distinguishability of four failure types is investigated by means of analysis and simulations. Detection of failures is found satisfactory, but problems in identifying correctly the mode of a failure may arise. These issues are closely examined as well as the sensitivity of GLR to modeling errors. The advantages and disadvantages of this technique are discussed and various modifications are suggested to reduce its limitations in performance and computational complexity.

THESIS SUPERVISOR: Alan S. Willsky

TITLE: Associate Professor of Electrical Engineering

THESIS CO-SUPERVISOR: Stanley B. Gershwin

TITLE: Lecturer in Electrical Engineering

ACKNOWLEDGEMENTS

I would like to express my gratitude to Prof. Alan Willsky for his continued support, direction, and patience up to the last minutes. Also to Dr. S.B. Gershwin for his frequent suggestions and encouragement and to Mr. Ed Y. Chow for the many discussions and his assistance during the earlier stages of this research.

I am also thankful to Ms. Fife Monserrate, Ms. Margaret Flaherty, Mr. Arthur Giordani and Mr. Darling for their valuable contribution to the preparation of this report.

Finally, my most warm appreciation to my family and friends without whose unfailing confidence and encouragement I would not have had the stamina to finish. Muchas Gracias!

This research was conducted at the Electronic Systems Laboratory, M.I.T. with support from the NASA Langley Research Center No. NSG-

TABLE OF CONTENTS		Page
ABSTRACT		2
ACKNOWLEDGEMENTS		3
TABLE OF CONTENTS		4
LIST OF TABLES		6
LIST OF FIGURES		7
CHAPTER 1 INTRODUCTION		12
1.1 Motivation		12
1.2 The GLR Method		16
1.2.1 The Generalized Likelihood Ratio (GLR) Technique		16
1.2.2 Application Problem: The Reduced F-8C Model		29
1.3 Summary of the GLR Approach and Overview		38
CHAPTER 2 CORRECT DETECTION UNDER MATCHED CONDITIONS		41
2.1 Matched Conditions and Correct Detection		41
2.2 The Detection Performance of the GLR		42
2.2.1 Performance Analysis		42
2.2.2 The Signature and Information Matrices		51
2.3 Simulation Results for Correct Detection		80
2.3.1 Detector Implementation		80
2.3.2 False Alarms: Detector Sensitivity		84
2.3.3 Failure Detection		89
2.3.4 The Failure Estimates: \hat{v} and $\hat{\theta}$		99
2.4 Summary of Correct Detection		117
CHAPTER 3 IDENTIFIABILITY OF FAILURES WITH THE GLR		122
3.1 Matched Conditions and the Cross-Detection Problem		122
3.2 Some Analysis of the Cross-Detection Problem		125

	<u>Page</u>
3.3 Cross-Detection Simulations	132
3.3.1 The Likelihood Ratios	132
3.3.2 The Failure Estimates	141
3.4 Distinguishability with Constrained GLR (CGLR)	144
3.5 Cross-Detection: Discussion	149
CHAPTER 4 SENSITIVITY TO MODELING ERRORS	151
4.1 Detection Under Mismatched Conditions	151
4.2 The Filter Matched to the Detector	152
4.2.1 Description of Mismatch	152
4.2.2 Complete Mismatch	155
4.2.3 Failure Detection Under Complete Mismatch	165
4.3 Partial Mismatch: ΔH	175
4.3.1 Compensation for ΔH	175
4.3.2 Approximate Analysis of Complete Mismatch	179
4.4 The Filter Matched to the System	182
4.5 Discussion of Modeling Errors and the GLR	186
CHAPTER 5 CONCLUSIONS AND RECOMMENDATIONS	189
5.1 Conclusions	189
5.2 Suggestions for Further Work	191
APPENDIX A	194
The Failure Signatures for Failure Models 1-4	194
APPENDIX B	197
Approximate Analysis: Partial Mismatch (in H) with the Filter Matched to the Detectors	197
APPENDIX C	206
Approximate Analysis: Partial Mismatch (in H) with the Filter Matched to the System	206
REFERENCES	215

LIST OF TABLES

Table 1.1	Set of Failures Simulated	37
Table 1.2	Noise Levels in Standard Deviations	38
Table 2.1	False Alarm Rates for Different Thresholds	88
Table 2.2	Delays in Detection for Different Thresholds: Jump Failures	91
Table 2.3	Delays in Detection for Different Thresholds: Step Failures	98
Table 3.1	Likelihood Ratio Profiles Over Window	139
Table 4.1	System and Filter Parameters for Aircraft Model at Flight Condition 12	154

LIST OF FIGURES

		<u>Page</u>
Figure 1.1	Full GLR Detector Scheme with Growing Bank of Filters	25
Figure 2.1	$P(\delta^2, \epsilon)$ for a Chi Squared Random Variable with 2 Degrees of Freedom	49
Figure 2.2	State Jump $G(r)$	55
Figure 2.3	Sensor Jump $G(r)$	56
Figure 2.4	State Step $G(r)$	57
Figure 2.5	Sensor Step $G(r)$	58
Figure 2.6	P_D and δ^2 for a 5σ State Jump Failure; $\epsilon=14$	65
Figure 2.7	P_D and δ^2 for a $5\sigma'$ Sensor Jump Failure; $\epsilon=14$	66
Figure 2.8	P_D and δ^2 for a 1σ State Step Failure; $\epsilon=14$	67
Figure 2.9	P_D and δ^2 for a $5\sigma'$ Sensor Step Failure; $\epsilon=14$	68
Figure 2.10	State Jump $C^{-1}(r)$	74
Figure 2.11	Sensor Jump $C^{-1}(r)$	75
Figure 2.12	State Step $C^{-1}(r)$	76
Figure 2.13	Sensor Step $C^{-1}(r)$	77
Figure 2.14	Simulation Flowgraph	82
Figure 2.15	Log-likelihood Ratios in the Absence of Failures: False Alarms	86
Figure 2.16	Log-likelihood Ratios for Jump Failures	94
Figure 2.17	Log-likelihood Ratios for Step Failures	96

	Page
Figure 2.18(a) $\ell(k, \theta)$ in State Jump Detector: 5 σ q State Jump	103
(b) $\ell(k, \theta)$ in State Jump Detector: 5 σ α State Jump	103
Figure 2.19(a) $\ell(k, \theta)$ in State Step Detector: 1 σ q State Step	104
(b) $\ell(k, \theta)$ in State Step Detector: 1 σ α State Step	104
Figure 2.20(a) $\ell(k, \theta)$ in Sensor Jump Detector: 5 σ' q Sensor Jump	107
(b) $\ell(k, \theta)$ in Sensor Jump Detector: 5 σ' α Sensor Jump	107
Figure 2.21(a) $\ell(k, \theta)$ in Sensor Step Detector: 1 σ' q Sensor Step	108
(b) $\ell(k, \theta)$ in Sensor Step Detector: 1 σ' α Sensor Step	108
Figure 2.22(a) $\ell(k, \theta)$ in Sensor Step Detector: 10 σ' q Sensor Step	110
(b) $\ell(k, \theta)$ in Sensor Step Detector: 10 σ' α Sensor Step	110
Figure 2.23 $\hat{v}(k)$ for 1 σ State Jump and State Step Failures	113
Figure 2.24 $\hat{v}(k)$ for 5 σ State Jump and State Step Failures	114
Figure 2.25 $\hat{v}(k)$ for 5 σ Sensor Jump and Sensor Step Failures	116
Figure 3.1 Wrong-time Cross-Detection Noncentrality Parameter: Step Failures	130

	<u>Page</u>
Figure 3.2(a) $\ell(k, \theta)$ in State Step Detector: $10\sigma'$ q Sensor Step	135
(b) $\ell(k, \theta)$ in State Step Detector: $10\sigma'$ α Sensor Step	135
Figure 3.3(a) $\ell(k, \theta)$ in State Step Detector: $10\sigma'$ q Sensor Step	137
(b) $\ell(k, \theta)$ in State Step Detector: $1\sigma'$ α Sensor Step	137
Figure 3.4(a) $\ell(k, \theta)$ in Sensor Step Detector: 1σ q State Step	140
(b) $\ell(k, \theta)$ in Sensor Step Detector: 1σ α State Step	140
Figure 3.5 $\hat{V}(k)$ for 1σ q State Step Failure: Sensor Step and State Step Detectors	142
Figure 3.6 $\hat{V}(k)$ for 1σ α State Step Failure: Sensor Step and State Step Detectors	143
Figure 4.1(a) Residuals under Mismatch: No Failure (α initial condition)	156
(b) Mean Value of Residuals in Figure 4.1(a)	156
Figure 4.2(a) $\ell(k, \theta)$ in State Step Detector Under Mismatch: No Failure	158
(b) $\ell(k, \theta)$ in Sensor Step Detector Under Mismatch: No Failure	158
Figure 4.3(a) $\hat{\theta}(k)$ in State Step Detector Under Mismatch: No Failure	160
(b) $\hat{\theta}(k)$ in Sensor Step Detector Under Mismatch: No Failure	160

	<u>Page</u>
Figure 4.4(a) $\ell(k, \theta)$ in State Step Detector Under Mismatch: No Failure	164
(b) $\ell(k, \theta)$ in Sensor Step Detector Under Mismatch: No Failure	164
Figure 4.5 $\ell(k, \theta)$ in State Step Detector Under Mismatch: 1σ q State Step	166
Figure 4.6 $\ell(k, \theta)$ in Sensor Step Detector Under Mismatch: 1σ q Sensor Step	168
Figure 4.7 $\ell(k, \theta)$ in State Step Detector Under Mismatch: 1σ α State Step	170
Figure 4.8 $\ell(k, \theta)$ in Sensor Step Detector Under Mismatch: $1\sigma'$ α Sensor Step	171
Figure 4.9(a) $\hat{\theta}(k)$ in State Step Detector Under Mismatch: 1σ q State Step	173
(b) $\hat{\theta}(k)$ in State Step Detector Under Mismatch: 1σ α State Step	173
Figure 4.10(a) $\ell(k, \theta)$ in State Step Detector Under H Mismatch: No Failure (α initial condition)	176
(b) $\ell(k, \theta)$ in Sensor Step Detector Under H Mismatch: No Failure (α initial condition)	176
Figure 4.11(a) $\ell(k, \theta)$ in State Step Detector with Partial Compensation: No Failure (α initial condition)	177
(b) $\ell(k, \theta)$ in Sensor Step Detector with Partial Compensation: No Failure (α initial condition)	177

	<u>Page</u>
Figure 4.12(a) $\ell(k, \theta)$ in State Step Detector with Partial Compensation: No Failure (q initial condition)	178
(b) $\ell(k, \theta)$ in Sensor Step Detector with Partial Compensation: No Failure (q initial condition)	178

CHAPTER 1

INTRODUCTION

1.1 Motivation

This thesis is concerned with the detection of abrupt changes in linear dynamic systems. The problem is one with many implications for estimation and control.

In many applications with Kalman-Bucy filtering [1], very good performance can be achieved when the modeling of the dynamics is sufficiently accurate. For a number of important applications, however, a linear model is only a good approximation to the actual system dynamics over short time intervals. In others, a filter may be desired of lower dimension than that of an accurate model. The use of the Kalman-Bucy filter in such cases frequently results in the divergence of the state trajectory of the system. Another important situation which also presents difficulties for estimation and control is the occurrence of failures in a system. Such events can be troublesome even for systems which are otherwise modeled with high accuracy.

One practical approach to the divergence problem is the use of adaptive estimation and filtering techniques which have been developed, as in [2], [3]. Although these techniques abound in the literature for cases where the changes in a system are slowly time-varying, only in the last few years has the problem of detecting sudden changes been addressed.

Linear models of systems subject to abrupt changes can be used to study these problems involving modeling inaccuracies. They can also

be used to represent systems which may undergo failures in some of its components, such as sensors or actuators. The word "failure", as used here, refers to abrupt changes in a system which need not be physical failures -- e.g., a sudden, but infrequent, acceleration in a system when the estimation is based on a constant velocity model.

With the greater availability and lower costs of digital hardware and software, more sophisticated design techniques can be studied in order to improve overall system performance and reliability. A recent survey by Willsky [4] describes a number of failure detection techniques and mentions some of the characteristics and tradeoffs involved in the various methods discussed. The class of problems considered in [4] is that of linear systems, where some analysis is possible. In this thesis we look in some detail at the Generalized Likelihood Ratio technique [5], [6].

The Generalized Likelihood Ratio (GLR) technique consists of performing hypothesis testing on the residuals, or innovations process, of the optimal (Kalman-Bucy) filter. The different hypothesis correspond to the behavior of the residuals assuming various failure models (including, of course, the 'no-failure' situation under normal conditions). The GLR formulation results in decision functions which allow us to extract a large amount of information (such as estimates of the failure mode, size and time of occurrence) about a failure, in addition to indicating whether a failure has occurred or not. A detailed analysis of the detection and identification performance of this technique is possible. In the following chapters we examine the performance of the GLR technique in a simple application. Finally, by offering a flexible set of implementations, the GLR

provides a practical method with which to study the many issues and trade-offs in the design of a reliable failure detection system. Some related work which this researcher has come across may be found in [7], [8], [9], [10].

Some important question concerning the capabilities of GLR for failure detection are of interest. In particular, we will look at several performance indices such as

- . false alarm rates
- . delays in detection
- . ability to distinguish among a variety of failure modes
- . sensitivity to modeling errors

The tradeoffs to be resolved in obtaining optimum performance, with these considerations in mind, are trade-offs between complexity (of the failure detection design) and performance. For a technique like the GLR, issues such as these must be understood before the relative merits of hardware redundancy vs. 'analytical redundancy' can be appreciated. The advantages and disadvantages of using hardware redundancy with voting decisions -- e.g., comparing two or more sensors to determine their reliability -- versus the use of more sophisticated (and generally, with higher computational costs) techniques must be considered. This work, hopefully, is a small step in such a direction. Further research should help identify situations where the added complexity is warranted.

Basic to the GLR method is the formulation of the alternative hypotheses by means of different failure models. It is therefore important to consider the relative distinguishability between failure of different modes when implementing the GLR detectors. It is important, for example,

to distinguish a sensor failure from an actuator failure. Correct identification of the mode of a failure is examined here, as well as the effects of errors in the modeling of the dynamic system. The performance of the detection system under various conditions is examined in the following sections in order to indicate where the main limitations of this approach lie. At this stage in the developments of a methodology for the design of GLR failure detection systems, it was felt that it is important to consider these questions in the context of a simple but meaningful example. We have chosen a simplified model of the F-8C aircraft.

It must be pointed out that in the following chapters the discussion is concerned with a "worst-case" kind of situation. We are assuming that there is only one set of measurements with which to work with. This is in contrast to the case where measurements are available from a set of redundant sensors [9], [10].

Note: Some inaccuracies exist in the reported values of various quantities in [12] and [13]. This work should correct such instances, until a further update is deemed necessary.

1.2 THE GLR METHOD

1.2.1 The Generalized Likelihood Ratio (GLR) Technique

We will now describe the generalized likelihood ratio (GLR) technique and the modifications which provide simpler formulations. The modifications allow some flexibility in the application of this approach to failure detection.

Consider the dynamical system model

$$x(k+1) = \Phi(k)x(k) + B(k)u(k) + w(k) + f_D(k, \theta) \quad (1.1)$$

$$z(k) = H(k)x(k) + J(k)u(k) + v(k) + f_S(k, \theta) \quad (1.2)$$

where $x(k) \in R^n$, $u(k) \in R^m$, $z(k) \in R^p$ are the state, input and output, respectively. The independent, Gaussian, white noise sequences $v(k)$ and $w(k)$ have statistics

$$E(w(k)) = 0, \quad E(w(k)w(j)') = Q(k)\delta_{kj} \quad (1.3)$$

$$E(v(k)) = 0, \quad E(v(k)v(j)') = R(k)\delta_{kj} \quad (1.4)$$

where δ_{ij} is the Kronecker delta (in this context representing the unit pulse at time $i=j$). The terms f_D and f_S are used to model a variety of abrupt system changes. By θ we denote the unknown time of occurrence of the failure.

The types of changes which are the subject of this thesis are:

1. Dynamics or State Jump

$$f_D(k, \theta) = v\delta_{k+1, \theta} \quad (1.5)$$

where $v \in R^n$ denotes the unknown direction and magnitude of the failure in state space and δ_{ij} is defined above. This model can be used to describe the occurrence of a brief disturbance in the system such as a sudden noise spike well outside the given statistics for $w(k)$ in (1.3), for example. Or, it can also describe momentary deviations of the control action away from the expected $B(k)u(k)$. Notice that this includes possible changes in either the control signal $u(k)$ or the gain through which it enters the system $B(k)$, or both. The modeling of v will be discussed later.

2. Dynamics or State Step

$$f_D(k, \theta) = v\sigma_{k+1, \theta} \quad (1.6)$$

Here σ_{ij} is the unit step function

$$\sigma_{ij} = \begin{cases} 0 & i < j \\ 1 & i \geq j \end{cases} \quad (1.7)$$

This model can be used to represent the effect of a failed actuator in a control system, whereby an additional constant, driving signal enters the system thus causing the state to move away from a prescribed trajectory. For example, a sudden increase in demand in a power generating system, where the demand acts as a control, due to an emergency situation perhaps due to a failure or shutdown in one of the plants. Or, a constant bias in the control signal of a digital flight control system of an aircraft which can be the result of some component malfunction. In addition, in many cases more complex failures (such as scale factor changes) resemble steps over periods of time, and a detector looking for a step can be used

to detect them .

3. Sensor Jump

$$f_s(k, \theta) = v\delta_{k\theta} \quad (1.8)$$

This can be used to represent bad data points (outliers) in the measurements of the output variables. By detecting such points one can prevent the error in the state estimates that otherwise results in the estimation subsystem (e.g., a Kalman filter).

4. Sensor Step

$$f_s(k, \theta) = v\sigma_{k\theta} \quad (1.9)$$

This models the onset of a bias in a measuring instrument. Again, this is an important situation where detection of the failure makes possible the removal of the consistent errors which result in the estimation of the process variables. The importance of being able to detect failures in the measuring system grows if one considers the use of feedback involving the estimated state variables. Again, one can use a model such as (1.9) to approximate a scale factor change over a period of time.

We recognize that these models are highly simplified descriptions of actual failure situations. However, they allow for detailed analysis in the GLR context and thus provide the opportunity to gain valuable insight into the workings of this technique. This is necessary before we can move on to analytically more complicated failure models.

Other failure modes that are of interest and that can be the subject of future study are:

5. Hard-over Actuator

$$f_D(k, \theta) = \Delta B(k)u(k)\sigma_{k+1, \theta} \quad (1.10)$$

Here $\Delta B(k)$ is an unknown change in the effectiveness of one or more of the actuators.

6. Increased Actuator Noise

$$f_D(k, \theta) = \xi(k)\sigma_{k+1, \theta} \quad (1.11)$$

where ξ is a zero-mean, white noise sequence with unknown covariance E .

7. Dynamics Shift

$$f_D(k, \theta) = \Delta \Phi x(k)\sigma_{k+1, \theta} \quad (1.12)$$

where $\Delta \Phi$ is an unknown shift in the plant dynamics.

8. Hard-over Sensor

$$f_S(k, \theta) = [\Delta Hx(k) + \Delta Ju(k)]\sigma_{k, \theta} \quad (1.13)$$

where ΔH , ΔJ represent scale factor changes in the sensors.

9. Increased Sensor Noise

$$f_S(k, \theta) = \xi(k)\sigma_{k, \theta} \quad (1.14)$$

This list of possible failure models is not intended to represent all cases of interest. Nevertheless, it is clear that these cases and combinations of them provide a fairly broad range of failure modes on which to base a study of failure detection techniques. It is also worth reiterating that failure models 1-4 may be viewed as having models 5-9 embedded in them, although admittedly not in a trivial way. For example,

models 5, 7 and 8 are, in some sense, 'dynamic biases' which are not radically different from models 2 and 4 when $x(k)$ and $u(k)$ remain approximately constant over the detection interval. This may be the case in systems with very slow dynamics. Alternatively, the increased noise models 6 and 9 can be thought of as random sequences of jumps, as in modes 1 and 3, given the proper distribution of their magnitudes v and times of occurrence θ .

Let us now consider the detection of such failures. The basis for the GLR approach, given in Willsky-Jones [6], is as follows. Design a Kalman filter based on the "no failure" hypothesis ($f_D = f_S = 0$). The filter equations are

$$\hat{x}(k+1|k) = \Phi(k)\hat{x}(k|k) + B(k)u(k) \quad (1.15)$$

$$\hat{x}(k|k) = \hat{x}(k|k-1) + K(k)\gamma(k) \quad (1.16)$$

$$\gamma(k) = z(k) - H(k)\hat{x}(k|k-1) - J(k)u(k) \quad (1.17)$$

where $\hat{x}(i|j)$ is the optimal (minimum mean-squared error) estimate of the state at time i based on measurements up to, and including, time j . The $\gamma(k)$ are the zero-mean, Gaussian innovations process (the residuals) associated with the optimal filter. The optimal filter gain sequence $K(k)$ is calculated from the equations

$$P(k+1|k) = \Phi(k)P(k|k)\Phi'(k) + Q(k) \quad (1.18)$$

$$P(k|k) = P(k|k-1) - K(k)H(k)P(k|k-1) \quad (1.19)$$

$$V(k) = H(k)P(k|k-1)H'(k) + R(k) = E(\gamma(k)\gamma'(k)) \quad (1.20)$$

$$K(k) = P(k|k-1)H'(k)V^{-1}(k) \quad (1.21)$$

with the appropriate initial conditions

$$\hat{x}(0|0) = E(x(0)) = \bar{x}_0 \quad (1.22)$$

$$P(0|0) = E([x(0) - \bar{x}_0][x(0) - \bar{x}_0]') = \Psi \quad (1.23)$$

Now suppose that an abrupt change, corresponding to one of the above models, takes place at time θ and from then on the system is described by (1.1) and (1.2).

By the linearity in the assumed system model one can then obtain an expression for the filter residuals

$$\gamma(k) = \tilde{\gamma}(k) + s(k, \theta) \quad (1.24)$$

where $\tilde{\gamma}$ is the "no failure" innovations with the statistics given above, and $s(k, \theta)$ is the effect of the failure on the residual. The form that $s(k, \theta)$ takes is different for each of the failure models. In particular, for models 1-4 which concern us here, one can write

$$s(k, \theta) = G_i(k, \theta)v \quad (1.25)$$

where $i \in \{1, 2, 3, 4\}$ denotes the failure type, and $G_i(k, \theta)$, called the failure signature matrix, can be precomputed (see Appendix A). This matrix gives the effect of the failure on the residual and, as one would suspect, plays a crucial role in the development and analysis of the GLR system.

The failure detection problem can be formulated as a decision to be made between competing hypothesis (Schweppe [14], Van Trees [15]).

$$H_0 : \text{no failure has occurred } (\theta > k)$$

H_i : a failure of type i has occurred ($\theta \leq k$)

where θ and v are regarded as unknown parameters (i.e., no prior distributions are postulated). The GLR approach attempts to isolate the different failures by using knowledge of the different effects that such failures have on the system residuals. An equivalent formulation of the problem, which is more useful for our purpose, is the following:

$$H_0 : \gamma(k) = \tilde{\gamma}(k) \quad (1.26)$$

$$H_i : \gamma(k) = \tilde{\gamma}(k) + G_i(k, \theta)v \quad (1.27)$$

The generalized likelihood ratio test (Van Trees [15]) can now be applied and we obtain a decision function based on the innovations sequence. Given this sequence, the procedure consists of computing the maximum likelihood estimates (MLE's) of θ and v assuming that a failure has occurred (for each type of failure). Substituting these values into a likelihood ratio test, one can then proceed to decide on the hypotheses. In other words, given the estimates

$$\hat{\theta}(k), \hat{v}(k) = \arg \max_{\theta, v} p(\gamma(1), \dots, \gamma(k) | H_i, \theta = \tilde{\theta}, v = \tilde{v}) \quad (1.28)$$

where p denotes the probability density function, the generalized likelihood ratio is defined by

$$L_i(k) = \frac{p(\gamma(1), \dots, \gamma(k) | H_i, \theta = \hat{\theta}(k), v = \hat{v}(k))}{p(\gamma(1), \dots, \gamma(k) | H_0)} \quad (1.29)$$

A decision between H_0 and H_i can then be made by means of a decision rule

$$\max_i L_i(k) \begin{matrix} H_1 \\ > \\ < \\ H_0 \end{matrix} \eta \quad (1.30)$$

where η is appropriately chosen as a design parameter in conformity with the goals of the detection system.

Without loss of information, one can simplify the decision function (1.29) by taking logarithms

$$\begin{aligned} \ell_i(k) &= 2 \ln L_i(k) \\ &= \sum_{j=1}^k \gamma'(j) v^{-1}(j) \gamma(j) \end{aligned} \quad (1.31)$$

$$- \sum_{j=1}^k [\gamma(j) - G_i(j, \hat{\theta}(j)) \hat{v}(j)]' v^{-1}(j) [\gamma(j) - G_i(j, \hat{\theta}(j)) \hat{v}(j)]$$

Here $\hat{v}(k)$ can be expressed explicitly as a function of $\hat{\theta}_i(k)$

$$\hat{v}_i(k) = C_i^{-1}(k, \hat{\theta}_i(k)) d_i(k, \hat{\theta}_i(k)) \quad (1.32)$$

where $C_i(k, \theta)$ is the deterministic matrix

$$C_i(k, \theta) = \sum_{j=1}^k G_i'(j, \theta) v^{-1}(j) G_i(j, \theta) \quad (1.33)$$

and $d_i(k, \theta)$ is a linear combination of the residuals

$$d_i(k, \theta) = \sum_{j=1}^k G_i'(j, \theta) v^{-1}(j) \gamma(j) \quad (1.34)$$

Finally, the log-likelihood ratio can be given by

$$\ell_i(k, \theta) = d_i'(k, \theta) C_i^{-1}(k, \theta) d_i(k, \theta) \quad (1.35)$$

and the MLE $\hat{\theta}_i(k)$ is the value of $\tilde{\theta} \leq k$ that maximizes

$$\ell_i(k, \hat{\theta}_i) = \max_{\tilde{\theta}} d_i'(k, \tilde{\theta}) C_i^{-1}(k, \tilde{\theta}) d_i(k, \tilde{\theta}) \quad (1.36)$$

The decision rule becomes

$$\max_i \ell_i(k, \hat{\theta}_i(k)) \underset{H_0}{\overset{H_i}{>}} \varepsilon = 2 \ln \eta \quad (1.37)$$

A failure of type i is declared if $\ell_i(k, \hat{\theta}_i(k))$ exceeds the threshold ε .

Some remarks are in order here.

- In equations (1.33) and (1.34) many of the terms are zero, since $G(j, \theta) = 0$ for $j < \theta$.
- The signature matrices $G_i(k, \theta)$ can be precomputed and stored, or they can be generated recursively (see Appendix A).
- The equations for $d_i(k, \theta)$ can be interpreted as defining matched filters for a failure of type i (for more on matched filters, see Van Trees [15], Schweppe [14]).
- As Figure 1.1 illustrates, this detection scheme involves the implementation of a growing bank of matched filters for different θ . To avoid this complication, the maximization in (1-36) is restricted to a "data window" such that θ lies in the interval

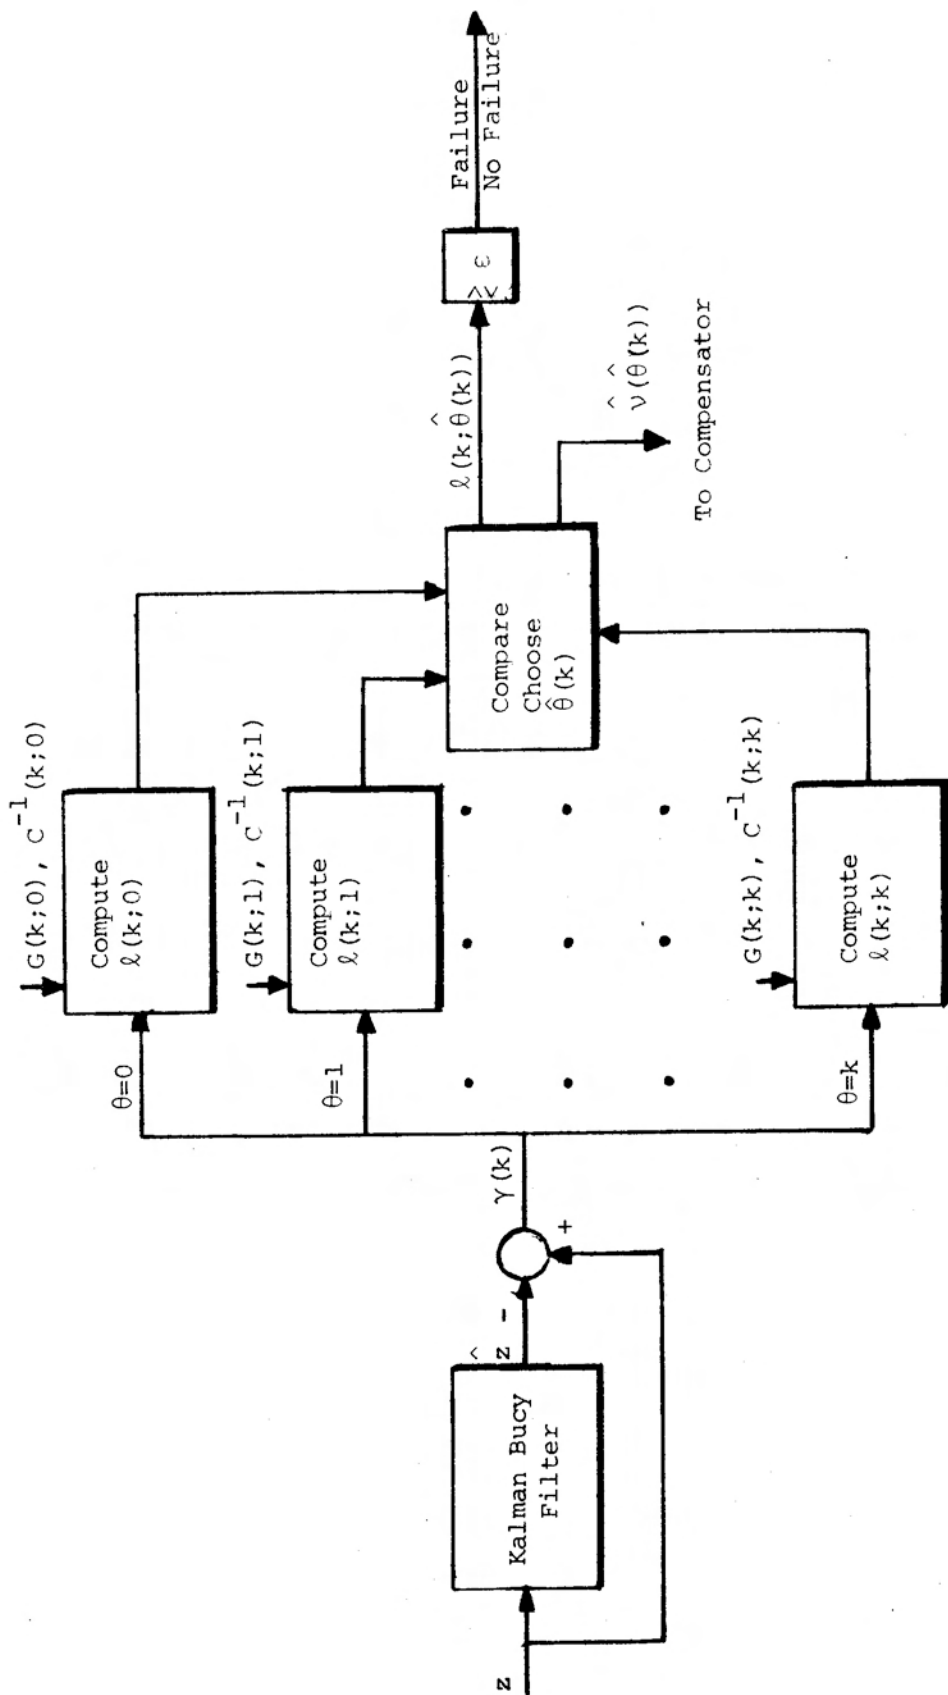


Figure 1.1 Full GLR Detector Scheme with Growing Bank of Filters

$$k-M < \theta \leq k-N \quad (1.38)$$

More on this and on the selection of M, N follow.

- The $\ell_i(k, \theta)$ can be updated, by means of a recursive algorithm, for every new residual $\gamma(k)$ produced in the Kalman filter.

The basic recursion is given by

$$C_i(k, \theta) = C_i(k-1, \theta) + G_i'(k, \theta) V^{-1}(k) G_i(k, \theta) \quad (1.39)$$

$$d_i(k, \theta) = d_i(k-1, \theta) + G_i'(k, \theta) V^{-1}(k) \gamma(k) \quad (1.40)$$

for each θ in (1.38). Note that the $C_i(k, \theta)$ does not depend on $\gamma(k)$ and can be precomputed and stored.

- If the system and filter are time-invariant, $G(k, \theta)$ and $C(k, \theta)$ depend only on $r = k - \theta$. This means that only $M - N + 1$ of each need to be generated when a window (1.38) is used. This reduces the computational burden significantly.
- The information generated by the GLR system can be very useful for compensation schemes to follow the detection of a failure. The estimates $\hat{\theta}$ and \hat{v} in particular can be used to determine updates for the filter estimate and covariance matrix, for example, in order that accurate functioning may continue. This would allow for the detection of multiple failures since the residuals are made to conform to the new situation which has been compensated for. More on the estimate and covariance incrementation can be found in Willsky-Jones [6].

As mentioned at the beginning of this section, there are modifications which can be incorporated into the GLR method. Each of these reflects different assumptions made on the types of failure of interest. In the preceding, the failure vector has been allowed to take on any value in R^n or R^p , for dynamics or sensor failures respectively. It is often the case however, that one has additional information concerning the possible failure modes.

In some cases one may, by physical considerations, conclude that a particular failure mode must necessarily be restricted to one of a particular set of directions in state or output space. This is the case, for example, when we know that only a particular actuator or sensor is subject to the assumed possible failures. This leads to the constrained GLR (CGLR) formulation, in which we constrain v in the form

$$v = \beta f_j \quad f_j \in \{f_1, \dots, f_r\} \quad (1.41)$$

where the f_j are the hypothesized failure directions and β is the unknown failure magnitude. One then computes

$$l_{ij}(k, \theta) = \frac{[f_j' d_i(k, \theta)]^2}{f_j' C_i(k, \theta) f_j} \quad (1.42)$$

which gives the log-likelihood ratio for a failure of type i in the direction f_j . (There may be different failure directions for different failure types, which implies that not all i - j pairs need be considered.) We then compute the maximum likelihood estimates (MLE's) \hat{i} , \hat{j} , $\hat{\theta}$ and the MLE of the failure magnitude

$$\hat{\beta}(k) = \frac{f'_j(k) d_i(k, \hat{\theta}(k))}{f'_j(k) C_i(k, \hat{\theta}(k)) f_j(k)} \quad (1.43)$$

By further constraining the failure such that the magnitude is also fixed

$$v \in \{v_1, \dots, v_s\} \quad (1.44)$$

we obtain the Simplified GLR (SGLR) formulation. Here one computes

$$\lambda_{ij}(k, \theta) = 2v_j' d_i(k, \theta) - v_j' C_i(k, \theta) v_j \quad (1.45)$$

and uses the largest of these to determine the presence of a failure and the best estimates of i , j and θ . In contrast to these two restricted GLR techniques, we will refer to the original version, where no information about v is assumed, as 'full' GLR.

This completes the introduction of the GLR-based approach to the detection of failures. The basic modeling assumptions, concerning the types of failure modes to which our attention will be directed, were presented. A hypothesis testing problem was formulated, a solution to which resulted in what we will refer to generically as the GLR system. It was also seen how variations of the main formulation result by imposing some further constraints on the failures to be detected. This provides a useful way to incorporate additional information that one might have on the possible failures in a manner which may simplify the procedure and alleviate the computational requirements.

We have not yet suggested what the expected performance of such a system might be or what appropriate measures of such performance are.

We will be dealing with questions like these throughout the rest of this thesis. But first we will present the system to which we will apply our failure detection scheme and on which our results are based. This problem has, in effect, been the "test-bed" for assessing GLR performance characteristics.

1.2.2 Application Problem: The Reduced F-8C Model

The essential concepts in the GLR technique for failure detection have been introduced in the previous section. In addition to our analysis of the performance, and limitations, of such a system, extensive use of simulations has been made. No attempt was made to generate complete and statistically significant data, or to provide any kind of final test of the system. The simulations did provide some confirmation of the analysis and helped us to develop intuition into GLR behavior. Considerable insight into the dynamics of failure detection was the result.

The simulations were made by applying the GLR system to a simplified model of the longitudinal dynamics of the F-8C aircraft. The model on which the detector design was based is a second-order, discrete-time version of the aircraft longitudinal dynamics at several particular flight conditions.

The motivation for using this model lies in the need to have a model of a concrete, physical system on which to implement the GLR detectors that would provide some common grounds for comparisons. This model provides a compromise in complexity between realism on the one hand and the amount of computation and ease of interpretation on the other. At this

stage in the development and study of the GLR approach to the detection of failures, some qualitative results were needed in order to understand its structure and performance characteristics.

Our model is derived from the continuous time representation of the F-8C longitudinal dynamics, linearized about flight condition 11. Flight condition 11 corresponds to level flight at an altitude of 20,000 feet, at a speed of Mach 0.6 in cumulus clouds. It can be represented by a 7-dimensional model:

$$\dot{x}(t) = Ax(t) + Bu(t) + G\xi(t) \quad (1.46)$$

where

- $x_1 = q$, pitch raterad/sec
- $x_2 = v$, velocity - V_0 [V_0 = Mach no. x speed of sound].....ft/sec
- $x_3 = \alpha$, (angle of attack) - (trim value).....rad.
- $x_4 = \theta$, pitch attitude.....rad.
- $x_5 = \delta_e$, (elevator deflection) - (trim value).....rad.
- $x_6 = \delta_{e_c}$, commanded elevator angle.....rad.
- $x_7 = w$, normalized wind disturbance rad.

State variables $x_5 = \delta_e$ and $x_6 = \delta_{e_c}$ take into account the actuator dynamics. The control variable is

$$u(t) = \dot{\delta}_{e_c}(t) \dots\dots\dots \text{rad/sec} \quad (1.47)$$

The matrices A, B and G are 7x7, 7x1 and 7x1, respectively. The wind disturbance, $x_7=w$, is modeled by the output of a first-order linear system driven by the white noise process $\xi(t)$. The wind model arises from a

given power spectral density

$$\Psi(\omega) = \frac{\sigma^2 h}{\pi V_0} \left\{ \frac{4}{4 + \frac{h}{V_0} \omega^2} \right\} \quad [\omega \neq \omega = x_7] \quad (1.48)$$

For flight condition 11, which is considered here, we have

$$h = 2,500 \text{ ft. (for altitudes } > 2500 \text{ ft.)}$$

$$V_0 = (0.6)(1,036.93 \text{ ft./sec}) = 622.15 \text{ ft./sec}$$

$$\sigma = 15 \text{ ft./sec (cumulus clouds)}$$

$$\omega \text{ (rad/sec)}$$

which determine the values of G and statistics of $\xi(t)$ in (1.46).

There are five sensor outputs given by z

$$z(t) = Cx(t) + \eta(t) \quad (1.49)$$

where

$$z_1 = z_q, \text{ pitch rate measurement}$$

$$z_2 = z_v, \text{ velocity error measurement}$$

$$z_3 = z_\theta, \text{ pitch attitude measurement}$$

$$z_4 = z_{\delta_e}, \text{ elevator angle measurement}$$

$$z_5 = z_{\omega_z}, \text{ normal acceleration measurement}$$

C is 5x7 matrix and $\eta(t)$ contains the five measurement noises which are white, mutually independent random variables with given statistics

$$E(\eta(t)) = 0 \quad (1.50)$$

$$E(\eta(t)\eta'(s)) = E\delta(t-s)$$

Reduction to a Two-Dimensional Model

Given this model, a simplified two-dimensional representation was obtained. (More on the F-8C model can be found in Athans & Dunn [16]). Some of the steps and considerations taken in the reduction of order of the model were:

- ignoring the input dynamics represented by x_5 and x_6 , as they are not the main variables of interest in an aircraft dynamics model.
- eliminating x_7 , the wind disturbance, as a variable and modeling its effects on the remaining ones by a white noise process (i.e., we ignore correlation in the wind).
- selecting the variables with the highest signal-to-noise ratios among the observations and ignoring the rest.
- using common sense and engineering intuition to correct and/or add for any other significant interactions.

The two state variables which were finally selected are

$x_1 = q$, the pitch rate

$x_2 = \alpha$, angle of attack - trim value

The last step was obtaining the corresponding discrete-time model to facilitate implementation on the digital computer as well as being consistent with our intended discrete-time analysis. A discretizing time step of $T = 0.03125$ sec (1/32 sec) was used. The resulting model for

flight condition 11 was

$$x(k+1) = \Phi x(k) + Gw(k) \quad (1.51)$$

$$z(k) = Hx(k) + Dv(k) \quad (1.52)$$

where Φ , G , H and D are the 2x2 matrices,

$$\Phi = \begin{bmatrix} 0.98258 & -0.14649 \\ 0.03059 & 0.97193 \end{bmatrix}, \quad G = \begin{bmatrix} 0.022596 & 0.0 \\ 0.004328 & 0.000226 \end{bmatrix}$$

$$H = \begin{bmatrix} 1.0 & 0.0 \\ 0.0 & 16.154 \end{bmatrix}, \quad D = \begin{bmatrix} 0.00873 & 0.0 \\ 0.0 & 0.06 \end{bmatrix} \quad (1.53)$$

The eigenvalues of Φ are

$$\lambda(\Phi) = 0.977 \pm j(0.0667) \quad (1.54)$$

and so the system is a stable one. The sequences $w(k)$ and $v(k)$ are zero mean, independent, white Gaussian sequences with unit covariance

$$E(w(k)w(j)') = I\delta_{kj}, \quad \delta_{kj} = \begin{cases} 0 & k \neq j \\ 1 & k = j \end{cases}$$

$$E(v(k)v(j)') = I\delta_{kj} \quad (1.55)$$

Note that the system (1.51), (1.52) is equivalent to that of (1.1), (1.2) if we set Q and R from (1.3), (1.4) to be

$$Q = GG', \quad R = DD' \quad (1.56)$$

The first step in designing a GLR system is obtaining the Kalman-Bucy filter (KBF) corresponding to the unfailed system. For the system given by (1.51)-(1.55), the steady-state version of the filter described

by equations (1.15) - (1.21) of the previous section gives us

$$\hat{x}(k|k-1) = \Phi\hat{x}(k-1|k-1) \quad (1.57)$$

$$\hat{x}(k|k) = \hat{x}(k|k-1) + \bar{K}\gamma(k) \quad (1.58)$$

with

$$\gamma(k) = z(k) - H\hat{x}(k|k-1) \quad (1.59)$$

$$\bar{K} = PH'V^{-1} \quad (1.60)$$

The filter error covariance matrix $P = \lim_{k \rightarrow \infty} P(k|k-1)$ is the steady-state solution of equations (1.18), (1.19) and gives \bar{K} , the constant, steady-state Kalman gain matrix. The covariance matrix for the residuals, V , is then given by (1.20)

$$V = HPH' + DD' \quad (1.61)$$

The resulting matrices \bar{K} , P and V are as follows for flight condition 11.

$$\begin{aligned} \bar{K} &= \begin{bmatrix} 7.5351 \times 10^{-1} & 4.6257 \times 10^{-2} \\ 1.3527 \times 10^{-1} & 1.2748 \times 10^{-2} \end{bmatrix} \\ P &= \begin{bmatrix} 5.6311 \times 10^{-4} & 1.0891 \times 10^{-4} \\ 1.0891 \times 10^{-4} & 2.2130 \times 10^{-5} \end{bmatrix} \\ V &= \begin{bmatrix} 6.3933 \times 10^{-4} & 1.7593 \times 10^{-3} \\ 1.7593 \times 10^{-1} & 9.3747 \times 10^{-3} \end{bmatrix} \end{aligned} \quad (1.62)$$

The predicted estimate, as given by equations (1.57)-(1.59), can be put in its recursive formulation

$$\begin{aligned}\hat{\mathbf{x}}(k+1|k) &= \Phi \hat{\mathbf{x}}(k|k-1) + \Phi \bar{\mathbf{K}} [z(k) - H \hat{\mathbf{x}}(k|k-1)] \\ &= \Phi (I - \bar{\mathbf{K}}H) \hat{\mathbf{x}}(k|k-1) + \bar{\mathbf{K}}z(k)\end{aligned}\quad (1.63)$$

with $\Phi(I - \bar{\mathbf{K}}H)$ given by

$$\Phi(I - \bar{\mathbf{K}}H) = \begin{bmatrix} 0.2620 & -0.8505 \\ -0.1239 & 0.7489 \end{bmatrix}\quad (1.64)$$

The eigenvalues of the filter are

$$\lambda_i(\Phi(I - \bar{\mathbf{K}}H)) = \begin{cases} 0.09966 & i=1 \\ 0.91188 & i=2 \end{cases}\quad (1.65)$$

Notice the absence of the control terms in the system (1.51), (1.52) and filter (1.57) - (1.59). They were not included since their effect on the residuals is cancelled in the KBF. The detector is therefore not affected and the analysis is simplified somewhat.

As was stated in section 1.2.1, our discussion on the performance of the GLR system will be limited to the detection of failures that can be represented by the failure models 1-4. Recall the four basic failure modes as applied to the simplified aircraft model.

1. State (Dynamics) Jump

$$\mathbf{x}(k+1) = \Phi \mathbf{x}(k) + Gw(k) + v\delta_{k+1,\theta}\quad (1.66)$$

2. State (Dynamics) Step

$$\mathbf{x}(k+1) = \Phi \mathbf{x}(k) + Gw(k) + v\sigma_{k+1,\theta}\quad (1.67)$$

3. Sensor Jump

$$z(k) = Hx(k) + Dv(k) + v\delta_{k,\theta}\quad (1.68)$$

4. Sensor Step

$$z(k) = Hx(k) + Dv(k) + v\sigma_{k,\theta} \quad (1.69)$$

where v is the failure and θ the time at which it occurs.

In the design of a GLR system under the assumptions of these failure models, some thought should be given to the proper use and interpretation of the failure vector v . For the cases where the constrained or simplified GLR are adequate or desirable this question is of particular importance. As their formulations implicitly depend on the hypothesized direction and/or magnitude of v , it is reasonable to expect that the careful selection of this failure vector will result in improved overall performance of the detection system. Ultimately, this is a question of physical and engineering considerations.

For the simulations, the failures were taken in orthogonal directions in failure space, $v \in R^n$ for 1, 2 and $v \in R^p$ for 3, 4 ($n=p=2$ for our simplified model). Thus, vectors of the form $(v_1 \ 0)$ and $(0 \ v_2)$ were considered for a range of values of v_1, v_2 thought to be of most interest. Let us take a closer look at the kinds of situations that these failures might represent in the case of the simplified F-8C model.

Recall that $x_1 = q$ is the pitch rate and that $x_2 = \alpha$ is the angle of attack deviation from the trim value. Thus a state jump of the form $(v, 0)^T$ or a state step of the form $(0 \ v_2)^T$ might be used to model the effect of a sudden wind shear that leads to an increasing angle of attack. On the other hand, a jump of the form $(0 \ v_2)$ could be used to model a relatively long-term upward or downward gust that initially manifests

itself as a shift in α . A step of the type $(v_1 \ 0)$ could arise from an elevator failure.

In the observation equation, failure models 3 and 4, we have a similar situation. A failure $(v_1 \ 0)$ may model a bad data point in the measurement of q for the jump case (3), or it may represent a permanent bias in that measurement for the case of a step (4). Such a bias might be the result of a component failure in a sensor, for example. By analogy, the same may be said about a $(0 \ v_2)$ failure, which then refers to the measurement of α .

Table 1.1 summarizes the failure schedule implemented in the simulations. State and sensor failure magnitudes are given in terms of the standard deviations of the noises affecting each variable (1σ equals one standard deviation of the noise). For clarification purposes σ_q, σ_α will denote the standard deviations of the noise in the dynamic equations for q and α , respectively, and $\sigma'_q, \sigma'_\alpha$ will denote the analogous noise level for the measurements of q and α . In table 1.2 the values of these 1σ noise levels are given.

DETECTOR AND FAILURE TYPE	$(v_1 \ 0)^T$	$(0 \ v_2)^T$
STATE JUMP	$1\sigma_q, 5\sigma_q, 10\sigma_q, 20\sigma_q$	$1\sigma_\alpha, 5\sigma_\alpha, 10\sigma_\alpha, 20\sigma_\alpha$
STATE JUMP	$.1\sigma_q, .5\sigma_q, 1\sigma_q, 5\sigma_q$	$.1\sigma_\alpha, .5\sigma_\alpha, 1\sigma_\alpha, 5\sigma_\alpha$
SENSOR JUMP	$1\sigma'_q, 5\sigma'_q, 10\sigma'_q, 20\sigma'_q$	$1\sigma'_\alpha, 5\sigma'_\alpha, 10\sigma'_\alpha, 20\sigma'_\alpha$
SENSOR STEP	$.1\sigma'_q, .5\sigma'_q, 1\sigma'_q$ $5\sigma'_q, 10\sigma'_q, 20\sigma'_q$	$.1\sigma'_\alpha, .5\sigma'_\alpha, 1\sigma'_\alpha$ $5\sigma'_\alpha, 10\sigma'_\alpha, 20\sigma'_\alpha$

TABLE 1.1 Set of Failures Simulated

STATE NOISE		SENSOR NOISE	
σ_q	σ_α	σ'_q	σ'_α
1 std. dev. = 2.2596×10^{-2}	1 std. dev. = 4.3335×10^{-3}	1 std. dev. = 8.7298×10^{-3}	1 std. dev. = 6.0×10^{-2}

TABLE 1.2 Noise Levels in Standard Deviations

For jump failures nothing under 1σ was looked at since those jumps would be undistinguishable from the noise. Such failure magnitudes were considered for step failures since they are detectable because their sustained presence provides more information as time passes. This will be seen in more detail when we examine the signature matrices $G_i(k, \theta)$ for the various failure modes.

This completes our presentation and description of the system model with which our numerical results of the simulations were generated. We will have more to say as we go along.

1.3 SUMMARY OF THE GLR APPROACH AND OVERVIEW

In the introduction we tried to define the goals of this effort: to evaluate the performance of the GLR-based technique for the detection of failures in the light of some analysis, and 'experience' via simulations of its application to a specific dynamical system.

In section 1.2.1 the basic modeling assumptions and issues that will concern us here were presented. A number of failure models were introduced, and what we will refer to as the GLR system was developed for the cases to which this work will be restricted. A simple but adequate (for our purposes) dynamic system model was discussed.

In the following we examine the performance of a GLR detection system as applied to a specific example. In combination with some analytic performance measures developed by Chow [12] and Willsky-Jones [5], the aim here is to develop insight into the practice of failure detection and the difficulties that arise.

Chapter 2 will present some results obtained by a straightforward application of the method in section 1.2.1 to the system presented in 1.2.2. In a sense, conditions are set to be ideal: the GLR detection system is based on exactly the same dynamics as the system which undergoes the failures. Furthermore, the correct detectors for the particular failure modes treated are implemented. After some discussion in section 2.2 on the question of performance measures and GLR characteristics, in section 2.3 some actual results are presented and interpreted, providing an initial evaluation of the failure detecting capabilities of the GLR.

Chapter 3 takes a look at some of the limitations on the capability of the GLR technique to correctly detect and identify the actual failure mode present in a system. The question of distinguishing among several possible failure modes is of central importance here. Some simulation results from the GLR implementation in Chapter 2 are discussed. Knowledge of these limitations obtained by increasing our understanding of the method, suggests ways of overcoming some of the difficulties and of realistically evaluating our ability to handle failures.

After some familiarity with the GLR system performance characteristics is developed in Chapters 2 and 3, in Chapter 4 our attention

focuses on the sensitivity of such performance to modeling errors. Important questions on the final relevance of the GLR system for reliable failure detection in applications where modeling errors may be substantial are raised and an attempt is made to look for a way out of the difficulties.

Finally, in Chapter 5, we try to put our results in perspective and to draw some conclusions. Discussion of the merits and drawbacks of the GLR approach is followed by specific suggestions on means to overcome the major limitations on the achievable performance levels. Areas of potential and of promising future work are pointed out.

CHAPTER 2

CORRECT DETECTION UNDER MATCHED CONDITIONS

2.1 MATCHED CONDITIONS AND CORRECT DETECTION

The motivation for the use of the GLR technique was explained in Chapter 1 and a framework for its use in the area of failure detection was developed. As a tool for detecting and identifying the events that we characterize as failures, the GLR technique offers some advantages. The estimation of θ and ν , for example, provides some essential information if compensation is required subsequent to the detection of a failure. The possibility of some analytical evaluation of the anticipated performance is another reason for the appeal of this approach.

In this chapter we examine the performance of an application of the full GLR technique of section 1.2.1 to the detection and identification of failures in the system presented in section 1.2.2. The terms 'matched conditions' will be used to denote the modeling assumptions on which the detection system design is based. We assume perfect knowledge of the system parameters and noise statistics. This is in contrast to the situation referred to as 'mismatched conditions', considered in Chapter 4, where modeling errors are allowed.

The main purpose here is to develop some intuition on the theoretically optimal performance that can be extracted from a simple application of the GLR method to the detection of failures. Our approach is qualitative in nature. It is meant to complement the work done by Chow [12] of a more quantitative kind. In addition to the overall understanding

we expect to develop, this will also provide a reference performance level with respect to which one can measure the degradation in detection and identification which takes place in a more realistic environment, where the system does not perfectly correspond to the model.

Furthermore we also explicitly assume that the failure mode is known and, consequently, that we are able to implement a GLR detector based on that mode. This is what is meant by correct detection, and no judgment on the resulting performance is implied. It is clear that if we are certain about the kind of failure anticipated, then we can do at least as well, if not better, as when the failure type must first be isolated. Chapter 3 considers the difficulties that may arise when the failure mode is unknown.

The results under these conditions indicate how well the GLR detection system can perform in under ideal conditions. Therefore expectations about what can, and cannot, be achieved will be firmly grounded. Section 2.2 presents some measures with which to evaluate the detection performance of the GLR technique. In addition, we begin to look at some questions on the detectability of the different failure modes. Then, in section 2.3, some results are shown from an application to the simplified F-8C model. Section 2.4 summarizes the results.

2.2 THE DETECTION PERFORMANCE OF THE GLR

2.2.1 Performance Analysis

We now begin our study of the performance of a GLR failure detection

system. In order to evaluate the performance and reliability of such a system, we need to select some indices that measure its quality. We look at the decision procedure in the GLR technique and at ways to characterize its dynamic behavior. This helps us to choose the detector parameters and to resolve tradeoffs which are inherent in this method. Our aim is to understand well how the GLR technique works.

If a condition of failure exists, changes are induced in the behavior of the innovations of the Kalman filter. The GLR detectors continually compute the correlation between the actual residuals and their anticipated behavior under each of the failure hypotheses. This concept is expressed by equations (1.26), (1.27) and (1.33)-(1.37), which describe the GLR detectors and procedure. The quantity $\hat{\ell}_i(k, \hat{\theta})$ is a normalized measure of that correlation between the system residuals and the residuals under the best estimate of the failure mode.

It is clear that for any configuration of the parameters $\{M, N, i, \epsilon\}$ which specify the GLR detectors, it is the statistical behavior of the random variable $\hat{\ell}_i(k, \theta)$ which defines the performance of the system. Some of the features or characteristics of this behavior, such as delay times to detection and false alarm rates, for example, can be studied by examining the evolution in time of the distribution of values that this variable takes under various conditions. Many interesting and important questions can be formulated in terms of certain events. These questions can then be transformed into the calculation of probabilities. Such an analysis allows us to study in detail the numerous

tradeoffs that one faces in designing an acceptable failure detection system.

Consider now a situation where a GLR detector is designed for the failure mode which coincides with that of the actual failure taking place. For simplicity we will drop $i \in \{1, 2, 3, 4\}$, the failure type, in the following discussion. In [12] Chow shows that for the full GLR $\ell(k, \theta)$ is a non-central, chi-squared (χ^2) random variable with n degrees of freedom for state failure detectors, or p degrees of freedom for sensor failure detectors (n is the dimension of the state and p the number of measurements). The mean or expected value of $\ell(k, \theta)$ is $n + \delta^2$ (or $p + \delta^2$ for the sensor cases) where δ^2 , the non-centrality parameter, is given by

$$\delta^2 = \delta^2(k, \theta) = v' C(k, \theta) v \quad (2.1)$$

when a failure of size v occurs.

This can be readily seen by examining the equations from the GLR. From section 1.2.1 we recall

$$\gamma(k) = \tilde{\gamma}(k) + G(k, \theta) v \quad (2.2)$$

$$d(k, \theta) = \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \gamma(j) \quad (2.3)$$

$$C(k, \theta) = \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) G(j, \theta) \quad (2.4)$$

and

$$\ell(k, \theta) = d'(k, \theta) C^{-1}(k, \theta) d(k, \theta) \quad (2.5)$$

Notice that in (2.3) and (2.4) the lower limit in the summation has been replaced by θ .

By expressing (2.3) in terms of $\tilde{\gamma}(k)$ one gets

$$d(k, \theta) = \tilde{d}(k, \theta) + C(k, \theta)v \quad (2.6)$$

$$\tilde{d}(k, \theta) = \sum_{j=\theta}^k G'(j, \theta)v^{-1}(j)\tilde{\gamma}(j) = \{d(k, \theta) | \gamma = \tilde{\gamma}\} \quad (2.7)$$

This follows by simple substitution of (2.2) into (2.3). Here $\tilde{d}(k, \theta)$ is a zero mean random vector with covariance matrix $C(k, \theta)$, as may be easily verified given the properties of $\tilde{\gamma}(k)$. When a failure is present, $d(k, \theta)$ has the same covariance, but the mean value becomes $C(k, \theta)v$ as (2.6) shows. Finally, using (2.6) in equation (2.5)

$$\begin{aligned} \mathcal{L}(k, \theta) &= [\tilde{d}(k, \theta) + C(k, \theta)v]'C^{-1}(k, \theta)[\tilde{d}(k, \theta) + C(k, \theta)v] \\ &= \tilde{d}'(k, \theta)C^{-1}(k, \theta)\tilde{d}(k, \theta) + v'C'(k, \theta)C^{-1}(k, \theta)\tilde{d}(k, \theta) \\ &\quad + \tilde{d}'(k, \theta)C^{-1}(k, \theta)C(k, \theta)v + v'C'(k, \theta)C^{-1}(k, \theta)C(k, \theta)v \\ &= \tilde{\mathcal{L}}(k, \theta) + 2v'\tilde{d}(k, \theta) + v'C(k, \theta)v \end{aligned} \quad (2.8)$$

where

$$\begin{aligned} \tilde{\mathcal{L}}(k, \theta) &= \tilde{d}'(k, \theta)C^{-1}(k, \theta)\tilde{d}(k, \theta) \\ &= \{\mathcal{L}(k, \theta) | \gamma = \tilde{\gamma}\} \end{aligned} \quad (2.9)$$

and where use is made of the fact that $C(k, \theta)$ is a symmetric matrix.

Using our knowledge that $\tilde{d}(k, \theta)$ has zero mean value and taking expectations on both sides of equation (2.8),

$$\begin{aligned}
 E(\ell(k, \theta)) &= E(\tilde{\ell}(k, \theta) + 2v'\tilde{d}(k, \theta) + v'C(k, \theta)v) \\
 &= E(\tilde{\ell}(k, \theta)) + 2v'E(\tilde{d}(k, \theta)) + v'C(k, \theta)v \\
 &= E(\tilde{\ell}(k, \theta)) + v'C(k, \theta)v
 \end{aligned}
 \tag{2.10}$$

The linearity property of the expectation operator has been used here. By identification with the result from [12] previously mentioned and (2.1) we conclude

$$E(\ell(k, \theta)) = n + \delta^2(k, \theta) \tag{2.11}$$

If we remember the decision rule which completes the GLR procedure from (1.37), we can make some immediate observations. Given the updated $d(k, \theta)$, once $\gamma(k)$ is computed, as in (1.40) for θ in the range $k-M < \theta \leq k-N$ as in (1.38), then $\hat{\theta}(k)$ is chosen to be the θ which maximizes the log-likelihood ratio in (1.36). A failure is declared to have occurred at time $\hat{\theta}$ if $\ell(k, \hat{\theta}(k))$ exceeds a threshold ϵ :

$$\begin{array}{l}
 \text{failure declared} \\
 \ell(k, \hat{\theta}(k)) < \epsilon \\
 \text{no failure}
 \end{array}
 \tag{2.12}$$

By studying the average behavior of $\ell(k, \theta)$ as given by (2.11) it is possible to develop some ideas on how well the decision rule (2.12) responds to failures of different kinds. This, in combination with the estimate of the failure in (1.32), can provide a picture of how the GLR system works and of its reliability. This means that one must study in detail the changes in $\delta^2(k, \theta)$. By assuming for the moment that the value of the threshold ϵ has been selected in some appropriate manner,

this kind of analysis will in the end provide the knowledge necessary to choose ϵ in a way that assures the best detection performance.

The two quantities that determine the evolution in time of $\delta^2(k, \theta)$, v and $C(k, \theta)$, illuminate some basic features of the GLR technique. The expression for δ^2

$$\delta^2(k, \theta) = v' C(k, \theta) v \quad (2.13)$$

already indicates two important facts. The way the failure vector v comes into play appeals to our intuition that somehow larger failures should be easier to detect. This quadratic dependence of δ^2 on the actual failure is modulated by the generally time-varying matrix $C(k, \theta)$. It is this quantity, the information matrix, which brings to light how the GLR approach uses information about the dynamic effect of failure on the system residuals. Here lies the essence of the flexibility that the GLR approach can bring to the failure detection problem: full use is made of the characteristic response of the physical system and Kalman filter to each failure mode.

The matrix $C_i(k, \theta)$, given by equation (2.4), measures the information available in the residuals at time k from a failure of type i occurring at time θ . It is therefore worthwhile to study how it changes for varying k and θ for the different failure modes. Such a study will show how much impact each failure mode has on the GLR detectors and, consequently, the degree of detectability of various failures. In fact, this quantity can be shown to be an observability matrix. For these reasons in the next section we take a closer look

at $C(k, \theta)$.

Complementary to the study of the average performance of the GLR by way of the non-centrality parameter δ^2 , one can define and compute various probabilities of the event that $\ell(k, \theta)$ takes on certain values or follows given patterns with k or θ varying. Three useful probabilities defined in [12] by Chow are: correct detection, wrong time and false alarm probabilities, given respectively by

$$P_D = \text{Prob} \{ \ell(k, \theta) > \varepsilon | v, \theta = \theta_t \} \quad (2.14)$$

$$P_{\theta/\theta_t} = P_{WT} = \text{Prob} \{ \ell(k, \theta) | v, \theta \neq \theta_t \} \quad (2.15)$$

$$P_F = \text{Prob} \{ \ell(k, \theta) > \varepsilon | v=0 \} \quad (2.16)$$

The correct detection probability P_D measures how sensitive the detectors are by giving the probability at time k of detecting a failure v which started at time θ_t . The wrong-time probability gives some idea of the detectability of a failure by indicating the accuracy of GLR in estimating the time of failure θ_t and the persistence of the failure effects for $\theta > \theta_t$. Finally, P_F gives the probability that a failure will be declared when no failure is present. It shows the sensitivity of the GLR detectors to the noises in the system.

Although the fact that for the full GLR the $\ell_i(k, \theta)$ are chi-squared random variables makes analysis difficult, one can compute the values of these probabilities for cases of interest (see the appendix in [12]). Figure 2.1 gives an example of how useful such computations can be. It is a plot of P_D for different values of the threshold and of




Fig. 2.1 $P_D(\delta^2, \epsilon)$ for a Variable with 2 Degrees of Freedom

the non-centrality parameter for a χ^2 random variable with two degrees of freedom. This is the case for our example, which consists of two states and two observations. Notice that the curve corresponding to $\delta^2 = 0$ is that of P_F , which is fixed once the threshold is specified. The correct detection probability, however, is an increasing function of δ^2 as shown. For a given threshold, it is the value of δ^2 that determines the probability of detection.

Figure 2.1 points out one of the basic tradeoffs in the GLR technique. By selecting ε large enough the false alarm rate can be reduced significantly, an appealing feature for failure detection. In doing so, however, P_D is also reduced for any value of δ^2 as the curves show. Thus one is only able to detect failures for larger magnitudes of v . This is true at a particular time step, since over a time interval $\delta^2(k, \theta)$ may grow enough to exceed the threshold. Two important concepts in interplay can be observed here. In eliminating false alarms via a higher threshold, the same degree of detectability is achieved for larger failures only. On the other hand, the dynamic character of $\delta^2(k, \theta)$ for a given failure mode may be such that the value of $\ell(k, \theta)$ required for detection is reached rapidly. This points out how studying $C(k, \theta)$ in (2.13) is valuable for understanding GLR failure detection and also, in particular, for selecting the optimal value of the threshold.

We have only mentioned some of the possible performance measures which are useful to evaluate and understand a GLR detection system. Different indices can be used to study other aspects of detector performance,

such as correlations of the $\ell_i(k, \theta)$ in time for various i , k and θ , for example. We will introduce some of these as we touch on the areas where they are most relevant. The next section examines some important elements of the GLR technique in more detail.

2.2.2 The Signature and Information Matrices

The last section focused on the characteristics of $\ell_i(k, \theta)$ and on the way these change under different failure situations. It was seen that this random variable can be studied by examining how the non-centrality parameter $\delta^2(k, \theta)$ responds to different failures. Thus one can develop some qualitative understanding of the dynamics of GLR-based failure detection and the detectability of the various failure modes. Also of practical interest is the development of some criteria for selecting those detector parameters, such as the threshold and the dimension of the sliding time window, which greatly affect the performance. With these concerns in mind, we will concentrate on the information matrix, $C_i(k, \theta)$.

We have already commented on why this quantity is of importance in relation to $\delta^2(k, \theta)$ and, consequently, to the statistical behavior of the log-likelihood ratio. In this section we will take a close look at the values of this matrix for different i , k and θ . We will proceed by first looking at $G_i(k, \theta)$, the failure signature matrix, which is central to a study of $C_i(k, \theta)$, as (2.4) makes evident.

The Failure Signature Matrix, $G_i(k, \theta)$

The basis for the GLR approach to failure detection lies, as seen in Chapter 1, in the changes that failures induce in the residuals of the filter in the system. These changes are specific to each failure mode and, for the cases being considered here, they can be expressed as the alternative hypotheses.

$$H_i : \gamma(k) = \tilde{\gamma}(k) + G_i(k, \theta)v \quad i = \{1, 2, 3, 4\} \quad (2.17)$$

where v is the failure vector and $G_0(k, \theta) \triangleq 0$ for all k and θ . Hence, the time histories of these signature matrices are very informative with respect to the degree of detectability of each failure mode. For a given failure v we can anticipate the capability of the GLR to accurately detect and estimate it if we understand how the $G_i(k, \theta)$ behave. Expressions for the signature matrices for the four failure modes are included in Appendix A. These show the dependence of $G_i(k, \theta)$ on the system and filter parameters.

We have computed the G_i matrices for failure modes $i = \{1, 2, 3, 4\}$ for the reduced F-8C model. They are plotted in Figures 2.2-2.5 as a function of $r=k-\theta$ for the steady-state filter case. Before discussing these curves, let us concentrate briefly on the elements of these matrices.

Consider $\Delta\gamma$, the deviation in the residual due to a failure v , from the expression (2.17)

$$\Delta\gamma(k) = \gamma(k) - \tilde{\gamma}(k) = G_i(k, \theta)v \quad (2.18)$$

or, in more detail,

$$\begin{bmatrix} \Delta\gamma_1(k) \\ \Delta\gamma_2(k) \end{bmatrix} = \begin{bmatrix} g_{11}(k-\theta) & g_{12}(k-\theta) \\ g_{21}(k-\theta) & g_{22}(k-\theta) \end{bmatrix} \begin{bmatrix} v_1 \\ v_2 \end{bmatrix} \quad (2.19)$$

The i^{th} row of $G(k-\theta)$, g_{i1} and g_{i2} , gives the effect of the failure on the i^{th} component of the residuals. Alternatively, the j^{th} column of the signature matrix, g_{1j} and g_{2j} , gives the effect of the j^{th} failure component, v_j , on the different elements of the vector of residuals.

For the two-dimensional F-8C model with failures taken in orthogonal directions, we get the following. For pitch rate failures of the form $v = [v_1 \ 0]^T$,

$$\Delta\gamma_1(k) = g_{11}(k-\theta)v_1 \quad (2.20)$$

$$\Delta\gamma_2(k) = g_{21}(k-\theta)v_1 \quad (2.21)$$

while for angle-of-attack failures, $v = [0 \ v_2]^T$,

$$\Delta\gamma_1(k) = g_{12}(k-\theta)v_2 \quad (2.22)$$

$$\Delta\gamma_2(k) = g_{22}(k-\theta)v_2 \quad (2.23)$$

Comparing the $g_{ij}(k-\theta)$ for different failure modes provides information about the distinguishability of failures of different modes. This is a very important consideration and will be discussed in more detail in Chapter 3.

Figure 2.2 and 2.3 are plots of the signatures for the aircraft example, where the failure modes correspond to state jump and sensor jump respectively. Notice, first of all, that the elements of both matrices die out as r , the elapsed time from the occurrence of a failure, increases. This is not surprising since both the system and the filter are stable. So, for the jump failures, the signatures reflect the impulse response of the system and filter. A sensor failure enters the filter directly while a state failure does so only in proportion to its effect on the output of the system. The figures indicate that the sensor jump signatures tend to zero much faster than those of a state jump failure. This is in agreement with the fact that the time constants of the filter are smaller than those of the system for the reduced aircraft model.

The signatures for the cases of state step and sensor step failures present a different picture however. Figures 2.4 and 2.5 contain plots of their respective elements (for a longer interval in r). Some of these curves now grow with r , rather than dying out. This is a reflection of the different nature of these failure modes. The sustained effect of a step failure on the residuals is in marked contrast to the case with jump failures. Given that the residuals are the inputs to the detectors, this implies that step failures are more detectable than jump failures. Recall that in the GLR detectors, weighted sums of the residuals are sequentially computed over the interval determined by




Fig. 2.2 State Jump $G(r)$




Fig. 2.3 Sensor Jump $G(r)$




Fig. 2.4 State Step $G(r)$




Fig. 2.5 Sensor Step $G(r)$

the moving window parameters M and N , as in equation (2.3). The persistent, if not increasing, deviations in the residuals apparent in Figures 2.4 and 2.5 suggest that higher values of $\ell_i(k, \theta)$ are achieved and that, in fact, these likelihood ratios grow in time. Consequently, better detection should follow, for a fixed threshold, than for the case of jump failures. By the same token, it is reasonable to expect that the estimation of v will be better for step failures than for jump failures.

It is important to realize what the shapes of the different signatures tell us about the system and the failures. Recall that the signatures indicate how the Kalman filter yields dynamically-related estimation errors when a failure is present. These properties of $G_i(k, \theta)$ translate directly, as will be emphasized, into the properties of the detection performance of the GLR.

Consider again the plots of $G_i(k, \theta)$. Notice how for the state failures (Figures 2.2 and 2.4) the dominant values correspond to g_{22} and g_{21} , in order of importance. As equations (2.20)-(2.23) indicate, these g 's give the components of the prediction error in α due to failures in α (g_{22}) and q (g_{21}). It is interesting that the main error caused by the q failures is in the estimate of α (as $g_{21} > g_{11}$). This reflects the way in which the effect of the failure propagates through the system and the filter. What the relative magnitudes of the curves for g_{22} and g_{21} imply is that state failures in α are more detectable than similar failures in q . Or, the same degree of detectability is achieved only for relatively larger failures in q than in α . In the case of sensor

failures (Figures 2.3 and 2.5), on the other hand, the converse is true: failures in q are more detectable than those in α . (Notice that, in general, $g_{i1} \geq g_{i2}$, $i = 1,2$). Here again, the off-diagonal elements of the signatures (g_{21}) play an important role for failures in q . For these failures the increase in q (the pitch rate) slowly integrates into α (angle of attack) in the filter. Thus, the estimate of α becomes incorrect to the extent that it leads to the dominant error. So the failure in q actually manifests itself mostly in the inconsistency between the true and predicted values of α . The GLR technique exploits these dynamic characteristics which are a consequence of the properties of the Kalman filter.

The failure signatures also provide information concerning the capability of the GLR to correctly identify and estimate θ_t , the true failure time. The plots suggest that jump failures might lead to better estimates $\hat{\theta}$ than step failures. This is due to the largely localized effect on the residuals at the initial times. For step failures, which may become more detectable with time, the initial values of the signatures are not the largest in general.

That is, the fact that the effects of some step failures persist in the residuals is clearly seen by a signature that persists. However, by its very nature, a persistent signature has an autocorrelation function that is not particularly peaked. Thus, for such failures we expect good detection performance but anticipate inaccuracies in our estimate of θ_t . To see this point, the reader may want to compute the

autocorrelation functions of g_{21} in Figure 2.3 and g_{21} in Figure 2.5. This clearly indicates the increased detectability of step failures (as indicated by the magnitude of the autocorrelation), and the increased problem in resolving the failure time θ (as indicated by the peakedness of the autocorrelation).

Finally, the signatures tell us about the relative distinguishability of the failure modes. If two different failure modes have similar signatures, potential difficulties arise in discriminating between them. In full GLR the estimate \hat{v} is free to take any value in failure space (e.g., a combined α and q sensor failure). Hence, the incorrect detector can select \hat{v} as that v which would best account for the obtained likelihood ratios although this \hat{v} might not be physically meaningful. This suggests that other GLR formulations (such as CGLR or SGLR, Chapter 1) might have less of a problem in choosing the correct failure mode for these cases. Since the failure estimate is constrained in direction and/or magnitude, different values of the likelihood ratios should result. Looking at the signatures, for example, one can anticipate possible difficulty in telling apart a sensor step in q from a state step failure in either q or α . The signature curves for these cases can look somewhat alike with proper scaling (which is done by \hat{v}). Also, there is a potential problem in distinguishing between a q state jump and an α sensor step. Notice the qualitative similarity between g_{22} , g_{12} in Figure 2.5 and g_{11} , g_{21} in Figure 2.2 (with the proper scaling).

This distinguishability problem clearly involves cross-correlations of the various signatures. We will discuss this topic further in Chapter 3.

The Information Matrix, $C_i(k, \theta)$

One way to interpret the statement that $C_i(k, \theta)$ measures the information in the residuals at time k from a failure of type i at time θ is to view it geometrically. $C_i(k, \theta)$ has the form of an inner product or projection operator. For example, the non-centrality parameter can be expressed, using (2.4),

$$\begin{aligned}
 \delta_i^2(k, \theta) &= v' C_i(k, \theta) v \\
 &= v' \sum_{j=\theta}^k G_i'(j, \theta) V^{-1}(j) G_i(j, \theta) v \\
 &= \sum_{j=\theta}^k v' G_i'(j, \theta) V^{-1}(j) G_i(j, \theta) v \\
 &= \langle G_i(\cdot, \theta) v, G_i(\cdot, \theta) v \rangle_{V^{-1}(\cdot)} \\
 &= \|G_i(\cdot, \theta) v\|_{V^{-1}(\cdot)}^2
 \end{aligned} \tag{2.24}$$

where the symbol (\cdot) represents the time index which runs from $j=\theta$ to $j=k$. In this case the projection of the failure signature onto itself gives a squared-norm measure of the deviation Δy in the residuals over the time interval from θ to k . It is reasonable then that a failure

mode resulting in a lasting presence in the residuals process, such as one of the step failures, will lead to large values of the non-centrality parameter. This corresponds to a high degree of detectability as the probability of correct detection P_D increases rapidly.

Suppose the failure vector v is in a given direction, $v=[v_1 \ 0]'$. Then $\delta^2(k, \theta)$ is just a particular element of the information matrix with a scaling factor

$$\begin{aligned} \delta^2(k, \theta) &= [v, \ 0] \begin{bmatrix} C_{11}(k, \theta) & C_{12}(k, \theta) \\ C_{21}(k, \theta) & C_{22}(k, \theta) \end{bmatrix} \begin{bmatrix} v \\ 0 \end{bmatrix} \\ &= v_1^2 C_{11}(k, \theta) \end{aligned} \quad (2.25)$$

Similarly for the other direction in the two-dimensional case. This means that the elements of $C_i(k, \theta)$ give the directional sensitivity of a detector for failure mode i to different failures v . The off-diagonal terms measure the effect on detectability of non-orthogonality in the failure directions. The study of the changing values of the information matrix is then an additional way of learning about the failure modes and the capability of the GLR method for detecting them.

If the signatures for the jump failures tend to zero as the elapsed time increases, then the information matrix reaches a steady-state at a fast rate. This is due to the quadratic form of $C_i(k, \theta)$ and to the fact that additional terms contribute less as r increases. The contrary is true in the case of step failures, where the signatures

either grow or remain at non-zero values. Here the information matrix does not necessarily reach a steady-state, as more information is accumulated with every term in (2.4).

Figures 2.6 - 2.9 illustrate this. They are plots of the non-centrality parameter and correct detection probability versus the elapsed time r . They correspond to jump failures (Fig. 2.6, 2.7) of magnitude 5σ and $5\sigma'$ for the state and sensor detectors and to step failures (Fig. 2.8, 2.9) of 1σ in the state and $5\sigma'$ in the sensors. Figures 2.6 and 2.7 are related to the failure signatures in Figures 2.2 and 2.3. We see that in fact δ^2 reaches steady state almost immediately for the sensor jumps while it grows at a diminishing rate for the state jump failures.

Similar comments can be made about the curves in Figures 2.8 and 2.9 in reference to the signatures in Figures 2.4 and 2.5. The differences and similarities among these curves are very informative of the performance of the GLR detectors. Careful analysis and observation reveals how the specific and characteristic dynamics of the system-filter combination are reflected in these varying degrees of detectability.

For example, consider the δ^2 curve for a sensor step in α , shown in Figure 2.9. While the δ^2 curves for the other state and sensor step failures show a non-diminishing rate of increase, this one reaches steady-state eventually, for the same time interval. The angle of attack α is approximately the integral of the pitch rate q . To the extent that this is true, a state jump in q is similar to a step increase in the




Fig. 2.6 5σ State Jump δ^2 and P_D with $\epsilon=14$ ($P_F=0.000912$)




Fig. 2.7 $5\sigma'$ Sensor Jump δ^2 and P_D with $\epsilon=14$ ($P_F=0.000912$)




Fig. 2.8 10 State Step δ^2 and P_D with $\epsilon=14$ ($P_F=0.000912$)




Fig. 2.9 $5\sigma'$ sensor step δ^2 and P_D with $\epsilon=14$ ($P_F=0.000912$)

value of α at the input to the filter. This can be seen qualitatively in the fact that the sensor step δ^2 curve for α in Figure 2.9 is not unlike that of a state jump in q in Figure 2.6. Analogous comments hold for the signature curves $g_{12}(r)$ and $g_{22}(r)$ in Figure 2.5 and the curves $g_{11}(r)$ and $g_{21}(r)$ in Figure 2.2, as was mentioned in discussing the signatures. This points out how the similarity in the signatures for two different failure modes, as suggested earlier, may lead to problems in distinguishing failures. The same comment can be made from these curves about sensor step failures in q and state step failures of both kinds. The similarity in the growth of δ^2 with time is remarkable, although not surprising considering the corresponding signatures.

For linear time-invariant systems, it makes sense to speak of a steady state. The rate of convergence of $C_i(k, \theta)$ to such a limit is a useful indicator of the length of the window to be chosen, over which $\hat{\theta}$ is to be selected in maximizing $\ell_i(k, \theta)$. For jump failures in the sensors it is clear that if they are not detected right away, the odds on doing so are not improved by waiting. The non-centrality parameter and P_D reach their final values almost immediately. In contrast to this, for step failures δ^2 is increasing in general with the elapsed time and, therefore, a longer window length would indeed enhance the possibility of detection significantly. It must be noted of course, that the value of the parameter M thus selected depends also on other

considerations. The actual size of the failure v must be taken into account. For a large enough v detection can be immediate for any failure mode and implementing a large sliding window (large M) is wasteful and unneeded computationally.

All references made in the above discussion to the non-centrality parameter $\delta^2(k, \theta)$ assume we know θ to be the time of failure. If so, then $\delta^2(k, \theta)$ is a measure of the average behavior of $\ell(k, \theta = \theta_t)$. In fact, however, θ_t is considered unknown by the GLR detectors until the decision is made by first maximizing $\ell(k, \theta)$ over $k-M < \theta < k-N$. For each θ in the moving window there is a matched filter $d(k, \theta)$ which will give the likelihood $\ell(k, \theta)$ that a failure began at that time. The $d(k, k-N)$, $d(k, k-N-1)$, ..., $d(k, k-M)$ "scan" the recent past for possible failures. In the maximization operation, (1.36), the GLR detector in effect 'slides' the signature $G(k, \theta)$ over the time history of the residuals in the interval given by the window. Then $\hat{\theta}(k)$ is selected as the 'starting' time for which the best correlation is achieved. This correlation is given by $\ell(k, \hat{\theta}(k))$.

Here we have exposed the dual nature of the dependence of $\ell(k, \theta)$ on its arguments. For a fixed θ , $\delta(k, \theta)$ for increasing k predicts the behavior of $\ell(k, \theta)$, as long as that θ remains in the moving window (the window $k-M < \theta < k-N$ is referenced to k , the 'present' time). So, if $\theta = \theta_t$ then $\delta^2(k, \theta)$ gives the growth pattern of the correlation as it evolves in time. However, for a fixed k , $\ell(k, \theta)$ with θ varying over the window

is a measure of that correlation for different possible 'starting' times (in a sense an autocorrelation, assuming a failure has occurred). The values of $\ell(k, \theta)$ for neighboring θ 's reflect the time structure of the failure signature. Therefore it will be different for the various failure modes. We will have more to say on this in section 2.3.4.

The non-centrality parameter for $\ell(k, \theta)$ when $\theta \neq \theta_t$ is given

[12] by

$$\delta^2(k, \theta | \theta_t) = v' c'(k, \theta | \theta_t) c^{-1}(k, \theta | \theta) c(k, \theta | \theta) v \quad (2.26)$$

with

$$c(k, \theta | \theta_t) = \sum_{j=\bar{\theta}}^k G'(j, \theta) V^{-1}(j) G(j, \theta_t) \quad (2.27)$$

and

$$\bar{\theta} = \max(\theta, \theta_t) \quad (2.28)$$

Notice that for $\theta = \theta_t$ this reduces to δ^2 as in (2.13). The expression in (2.26) is in effect an autocorrelation function for the residuals. It gives the sensitivity of the detector to the error in the estimate of θ_t . If $\delta^2(k, \theta | \theta_t)$ does not vary much for small $\tau = \theta - \theta_t$, the GLR detector is likely to choose the wrong θ as often as the correct one. We have already suggested that this might be the case for step failures in general. But if $\delta^2(k, \theta | \theta_t)$ drops sharply in value for small τ , the detector will have little trouble in consistently choosing $\hat{\theta} = \theta_t$. By selecting $\hat{\theta}$ as the θ for which $\ell(k, \theta)$ is largest, this is almost assured.

Consider the expression for $C(k, \theta | \theta_t)$ in (2.27). If we let $r=k-\theta$ and $s=k-\theta_t$, we can write

$$\begin{aligned} C(k, \theta | \theta_t) &= \sum_{l=0}^{\min(r, s)} G'(l) V^{-1} G(l+\tau) \\ &= C(r, s) \end{aligned} \tag{2.29}$$

for the time-invariant case, with $\tau=s-r$. This gives the value of the autocorrelation in the signatures for a fixed τ . When $\tau=0$ we get back the original information matrix

$$C(k, \theta) = C(r, r) \tag{2.30}$$

Thus the non-centrality parameter in (2.26) is a normalized autocorrelation function for the failure v when viewed as a function of τ (i.e., with θ as the variable since $\tau=\theta-\theta_t$).

$$\begin{aligned} \delta^2(k, \theta | \theta_t) &= v' C'(r, s) C^{-1}(r, r) C(r, s) v \\ &= \delta^2(r, s) \\ &= \delta^2(r, r+\tau) \end{aligned} \tag{2.31}$$

Going back briefly to the failure signatures with equation (2.29) in mind, we can see how the wrong-time information matrix reveals the time structure of the different failure modes. Clearly the jump failures (Figures 2.2 and 2.3) lose correlation rapidly as τ , the time shift, increases. This is more pronounced for the sensor jump, where the transients disappear very quickly. Then $\delta^2(k, \theta | \theta_t)$ drops sharply in value,

for a fixed k , as θ varies. For the step failures (Figures 2.4 and 2.5) this is no longer the case. As τ increases $C(k, \theta | \theta_t)$ in (2.29) changes little because of the larger, if not increasing, values of the signatures. Thus $\delta^2(k, \theta | \theta_t)$ does not change drastically for small $\tau = \theta - \theta_t$. Since δ^2 is a function of the failure, these results are modified depending on v . For example, the signature for sensor step failures (Figure 2.5) indicates that the autocorrelation for a failure in q is higher than for one in α . For a q failure the situation is much like that for the state steps while for α it is more like for the state jumps. We will discuss this further in the next sections and in Chapter 3.

One other consideration in deciding on the selection of the detector parameters is the desired quality of the failure estimate \hat{v} . In some cases this might not be an issue at all, but if, for example, the value of \hat{v} will be used in compensation adjustments following detection, the estimation accuracy gains in importance. In [12] Chow showed that the error covariance of the estimate $\hat{v}(k)$, which has mean value v , is given by

$$E([\hat{v}-v] [\hat{v}-v]') = C^{-1}(k, \theta) \quad (2.32)$$

under the assumption that the full GLR detector has selected the true failure type and time of failure θ .

In Figures 2.10-2.13 the elements of $C^{-1}(k, \theta)$ are plotted for $r=k-\theta$, based on the simplified F-8C model. One can see that for the jump failures, in both the state and the sensors, the best estimation accuracy




Fig. 2.10 State Jump $C^{-1}(r)$




Fig. 2.11 Sensor Jump $C^{-1}(r)$




Fig. 2.12 State Step $C^{-1}(r)$




Fig. 2.13 Sensor Step $C^{-1}(x)$

is achieved in less than five time steps. For step failures however, Figures 2.12 and 2.13 indicate that waiting up to twenty time steps can result in improved estimation. Again, this all agrees with the notion that those failure modes leading to more persistent effects on the filter residuals provide more information with time. Since the GLR detectors are computing the correlations between these residuals and their hypothesized values for each failure mode, better estimates result from having more data available. The signal-to-noise ratios are inherently much higher in some failure modes than in others. Notice, for example, that for the sensor failures (Figures 2.11 and 2.13) the element C_{22}^{-1} of the covariance matrix is the largest consistently. This means that less confidence exists on the estimate of the second component, v_2 , of the failure: the angle-of-attack component. In equation (1.52) in Chapter 1 we can see that the measurement of α has a higher noise content than that of q . This, incidentally, is reflected in the Kalman gain matrix in the filter, where the larger gains are in the first column (see (1.61)). More weight is given to the information in the q residual than to that of α . Going back to the signatures, Figures 2.3 and 2.5, the sensor failures in q lead to the largest values of the g_{ij} for both jumps and steps. On the other hand, Figures 2.2 and 2.4 indicate that α state failures have the largest signatures. Figures 2.10 and 2.12 show C_{22}^{-1} to be the smallest values accordingly.

Before leaving this section, some remarks on the invertibility of the information matrix point out some interesting connections to

other results. Given the GLR technique, for a failure to be detectable it must be able to influence the output of the system. For this reason all sensor failures are detectable, although the degree of detectability may be small for a particular failure. The noise disturbances in the system in fact impose a lower limit on the detectability of some sensor failures. The question is really one concerning the observability of the system in the absence of failures. Clearly, failures originating in an unobservable region of the state space will go undetected by a detection system that uses the output measurements, or the residuals of the Kalman filter (which is equivalent), as the information input.

In [12] it is shown that the information matrix is invertible if the system is observable and, in fact, $C(k, \theta)$ can be interpreted as an observability matrix. For sensor failure modes $C(k, \theta)$ is always invertible while for state failure $C(r)$ is invertible if the system is observable in r steps or less. If the system is time-varying, the relation to observability still holds, although with some modifications (see Chow [12]). The result on invertibility may be used in the selection of the other window parameter, N . One can choose the value of N in the window $k-M < \underline{k} < k-N$ such that $C^{-1}(N)$ exists.

Finally, it must be mentioned that one may study the detectability of different failure directions by looking at the behavior of the eigenvalues and eigenvectors of $C_i(k, \theta)$. It has already been mentioned that the elements of $C_i(k, \theta)$ give the directional sensitivity of the GLR detector for the i^{th} failure mode. The eigenvalues and eigenvectors

of this matrix also provide useful information in this area. For a failure aligned with one of its eigenvectors, the non-centrality parameter of interest is proportional to the corresponding eigenvalue while the covariance of the estimate is inversely related to it. For more on this see Chow [12].

This concludes our comments on the failure signature and information matrices. We have tried to understand some of the structure of the GLR technique in order to develop some intuition into the performance that one can expect under various failure situations. We will utilize this intuition again when we discuss the problem of cross-detection in Chapter 3.

2.3 SIMULATION RESULTS FOR CORRECT DETECTION

At this point we are in the position to study some simulations of a GLR system. The results in this section conform to the assumptions behind the correct detection formulation under matched conditions.

2.3.1 Detector Implementation

The simulations were carried out on an IBM 370 digital computer, under the two-dimensional, discrete-time model of the F-8C aircraft. The noise disturbances were taken from a random number generator with standard normal distribution [13].

Since numerous simulations were done, a specially designed Fortran program, the Multiple Detector Simulation Program (MDSP), was used for these purposes. MDSP simulates linear, time-invariant, discrete-time systems with their steady-state Kalman-Bucy filters. Up to four GLR detectors can be implemented at a time for failure modes 1-4, in any combination and window sizes. Once the system, filter and detectors are specified, any of the four kinds of failures can be introduced. The output can be selected to display different quantities of interest. These include the state, its estimate, the residuals and the output of the detectors. The detector performance is available in terms of the $\ell_i(k, \theta)$'s, the detection decisions $\hat{\theta}(k)$ and $\hat{v}(k)$. Different combinations of this information can be displayed or suppressed for any run. A description of MDSP and its user options is contained in [13].

In Figure 2.14 a flowgraph illustrates the implementation and operation of the detection scheme. Once the system data (in the form of parameter matrices, noise statistics and filter gains) are read in, MDSP proceeds to compute and store the $G_i(k, \theta)$, $C_i(k, \theta)$, $C_i^{-1}(k, \theta)$ and V^{-1} for the length of the windows to be implemented, $M-N+1$. The simulation then begins with specified initial conditions and thresholds for the detectors. At the chosen θ_t , the failure v from the selected failure mode is introduced into the system. Meanwhile, at every time k the filter residuals enter the detector algorithms. The matched filters are updated and the log-likelihood ratios are computed for each θ in

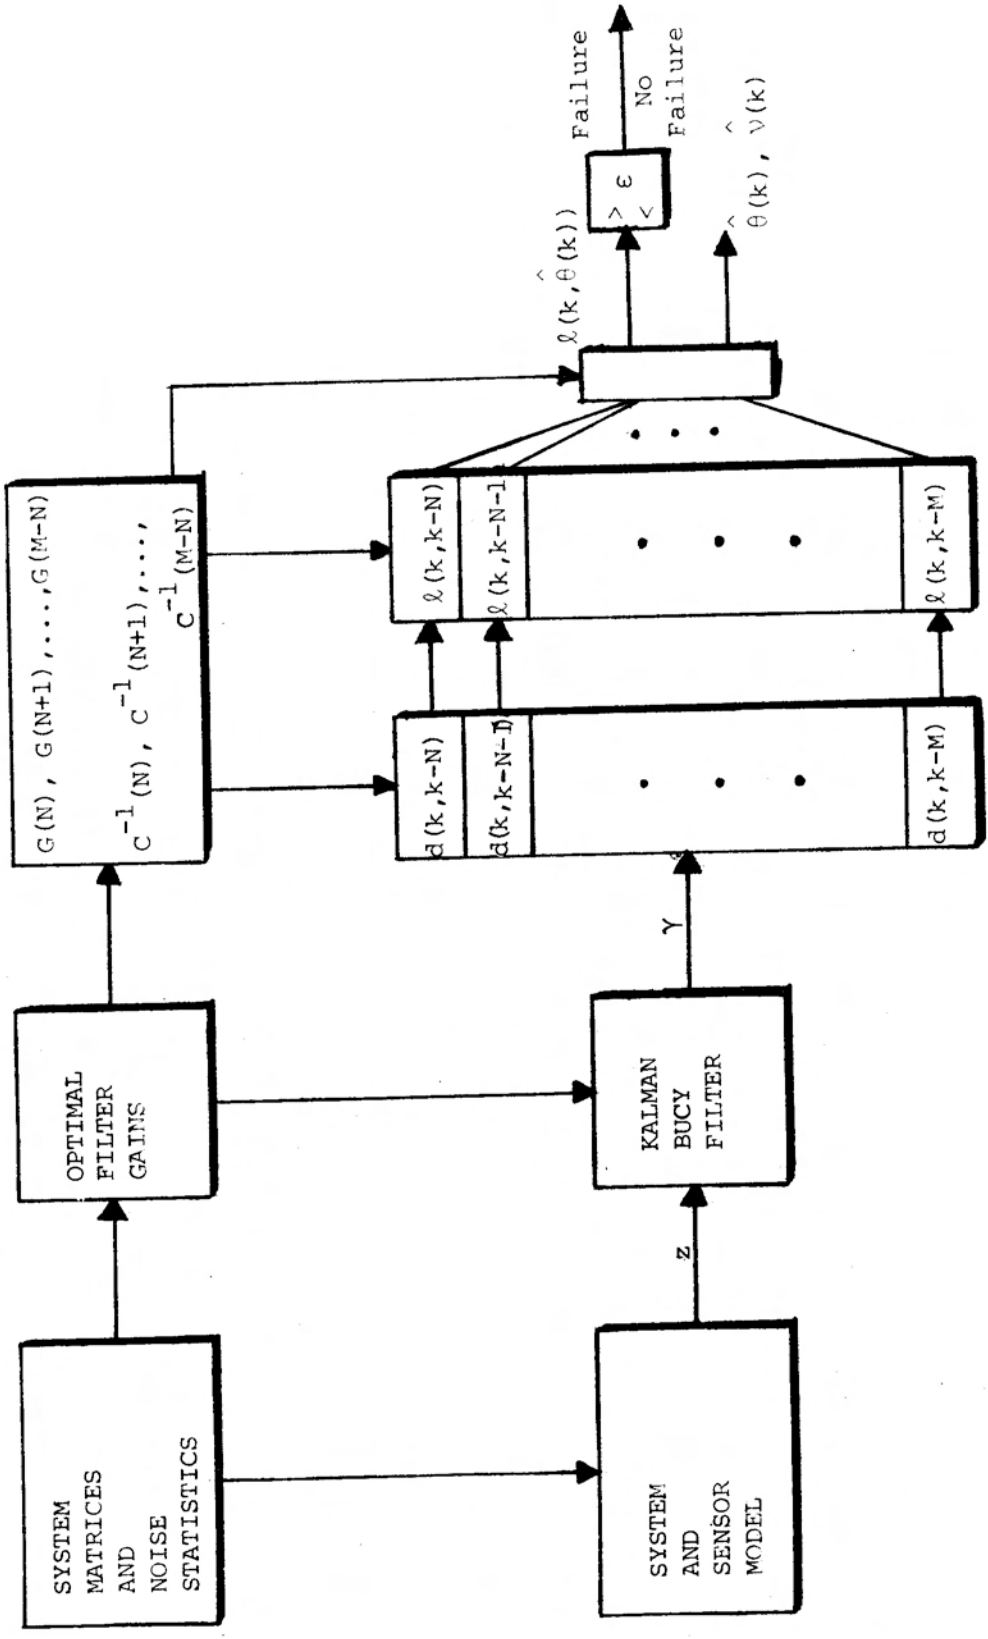


Figure 2.14 Simulation Flowgraph

the range of the moving window.

$$d_i(k, \theta) = d_i(k-1, \theta) + G_i'(k, \theta) V_i^{-1}(k) \gamma(k) \quad (2.33)$$

$$\theta \in S = \{\theta | k-M < \theta < k-N\}$$

$$\ell_i(k, \theta) = d_i'(k, \theta) C_i^{-1}(k, \theta) d_i(k, \theta) \quad (2.34)$$

Detection is declared when

$$\max_{\theta \in S} \ell_i(k, \theta) > \varepsilon \quad (2.35)$$

for each detector, in which case the MLE's of θ and v are produced

$$\hat{\theta}_i(k) = \arg(\max_{\theta \in S} \ell_i(k, \theta)) \quad (2.36)$$

$$\hat{v}_i(k) = C_i^{-1}(k, \hat{\theta}_i(k)) d_i(k, \hat{\theta}_i(k)) \quad (2.37)$$

The threshold ε was set to $\varepsilon=5$, a relatively low value, in order to study the effect of the noises on detection (i.e., false alarms, estimation accuracy, delays, etc.). In addition, this allowed the performance for higher thresholds to be inferred from the data since all the necessary information (the ℓ 's) was available. The moving window for θ was chosen as follows:

Jump Detectors (i=1,3)

$$M = 10, \quad N=0 \quad (2.38)$$

Step Detectors (i=2,4)

$$M = 30, \quad N=0 \quad (2.39)$$

The value of N is justified since this is an observable system. The parameter M was chosen with the discussion in 2.2.2 taken into consideration. That is, for jump failures a longer window in the detectors would provide no further information. Figures 2.2-2.3, 2.6-2.7 and 2.10-2.11 indicate that the short-lived effects of these failures are detectable immediately following their occurrence or not at all. On the other hand, a look at Figures 2.4-2.5, 2.8-2.9 and 2.12-2.13 shows how the prolonged presence of the step failures leads to improved detection. The detectability and accuracy in the estimates for these failures benefits with the passage of time. The value of M selected represents a compromise between these factors and the computational and storage requirements.

The failures were simulated one at a time for the range of magnitudes and directions shown in Table 1.1. The motivation and interpretation of the failure events modeled by these vectors were discussed in Chapter 1. Consideration of these failures in the different state variables and sensors enables us to appreciate the way the system dynamics and detector response are related. For each failure simulated in this section two different realizations of the noise sequences were considered. While this does not statistically validate the results, it does provide insight into GLR performance.

2.3.2 False Alarms:Detector Sensitivity

We begin the evaluation of the simulation results by looking at false alarms. This is an important issue in the design of the GLR

detectors since detection cannot be reliable if they are overly sensitive to new data. 'Detection' in the absence of failures is no better than neglecting to detect one when it is present. A study of the resulting false alarm rates for a given implementation of the GLR detectors is therefore crucial. Dealing with this problem involves one of the main tradeoffs in the design of a failure detection system. Its resolution lies in the selection of the threshold. A large value of the threshold can eliminate most or all false alarms, but at the expense of other factors in the quality of detection.

Figure 2.15 shows a plot of

$$\ell_i^*(k, \theta) = \max_{\theta} \ell_i(k, \theta) \quad (2.40)$$

for the state jump and state step ($i=1,3$) detectors for a run with no failure. At a time k , $\ell^*(k, \theta)$ is the largest log-likelihood ratio for all the θ 's included in the window. The data corresponding to the state jump detector stops at $k=20$. Two facts are noteworthy here:

- . a relatively small change in the threshold can reduce the number of false alarms very effectively
- . the jump detector is more sensitive to the noise on the average, than the step detector

The first point is encouraging since most failures of interest are likely to still be detectable after a small change in the threshold. Large values of the non-centrality parameter are not uncommon (see the δ^2 plots in section 2.2.2), thus assuring a good probability of detection.




Fig. 2.15 False Alarms: Plot of $\ell_1^*(k, \theta)$ vs. k for $\nu = 0$

With respect to the second remark, jump failures can be thought of as noise spikes or outliers outside their statistical distributions. Consequently, the jump detectors can be "fooled" by a noise sequence with large spikes. By the same token, step failures involve persistent effects and thus are less susceptible to short bursts of noise.

In Table 2.1 we compare some results from false alarm ($\nu=0$) simulations with previously computed values of the false alarm probabilities for different thresholds. Recall the definition of P_F :

$$\begin{aligned}
 P_F &= \text{Prob} \{ \ell(k, \theta) > \epsilon | H_0 \} \\
 &= \int_{\epsilon}^{\infty} p(\ell(k, \theta) | H_0) d\ell
 \end{aligned}
 \tag{2.41}$$

The numbers offered for comparisons are ratios of times when $\ell(k, \theta)$ exceeded the threshold to the total number of time steps, NTS, for each particular run. NTS equals 20 for the jump detectors and 40 for the step detectors. By ND we represent the number of times that detection was declared, i.e., times k for which $\ell_i^*(k, \theta) > \epsilon$. NDD however, is a more meaningful measure of false alarm rates than ND. It is the total number of distinct detections: detections declared as different events. This was determined by the estimate $\hat{\theta}$ of the "failure" times (different $\hat{\theta}$'s corresponding to distinct detection). So, for example, three detection decisions in a row at times $k-1, k-1, k$ declaring a failure as having occurred at some $\hat{\theta}_1$ is counted as 1 for NDD but as 3 for ND. While ND/NTS

Threshold	Computed P_F	State Jump		Sensor Jump		State Step		Sensor Step	
		ND/NTS	NDD/NTS	ND/NTS	NDD/NTS	ND/NTS	NDD/NTS	ND/NTS	NDD/NTS
$\epsilon=5$	0.082085	0.575	0.125	0.700	0.100	0.475	0.150	0.7625	0.150
$\epsilon=7$	0.030197	0.175	0.100	0.575	0.075	0.3375	0.0875	0.300	0.125
$\epsilon=10$	0.006738	0.0	0.0	0.300	0.050	0.2625	0.025	0.1125	0.0625
$\epsilon=14$	0.000912	0.0	0.0	0.175	0.025	0.1625	0.0125	0.0	0.0

Table 2.1 False Alarm Rates for Different Thresholds

says something about the accumulated effect of the noises on the sensitivity of the detectors, NDD/NTS in a sense sifts out the effect of particularly large noise values once they result in a detection decision.

Although admittedly primitive, these counts qualitatively confirm the anticipated relationship to the value of the threshold. One would need a much larger data set over which to obtain a truer image of the noise handling capabilities of the GLR detectors. As previously suggested, state jump and step detectors show a tendency to be less affected by the noise than the sensor detectors. The reason is that the system acts as a low-pass filter on the plant noise. In this sense then the measurement noise is potentially the bigger problem for detection. A different kind of failure detection criterion - one which checks for persistence in the $\ell(k, \theta)$'s - would show greatly reduced false alarm rates. One such test might, for example, count only when $\ell(k, \theta)$ remains above the threshold or grows in time. This would take advantage of the patterns in the likelihood ratios, as a function of k and θ . We will talk about this further in section 2.3.4.

2.3.3 Failure Detection

We have seen the response of the GLR detectors in the absence of failures. The next question is, how well can the GLR system spot and track a failure once it has taken place? With good performance we associate fast detection following a failure and reasonable accuracy in the estimates for the time of failure and of the failure itself. The

identification of the correct failure mode is not an issue here, but it is the topic in Chapter 3. Complementing the discussion on the detectability of different failures, we now look at the simulations with failures. One must keep in mind that the failure size is measured here in proportion to the intensity of the corresponding noise disturbance that affects the system in the same way that the failure does.

One good place at which to begin an evaluation of the detection performance of the GLR in the simulations is with the delay times. If a failure occurs at time θ_t , then the delay time, $\tau = k - \theta_t$, until the detector declares a condition of failure is a good indication of the sensitivity of the detection system. If the delays are too long the detection system may be useless. This presents the other side to the trade-off in the selection of the threshold. For a given application one must reach a decision which is a compromise between tolerable false alarm rates and delay times in detection. No general rules can be put forth since the criteria change with the system and the failures that are considered important.

We have computed the delay times in the detection of the failures simulated. Table 2.2 includes the detection delays for all the state and sensor jump failures simulated for the threshold set to the values 5, 7, 10 and 14. They are grouped by failure mode and by the direction of the failure vector: pitch rate and angle-of-attack (and their measurements). The columns in Table 2.2 show the delays for all the failures

STATE JUMP: (v,0)				SENSOR JUMP: (v,0)					
v	$\epsilon=5$	$\epsilon=7$	$\epsilon=10$	$\epsilon=14$	v	$\epsilon=5$	$\epsilon=7$	$\epsilon=10$	$\epsilon=14$
1 σ	2,4	3,4	3,5	7,9	1 σ	0,0	0,0	0,0	4, ∞
5 σ	0,0	0,0	0,0	0,1	5 σ'	0,5	0, ∞	1, ∞	∞ , ∞
10 σ	0,0	0,0	0,0	0,0	10 σ'	0,1	0,6	0, ∞	0, ∞
20 σ	0,0	0,0	0,0	0,0	20 σ'	0,0	0,0	0,0	0,0
STATE JUMP: (0 v ₂)				SENSOR JUMP: (0,v ₂)					
1 σ	0,0	0,0	15,0	∞ ,2	1 σ'	0,0	0,0	0,0	7, ∞
5 σ	0,0	0,0	0,0	0,0	5 σ'	0,0	0,0	0,0	0,0
10 σ	0,0	0,0	0,0	0,0	10 σ'	0,0	0,0	0,0	0,0
20 σ	0,0	0,0	0,0	0,0	20 σ'	0,0	0,0	0,0	0,0

Table 2.2 Delays in Detection for Different Thresholds:
 Jump Failures. Measured in time steps from θ_T .
 (One time step = 1/32 of a second).

with a fixed threshold while the rows select a particular failure for the various thresholds. The two entries at each point in the table indicate the delays for two different sample runs, i.e., with distinct noise sequences. An entry of ∞ symbolizes no detection for the length of the simulation and the windows implemented. This does not necessarily mean that the failure is undetectable, although for the jump failures it is doubtful that a longer window would improve matters much because of their transient nature.

Withholding judgment on the statistical significance of the data, we may observe trends which are not totally unexpected. If, for whatever reason is necessary, the threshold must have a higher value than initially thought, the price is paid in added delay before detection takes place. This, however, is true mainly for the smaller failure sizes. If the failure is sufficiently strong, its impact on the residuals rapidly brings about large values in the likelihood ratios, thus ensuring quick detection. The differences in delays for the various thresholds in Table 2.2 between q and α failures depend on the signatures and on the failure size. Figure 2.2 suggests that α failures (g_{i2}) should be quicker to detect than q failures (g_{i1}) because of the higher initial values. For sensor jumps (Figure 2.3) it is not so clear. However, as equation (2.17) illustrates the failure vector is equally important since it scales up or down the effect shown by the signatures. Since the failure magnitudes have been referenced to the noises, the relative values of the signatures

for q and α failures must be seen in light of the relative values of the the failures in q and α also. This, for example, accounts for the marked difference in the values of δ^2 for q and α sensor jumps in Figure 2.7. The corresponding signature in Figure 2.3 shows somewhat larger values for q failures (g_{i1}) than for α (g_{i2}). The noise intensity in α , however, is distinctively larger than that in the q sensor, as Table 1.2 points out.

This is perhaps easier to see by looking at the growth of $\ell_i^*(k, \theta)$ (see (2.40)) after a failure takes place. Figure 2.16 is a logarithmic plot of $\ell_i^*(k, \theta)$ for state and sensor jumps of 5σ and $5\sigma'$, respectively, in both q and α . Here we can observe that

- . State jumps give rise to higher values in the GLR than sensor jumps. This is in agreement with the more prolonged effect on the residuals. The signatures, we have seen, clearly illustrate this (see Figures 2.2 and 2.3).
- . While $\ell_i^*(k, \theta)$ grows with time for the state jump failures, the maximum values for the sensor jumps is reached almost immediately. Notice the similarity to the growth in δ^2 shown in Figures 2.6 and 2.7.
- . The failure in the q sensor suffers more severely than one in α in loss of detectability for an increase in the threshold. We have just discussed this (see Figure 2.7).
- . $\ell_i^*(k, \theta)$ drops in value much more sharply for the sensor failures once the time of failure θ_t is no longer a candidate for $\hat{\theta}$ in the window.

The last observation points out again how the kind of correlation present in the residuals from time to time, which differs for the various failure modes, plays an important role in determining the detector




Fig. 2.16 Evolution of $l_i^*(k, \theta)$ for Jump Failures, Average over Two Runs

response. For the sensor jump failures the signature (Fig. 2.3) shows a quick transient, well localized in time. This implies that the information in the residuals for two overlapping time intervals becomes uncorrelated for very small shifts in region of overlap. This is not the case for state jumps where the failure is integrated into the state of the system and its effects last longer. These kind of structural differences in the impact of different failure modes can and should be exploited in seeking the most reliable detection system.

Figure 2.17 is a similar plot of $\ell_1^*(k, \theta)$ for state and sensor step failures. Notice that the data for the state step failures is for a magnitude of $.5\sigma$ in q and α while for the sensor steps it is of 5σ in the measurements of q and α . The state step failures, in general, have a much stronger impact on the residuals, and consequently on the GLR, than comparable sensor failures. The reasons are the same discussed in connection to the jump failures. The steps in the state are integrated into ramp-like behavior in the output of the system (see the signatures in Figures 2.4 and 2.5).

It is worth comparing the curve for $\ell_1^*(k, \theta)$ corresponding to a sensor step in α with that for the non-centrality parameter for the same failure in Figure 2.9. This corroborates our analysis of δ^2 as a measure of the dynamic behavior of the log-likelihood ratio and the performance of the GLR system. One might also note the difference in the growth rates and actual values of $\ell_1^*(k, \theta)$ between jump (ℓ_1^* , ℓ_2^*) and step (ℓ_3^* , ℓ_4^*) failures. The sustained presence of step failures results




Fig. 2.17 Evolution of $\hat{\theta}_1^*(k, \theta)$ for Step Failures, Average over Two Runs

in an order of magnitudes difference in these values.

Figures 2.16 and 2.17 verify our analysis of the wrong-time non-centrality parameter in Section 2.2.2. Consider the behavior of $\ell_i^*(k, \theta)$ when θ_t drops from the moving window ($k=15$ for jumps and $k=35$ for steps). After this time, $\hat{\theta}$ is chosen as the most likely of the remaining θ 's. In general the next θ chosen is θ_t+1 , θ_t+2 , etc. So we are looking at the ℓ^* 's corresponding to $\delta^2(r, r+\tau)$, where both r and τ are increasing. The fact that r is also increasing in addition to $\tau=\theta-\theta_t$ accounts for the continuing increase in ℓ_i^* for some of the curves in Figure 2.17. For step failures $\delta^2(k, \theta|\theta_t)$ increases with k for all the θ 's in the window (although, for a fixed k , δ^2 decreases with τ). As discussed earlier, the sensor jumps lose correlation in the ℓ 's rapidly for the wrong θ 's. The same is true for state jumps, although it is less pronounced. Finally, notice how the ℓ 's for an α sensor step indicate wrong-time correlations not unlike the state jumps. This is further evidence of the way in which the differences among the failure modes seen in the signatures translate into distinctive patterns in the likelihood ratios. We will see this in more detail in the following section.

Finally, Figure 2.17 tells us that we can expect for step failures

- . shorter delays in detection in general
- . less sensitivity of detection to changes in the threshold

In Table 2.3 the delays in detection are presented for two sample runs, for state and sensor step failures. Notice that data is included now for

STATE STEP: (v, 0)				SENSOR STEP: (v, 0)					
v	$\epsilon=5$	$\epsilon=7$	$\epsilon=10$	$\epsilon=14$	v	$\epsilon=5$	$\epsilon=7$	$\epsilon=10$	$\epsilon=14$
1/10 σ	3,0	4,0	7,5	17,13	1/100'	0,0	3,0	4, ∞	∞, ∞
1/2 σ	2,0	3,4	3,4	3,4	1/20'	0,0	2,0	4,9	$\infty, 26$
1 σ	1,1	2,1	2,3	2,3	1 σ '	0,0	2,5	3,5	7,9
5 σ	0,0	0,0	0,0	0,1	5 σ '	0,1	1,3	1,3	2,4
10 σ	-	-	-	-	10 σ '	1,0	0,1	0,0	0,1
20 σ	-	-	-	-	20 σ '	0,0	0,0	0,0	0,0
STATE STEP: (0 v ₂)				SENSOR STEP: (0 v ₂)					
1/10 σ	13,0	14,0	14,33	14,33	1/100'	0,0	3,0	15, ∞	27, ∞
1/2 σ	5,0	5,0	6,0	6,2	1/20'	0,0	0,0	15,0	27, ∞
1 σ	1,0	2,0	2,0	3,2	1 σ '	0,0	0,0	15,0	27,2
5 σ	0,0	0,0	0,0	0,0	5 σ '	0,0	0,0	0,0	0,0
10 σ	-	-	-	-	10 σ '	0,0	0,0	0,0	0,0
20 σ	-	-	-	-	20 σ '	0,0	0,0	0,0	0,0

Table 2.3 Delays in Detection for Different Thresholds:
Step failures. Measured in time steps from θ_T .

failures of magnitudes $.5\sigma$, $.1\sigma$ and $.5\sigma'$, $.1\sigma'$. The two points mentioned above are clearly indicated by the (limited) data for step failures of 1σ (or $1\sigma'$) and larger, in comparison to the jump failures. Even for the case of most small failures here, detection is fast for high thresholds. Although somewhat sensitive to increases in ϵ , the growing nature of δ^2 for these failures means that for a larger window detection is assured after some time. In order to provide some additional insight into these results, we point out that with a threshold set at $\epsilon=7$ the longest delay in the simulations for $.1\sigma$ and $.1\sigma'$ step failures was 14 time steps (about .44 seconds). This was the case for a $.1\sigma$ state step in α , which translates into a magnitude of 4.3335×10^{-4} radians, or approximately 2.4831×10^{-2} degrees. The same failure with $\epsilon=14$ doubles its delay time, still under one second, and eliminates almost all false alarms, as seen before. Thus we see that the GLR performance under perfectly matched conditions can be extremely good. The limitations come in, of course, as we consider some of the more realistic constraints -- i.e., when the type of failure is unknown or when the system model is in error. These issues will be addressed in Chapters 3 and 4.

2.3.4 The Failure Estimates: \hat{v} and $\hat{\theta}$

The simulation results seen so far tentatively corroborate our prior qualitative analysis of the GLR performance. Another component of the performance of a failure detection system is the estimation of the failure event. For a technique such as the GLR which takes into account the dynamic propagation of the failure through the system, the

estimates of the failure and of its time of occurrence are good indicators of the quality of the detection. These estimates can provide crucial information, as we will see, in the event that complications arise in detection. Specifically, they will be useful in providing information that will allow us to determine the failure type (Chapter 3), or to identify the presence of significant modeling errors (Chapter 4). For the moment, let us take a look at the performance of the GLR detectors under matched conditions and correct detection of failure modes, by way of these estimates.

From (2.37) the estimate of the failure at time k is given by

$$\hat{v}(k) = C^{-1}(k, \hat{\theta}(k)) d(k, \hat{\theta}(k)) \quad (2.42)$$

once $\hat{\theta}(k)$ has been selected. Assuming that $\hat{\theta}$ is the correct estimate, $\hat{\theta} = \theta_t$,

$$\begin{aligned} E(\hat{v}(k)) &= C^{-1}(k, \hat{\theta}) E(d(k, \hat{\theta})) \\ &= C^{-1}(k, \hat{\theta}) E(\tilde{d}(k, \hat{\theta}) + C(k, \hat{\theta})v) \\ &= C^{-1}(k, \hat{\theta}) E(\tilde{d}(k, \hat{\theta})) + C^{-1}(k, \hat{\theta})C(k, \hat{\theta})v \\ &= 0 + Iv \\ &= v \end{aligned} \quad (2.43)$$

So \hat{v} is an unbiased estimate of the true failure if $\hat{\theta}$ is accurate. Notice that $\hat{v}(k)$ in (2.42) is the unique solution of

$$C(k, \hat{\theta})v = d(k, \hat{\theta}) \quad (2.44)$$

where v is considered unknown. The GLR detectors actually solve for the failure which best fits the mean value of the matched filter for the selected time $\hat{\theta}$.

If the estimated failure time is incorrect, $\hat{\theta} \neq \theta_t$, then

$$\begin{aligned} E(\hat{v}(k)) &= C^{-1}(k, \hat{\theta}) E(\tilde{d}(k, \hat{\theta}) + C(k, \theta_t)v) \\ &= C^{-1}(k, \hat{\theta})C(k, \theta_t)v \end{aligned} \quad (2.45)$$

The bias in \hat{v} introduced by the error in $\hat{\theta}$ can be examined for different $\hat{\theta} - \theta_t$ in much the same way as the wrong-time non-centrality parameter $\delta(k, \theta | \theta_t)$.

With these ideas in mind we now look at what the simulations show about the GLR estimates. First we consider the estimate $\hat{\theta}$ of the time of failure. The general trends in the simulations of the GLR detectors for the four failure modes are discussed, first for the state jump and state step failures.

State Failures

- . for state jump failures of size 1σ and higher, the estimate $\hat{\theta}$ is accurate in general.
- . for state step failures larger than $.5\sigma$, $\hat{\theta} \approx \theta_t$ almost immediately; for $v < .5\sigma$, $\hat{\theta}$ undergoes an initial transient but soon converges to θ_t (+ 1 or 2 time steps).
- . once θ_t is no longer in the range for θ in the window, GLR chooses the point farthest in the past that is in the window.

Looking at the growth of $\delta^2(k, \theta)$ for these failure modes in Figures 2.6 and 2.8, the initial rate of growth for the step cases is not very large. So for small failures it is possible for the noises to incorrectly influence $\hat{\theta}$ until enough time has passed and larger values in δ^2 are achieved. For the state jumps the situation is different and there is, in general, good estimation of θ_t (but for larger size failures). Figures 2.18 and 2.19 illustrate this. They are three-dimensional plots of $\ell(k, \theta)$ as a function of its two time arguments. For these the detector windows were set at $M=50$, $N=0$; the total length of the run is 60 time steps and $\theta_t = 5$. For each time k , the values in the (ℓ, θ) plane are the $\ell(k, \cdot)$ for $k-50 < \theta \leq k$. Alternatively, for a fixed θ the projection on the (ℓ, k) plane gives the growth in time of $\ell(\cdot, \theta)$, $k=0, 1, \dots, 60$. The values along the diagonal $k=\theta$ correspond to $\ell(k, k)$ and moving parallel to the θ axis toward $\theta=0$ shows the output of the detectors at a time k for the whole window. Figure 2.18 is the output of the state jump detector for 5 σ jumps in q and α . The output of the state step detector is shown in Figure 2.19 for 1 σ failures in both state variables also.

Consider first the state jump failures in Figures 2.18. We see that after the failures take place at $k=5$ the $\ell(k, \theta)$ quickly grow, reaching the final values after about 20 time steps (5/8 sec.). Notice how well the failure time is isolated, with $\ell(k, \theta | \theta_t)$ for $\theta \neq \theta_t$ having much smaller values than $\ell(k, \theta_t)$. As k increases $\ell^*(k, \theta_t)$ recedes from the diagonal




Fig. 2.18(a) Likelihood Ratios: $50 q$ State Jump Failure; State Jump Detector




Fig. 2.18(b) Likelihood Ratios: 50α State Jump Failure; State Jump Detector

$$l(k, \theta)$$

$$l_{\max} : 5,145$$




Fig. 2.19(a) Likelihood Ratios: 1σ q State Step Failure; State Step Detector

$$l(k, \theta)$$

$$l_{\max} : 5,615$$




Fig. 2.19(b) Likelihood Ratios: 1σ α State Step Failure; State Step Detector

as the correct estimate $\hat{\theta} = \theta_t$ is maintained. Notice what happens after time $k=55$. The true failure time θ_t is no longer in the window so θ_t+1 , θ_t+2 , etc. are chosen as the values of $\hat{\theta}$. The GLR detector chooses the earliest time in the window as the most likely starting time for the effects of the failure on the residuals. It is clear in Figure 2.18 that, for the jump failures, after some time the $\ell(k, \theta)$ for the remaining θ 's in the window will be smaller. Eventually the detector will not see the failure at all any longer. As the signatures show, the effect of a jump failure in the state diminishes with time.

The situation for the state step failures is different in many ways, as Figure 2.19 indicates. Here the growth of the log-likelihood ratios is slower initially, but much larger values are reached and no steady-state is seen. The distribution of $\ell(k, \cdot)$ over the window shows a very different response. While the wrong-time $\ell(k, \theta | \theta_t)$ grow in the same way as $\ell(k, \theta_t)$, for a given interval in θ the drop in the ℓ 's is not as dramatic as in the state jump detector. Although the same behavior is observed for $\hat{\theta}$ once θ_t leaves the window at time $k=55$ (hidden by the angle in which the plots are done), the increasing values of $\ell(k, \theta | \theta_t)$ with k mean that detection can occur long after, for $k \gg \theta_t$. This illustrates the central difference between the two failure modes. The continued presence of a step failure leads to a very different response pattern of the $\ell(k, \theta)$. As a function of k , the composition of $\ell(k, \theta)$ tells us about the detectability of the failures as time progresses.

As a function of θ the same response shows the kind of correlations in time which characterize the temporal evolution of a failure mode in the system. That is, the transient nature of the jump leads to highly peaked ℓ 's as a function of θ , while the persistence of a step leads to large ℓ 's over a range of θ 's.

Sensor Failures

The following summarizes general trends in the simulations corresponding to sensor jump and step failures:

- . much more sensitivity to the noises in the system than for comparable state failures
- . much lower values of $\ell^*(k, \theta)$, implying greater sensitivity also to changes in the threshold for the detection of small failures
- . for sensor jumps $\hat{\theta}$ is very erratic once θ_t leaves the moving window (as one might expect for such a transient failure)

We can discuss these features by referring to Figures 2.20 and 2.21. They show the corresponding plots of $\ell(k, \theta)$ for 5 σ ' sensor jumps and 1 σ ' sensor steps, respectively. Several things stand out immediately. First consider Figure 2.20. As a function of k , $\ell(., \theta)$ reaches its maximum value almost instantly. This is just as suggested by the non-centrality parameters in Figure 2.7. The magnitudes of these likelihood ratios are relatively low and keep the proportions seen for the δ^2 , with higher values for the 5 σ ' sensor jump in α . Notice that these plots are scaled to the maximum value of $\ell^*(k, \theta)$ (this must be kept in mind




Fig. 2.20(a) Likelihood Ratios: $50 \sigma_q$ Sensor Jump Failure; Sensor Jump Detector




Fig. 2.20(b) Likelihood Ratios: $50 \sigma_\alpha$ Sensor Jump Failure; Sensor Jump Detector




Fig. 2.21(a) Likelihood Ratios: $1\sigma'$ q Sensor Step Failure; Sensor Step Detector




Fig. 2.21(b) Likelihood Ratios: $1\sigma'$ α Sensor Step Failure; Sensor Step Detector

since, for example, the values of $\ell(k, \theta)$ for the jump in q are of the order of the noise level in $\tilde{\ell}(k, \theta)$ which are therefore scaled up in the plot). With respect to θ , $\ell(k, \theta)$ behaves similarly to the jump in the state case, although the change in value for θ 's neighboring θ_t is more pronounced. The ratio $\ell(k, \theta_t) / \ell(k, \theta | \theta_t)$ can be large even for $\theta - \theta_t$ small. This reflects the very short duration of the failure effects on the residuals, as is clear in its signature. Once θ_t leaves the window $\hat{\theta}$ is randomly directed by the noise.

In Figure 2.21 we have the detector output for $10\sigma'$ sensor failures in q and α . Again we find that the values of $\ell(k, \theta)$ are much smaller than those for state step failures. Furthermore, the $\ell(k, \theta)$ look very different for a q sensor step and one in α . All this was indicated by the non-centrality parameters and the signatures in section 2.2.2. Finally, the response to a sensor step in α is qualitatively similar to that for the state jump in q (see Figure 2.18) as was argued earlier in the Chapter. This similarity is even more striking if we consider larger sensor failures where the noise effects are small in comparison. Figure 2.22 makes this more evident. It gives the response to sensor step failures of magnitude $100\sigma'$ in both q and α . Except for the values reached by $\ell(k, \theta)$, notice the resemblance of the ℓ 's for a sensor step in q to those for state steps in Figure 2.19 and between the ℓ 's for a sensor step in α and those for state jumps in Figure 2.18.

This last fact already anticipates a source of possible trouble. If we do not assume, as we have done in this chapter, that the failure

$l(k, \theta)$

$l_{\max} : 2,636$




Fig. 2.22(a) Likelihood Ratio: $10\sigma'$ q Sensor Step Failure; Sensor Step Detector

$l(k, \theta)$

$l_{\max} : 587$




Fig. 2.22(b) Likelihood Ratio: $10\sigma'$ α Sensor Step Failure; Sensor Step Detector

mode is known, then these plots suggest that other failure modes might trigger similar responses in the detectors. It must be pointed out, however, that the shapes of correct detector ℓ 's being the same does not necessarily mean that there will be cross-detection problems. These important issues are the topic of Chapter 3.

We now come to the estimates of the actual failures, \hat{v} . We have seen that if $\hat{\theta}$ is a correct estimate, \hat{v} has as mean value v .

Let us consider first the state jump and step failures.

State Failures

The estimation performance of the GLR detectors observed in the simulations, for all the state failures tried, show:

- for jumps, the best estimates are obtained in less than 5 time steps ($\sim 1/6$ sec.). When θ_t leaves the window v degrades soon after,
- for step failures of magnitude .5 σ or less, the estimates are not very accurate and lose quality with time. For larger failures there is improvement in v , with the best estimate occurring after a wait of about 15 time steps ($\sim 1/2$ sec.).

It is not surprising that for state jump failures the estimate \hat{v} does not improve much after its initial values. The covariance of the estimate error, C^{-1} , remains near constant in time (Figure 2.10). The initial estimates are as accurate as any others which may follow. When θ_t is outside the window -- so that $\hat{\theta} \neq \theta_t$ necessarily -- the estimate

$$\hat{v}(k) = C^{-1}d(k, \hat{\theta}), \quad \hat{\theta} \neq \theta_t \quad (2.46)$$

degrades rapidly. The reason is that $d(k, \theta_t)$ and $d(k, \hat{\theta})$ in (2.46) correlate poorly except for very small $|\hat{\theta} - \theta_t|$. This is not the case for state step failures. As the 3-D plots just seen indicate, for state steps the $d(k, \theta)$ for θ 's after θ_t are highly correlated with $d(k, \theta_t)$. This is evidenced by the relatively close values of $\ell(k, \theta_t)$ and $\ell(k, \theta | \theta_t)$, as long as $|\theta - \theta_t|$ remains moderate. Furthermore, Figure 2.12 shows how much the error covariance for \hat{v} is reduced by a wait of about 15 time steps. After that there is little improvement. These results are thus consistent with our prior analysis.

Figures 2.23 and 2.24 illustrate these ideas. The first one is a plot of the estimates produced by the jump and step detectors for 10 jump and step state failures in q . The jump detector has window parameters $M=10, N=0$ while the step detector has a window $M=30, N=0$. Notice that for the jump detector the estimate \hat{v} after $k=15$, the last time that θ_t is in the window, quickly moves in the wrong directions. In addition, there is little change in \hat{v} while θ_t remains in the window. The state step estimate, however, does tend to converge slowly toward the actual failure. Convergence is faster initially, coinciding with the changing $C^{-1}(k, \theta)$ in Figure 2.12. For very large jump and step failures, the \hat{v} 's are much better in that we have the same standard deviation for the estimate errors, but these are smaller in relation to the larger sizes of the actual failures.

Sensor Failures

For the sensor jump and step failures, the simulations allow for




Fig. 2.23 Phase Plane Plot of the Estimate: 1σ State Failures




Fig. 2.24 Phase Plane Plot of the Estimate: 50 State Failures

the following observations:

- the difference in the quality of \hat{v} between small and large failures is more significant than for the state failures
- sensor jumps of up to $5\sigma'$ in size have estimates which are very sensitive to the noises
- for step sensor failures of $5\sigma'$ or larger in magnitude, the accuracy achieved can be very good. One can achieve improvement in the estimates by waiting: the best estimates are obtained after more than 10 time steps ($> 1/3$ sec).
- rapid degradation in \hat{v} occurs when θ_t drops out of the window in the jump detector. For the step detector the degradation is gradual.

At this point it should not be unexpected that for small failures in the sensors the estimates are very sensitive to the noise disturbances. The information matrices do not reach as large values as their counterparts for state failures. Therefore in

$$d(k, \theta) = \tilde{d}(k, \theta) + C(k, \theta_t) v \quad (2.47)$$

$k-M < \theta < k-N$

the first term, $\tilde{d}(k, \theta)$, can greatly affect the estimate \hat{v} in (2.42) for the failures which are not very large. Since \hat{v} is computed to solve equation (2.44), the effect of that term is negligible by comparison for large v . Another factor involved in this question has to do with the dynamic nature of the system response to the failure. Clearly, for a sensor step for which $C(k, \theta_t) v$ is increasing, one can expect improved estimates after some time. The error covariance for the estimate in




Fig. 2.25 Phase Plane Plot of the Estimate: 50 Sensor Failures

this case, shown in Figure 2.13, indicates the improved confidence in the estimate after a waiting period. In Figure 2.25 a plot of \hat{v} is given for $5\sigma'$ jump and step failures in the q sensor. Notice how the estimate for the step finally converges near v after a very erratic beginning.

2.4 SUMMARY OF CORRECT DETECTION

The GLR technique for failure detection has been examined in detail in this Chapter. Under the assumption of perfect knowledge of the model parameters and of the failure mode, the performance of the GLR detectors is quite good. We have seen that the detectors are very sensitive and, therefore, that detection is guaranteed for all meaningful failures (a sensor jump, for example, of less than $3\sigma'$ does not make much sense).

While this kind of sensitivity might suggest difficulties arising from the noise disturbances in the system, we have seen that this is not the case. The reason for this is the very concept of hypothesis testing inherent in the GLR -- i.e, the GLR detectors take full advantage of the structured and dynamic effect of the failures by correlating the observed residuals with the failure signatures.

The differences observed in the performance of the GLR detectors for the four failure modes studied here have been linked to the characteristics of the failure signatures for those modes. A detailed analysis of the signatures $G(k,\theta)$ and of the information matrix $C(k,\theta)$ associated with them has shown that a study of the average behavior of the likelihood

ratios and failure estimates $\hat{\theta}$ and $\hat{\nu}$ is feasible and reasonably accurate in predicting the performance of the detectors. The simulations presented corroborate the previous analysis based on knowledge of the signatures.

One of the key quantities in our analysis is the non-centrality parameter δ^2 , which gives the failure-induced component of the expected value of the likelihood ratios. This δ^2 captures all the information concerning the failure and its dynamic effect on the filter residuals. Its study for different failures in the four modes of interest allows us both to predict the detection performance for a given detector configuration and to systematically select those detector parameters to achieve optimum performance.

By means of $\delta^2(k, \theta)$ we can look at the patterns in the behavior of the likelihood ratios, for different failure modes, as a function of k and θ . We saw how, ultimately, all these features can be related back to the signatures of each failure mode. For example, we saw that for jump failures, especially sensor jumps, δ^2 grows toward a steady-state value (as does the probability of detection) for the correct θ , while falling sharply in value for the wrong time θ in the detector window. Figures 2.18 and 2.20 in section 2.3.4 illustrate this vividly. This reflects the transient nature of the filter response to these failures (Figures 2.2 and 2.3).

In the case of most step failures δ^2 increases indefinitely with k after the failure takes place. Furthermore, the wrong-time δ^2 also

attain large values with time, especially near θ_t . The exception is a sensor step failure in α which looks like a state jump (Figures 2.2 and 2.5), a fact which is clearly indicated by comparing Figures 2.19 and 2.22 with Figure 2.18. In contrast to the jump cases, the rest of the signatures for these step failures grow with the elapsed time thus providing more information in the residuals as time passes. Since the residuals over any time interval (after θ_t) show increasing values, especially near θ_t , the GLR detectors for these modes in general have more of a problem with wrong-time estimates $\hat{\theta}$, although detection for some θ is far easier than in the jump case. This is also shown by Figures 2.19 and 2.21(a), 2.22(a).

All these observations about the failure modes and the likelihood ratios point out the role played by the selection of the threshold. As mentioned earlier, false alarms can be effectively removed by a raised threshold. This, in turn, implies longer delays in the detection of failures since the likelihood ratios must reach the new value. The analysis and simulations show, however, that for most failures of interest the delay is minimal. In general, failures of all four modes are quickly detectable if their magnitudes are of the order of the noise intensity or greater. The most troublesome failures are sensor jumps, which are nevertheless easily detectable for larger failures (e.g., sensor jumps of magnitude $5\sigma'$ or greater).

While the probabilities used as performance measures, P_D and P_F , are valuable indicators of the detection quality of a specific GLR scheme,

they are by no means the only important ones or the definitive ones. Other probabilities can be defined on events which relate to the evolution of the likelihood ratios over intervals of time, rather than for given values of k and θ . In this sense P_D and P_F are static measures which must be computed for a sequence of times. One can, for example, define the time-to-detection probability

$$P_{TD}(T, \epsilon, \theta, \theta_t, \nu) \triangleq \text{Prob}(\ell(k, \theta) > \epsilon \text{ for some } k \leq T | \theta_t, \nu) \quad (2.48)$$

of the event that detection takes place for a given θ in T time steps or less. This quantity is clearly of importance in determining detection delay as a function of the threshold. The only drawback to this kind of analysis is that correlations between likelihood ratios are involved and for full GLR (and CGLR) it becomes difficult. The reason is the chi-squared nature of these random variables. Thus, a shift in focus to the simplified GLR (SGLR) might be rewarding in trying to obtain this kind of information. For the SGLR, recall, the likelihood ratios become Gaussian random variables. The evaluation of interval-related quantities such as (2.48) is still involved in this case but is somewhat more tractable. Further work in this area is needed. It is felt that such SGLR analysis will also shed light on full GLR and CGLR performance, since the detectors have performance characteristics that are similar to SGLR (e.g., wrong-time and detectability conditions are the same for all three).

The tools with which we have examined the performance of the GLR detectors have been adequate, nevertheless, for the task of uncovering

the basic mechanisms at work. On the one hand we have seen the interaction between failures of different modes and the responses they induce in the Kalman filter residuals. On the other, this interplay between failure modes and residuals has been related to the relative detectability of these modes via the GLR technique. These tools allow, at this stage, for a systematic consideration of the different tradeoffs which must be taken into account in implementing a reliable detection system. Once the threshold has been chosen (a compromise in delays and false alarms), the length of the windows can be arrived at (for example, by computing the waiting times necessary for a specified probability of detection: e.g., $P_D = .95$ or $P_D = .99$).

In the next chapter we relax one of the assumptions held up to now. We consider the problem of isolating correctly the mode of a failure, if and when detection has taken place. This will provide an opportunity to determine some of the limitations of the GLR technique in a more realistic environment.

CHAPTER 3

IDENTIFIABILITY OF FAILURES WITH THE GLR

3.1 MATCHED CONDITIONS AND THE CROSS-DETECTION PROBLEM

This chapter considers an important question concerning failure detection which has not been directly addressed yet: the distinguishability of the various failure modes. The discussion of the performance of the GLR technique in Chapter 2 was restricted to the situation where the failure mode is known. Consequently, the emphasis was on the capability of a GLR detector to correctly identify when a failure in a specific mode occurs and to estimate it. We will now consider the case where the mode of the failure is not known in advance. Given the decisions and output of a set of GLR detectors corresponding to different failure modes, the problem is to isolate the correct decision with the highest possible degree of certainty. The assumption is made that if a failure does take place, its mode is that of one of the detectors.

This problem may be viewed in a different way. Suppose that a GLR detector has been implemented based on the hypothesis of a particular failure mode. While in operation, detection is declared for a number of consecutive times over an interval. Complementing these decisions are sequences of estimates of a failure and its time of occurrence, $\hat{v}(k)$ and $\hat{\theta}(k)$. We can assume that the threshold has been chosen to eliminate the possibility of false alarms. The question is the following: with what degree of confidence can we accept these decisions, i.e., how certain can we be that in fact a failure of that type has occurred? The

problem is an important one, especially if compensating action is to follow such a decision. Clearly, compensation for the wrong type of failure will, at best, not correct the source of failure.

For example, in Chapter 2 it was pointed out that failures which can be modeled by a sensor step in α alter the residuals in a way not unlike a state jump failure in q . If, in response to detection declared by the sensor step detector, a new sensor is activated when the actual failure originates in the actuator affecting the pitch rate, nothing is achieved. The question involves looking at some of the fundamental limitations of the GLR method.

In many situations additional information may be available which may resolve the ambiguity about the failure mode. One may, for example, have sufficient reason to believe that the α sensor in the above hypothetical situation is functioning properly (e.g., we may have two α sensors that agree). For the rest of this chapter the assumption is made that no such extra information is available, and thus we are considering what is decidedly a worst-case situation.

Given the way the GLR technique works, one can expect that any non-random development in the innovations sequence $\gamma(k)$ will result in increased values of $\ell_i(k, \theta)$ for the different i . Depending on the selection of parameters in the detectors, this may be followed by detection and estimation of the possible failures. In Chapter 1 the decision rule

$$\max_i \ell_i(k, \hat{\theta}_i(k)) > \epsilon \quad (3.1)$$

was given, for a set of GLR detectors based on different failure mode hypotheses. This decision rule is very straightforward and appealing in its simplicity. Simulations for different failure conditions indicate, however, that it does not always choose the correct failure mode. In a number of cases the $\ell_i^*(k, \theta)$ for two different detectors are close enough in value that the random influences from the noise become decisive. Although the correct failure mode may be selected, that decision carries a degree of uncertainty which might not be tolerable for many applications.

The source of this cross-detection problem lies in the fact that failures of different modes may have signatures which are close, in some sense, to one another. An example has already been given. A decision rule such as (3.1) will work adequately in separating most jump failures from step failures. A step detector will, in general, indicate less correlation (i.e., lower values of $\ell_i(k, \theta)$) with a jump failure than a jump detector would. Conversely, a step-like failure will correlate much better with both step detectors than with any jump failure signature. But one step failure of a particular mode may correlate almost as well with signatures for either state or sensor step failures. Thus, a real distinguishability problem exists between these modes.

We saw in Chapter 2 how the patterns in the log-likelihood ratios $\ell_i(k, \theta)$ with respect to both arguments reflect the characteristics of the failure responses (the signatures) of the system to the various

modes. The analysis of the performance of the GLR system via the estimates and $\delta^2(k, \theta)$ for the correct and wrong times shows how those characteristics relate to the detectability of the failures. In section 3.2 we observe how these quantities change when the detector does not correspond to the correct failure mode. Then in section 3.3 some results from the simulations of the aircraft model are presented. A thorough study of the cross-detection problem is beyond the scope of this work, but we are able to shed light on this crucial issue and on some of the fundamental limitations of GLR. In section 3.4 is a brief discussion evaluating what other GLR formulations, the Constrained GLR (CGLR) in particular, may offer as a way out of this difficulty. Finally, section 3.5 sums up our treatment of the identifiability factor in the GLR and suggests the need for better decision rules.

3.2 SOME ANALYSIS OF THE CROSS-DETECTION PROBLEM

We begin our brief analysis of the identifiability of the different failure modes by considering the response of the incorrect detector to a failure. Consider a GLR detector based on the hypothesis

$$H_i : \gamma(k) = \tilde{\gamma}(k) + G_i(k, \theta_t) \nu \quad (3.2)$$

for a failure of type i starting at $k = \theta_t$. Suppose that the actual failure which occurs is of a different mode, j . Then the actual innovations in the filter will be

$$\gamma(k) = \tilde{\gamma}(k) + G_j(k, \theta_t) \nu \quad (3.3)$$

In order to understand what takes place in this situation, let us follow these residuals into the detectors and observe how their outputs are affected.

First, we look at the matched filters for the times in the window. The subscript i/j will denote quantities in the i^{th} detector when the residuals contain information from a failure of type j . So we have:

$$\begin{aligned}
 d_{i|j}(k, \theta | \theta_t) &= \sum_{m=\theta}^k G_i'(m, \theta) V^{-1}(m) \gamma(m) \\
 &= \sum_{m=\theta}^k G_i'(m, \theta) V^{-1}(m) [\tilde{\gamma}(m) + G_j(m, \theta_t) \nu] \\
 &= \tilde{d}_i(k, \theta) + c_{i|j}(k, \theta | \theta_t) \nu
 \end{aligned} \tag{3.4}$$

for each θ in the window, $k-M < \theta < k-N$. This has the same form as in the correct detector situation (see (2.6)) but now

$$c_{i|j}(k, \theta | \theta_t) = \sum_{m=\max(\theta, \theta_t)}^k G_i'(m, \theta) V^{-1}(m) G_j(m, \theta_t) \tag{3.5}$$

is the wrong-time, cross-detection information matrix. The mean value of each $d_{i|j}$ for all the θ 's in the window at a given time k is shaped by the values of this matrix with θ varying. For any θ

$$E(d_{i|j}(k, \theta | \theta_t)) = c_{i|j}(k, \theta | \theta_t) \nu \tag{3.6}$$

and the log-likelihood ratios become

$$\ell_{i|j}(k, \theta | \theta_t) = d_{i|j}'(k, \theta | \theta_t) C_i^{-1}(k, \theta) d_{i|j}(k, \theta | \theta_t) \quad (3.7)$$

where $C_i^{-1}(k, \theta)$ is the same as in (3.5) with $\theta = \theta_t$ and $i=j$. We will use this notation for simplicity

$$C(k, \theta) \triangleq C(k, \theta | \theta) \quad (3.8)$$

$$\delta^2(k, \theta) \triangleq \delta^2(k, \theta | \theta) \quad (3.9)$$

when $\theta = \theta_t$.

The wrong-time, cross-detection non-centrality parameter for the likelihood ratios (3.7) in the detector window, using (3.6), is then

$$\begin{aligned} \delta_{i|j}^2(k, \theta | \theta_t) &= E(d_{i|j}(k, \theta | \theta_t))' C_i^{-1}(k, \theta) E(d_{i|j}(k, \theta | \theta_t)) \\ &= v' C_{i|j}'(k, \theta | \theta_t) C_i^{-1}(k, \theta) C_{i|j}(k, \theta | \theta_t) v \end{aligned} \quad (3.10)$$

which for $i=j$ reduces to the wrong-time (correct detector) δ^2 in (2.26) and if $\theta = \theta_t$ it equals δ^2 in (2.13). The evolution of this term under different conditions gives us the sensitivity of the detectors to failures in other modes. As in the correct detection situation, such performance indices as delay times, false alarms and correct time detection are related to the behavior of $\delta_{i|j}^2$.

In a similar fashion, the estimate v of the failure under these conditions is given by

$$\begin{aligned}\hat{v}_{i|j}(k) &= C_i^{-1}(k, \hat{\theta}_{i|j}) d_{i|j}(k, \hat{\theta}_{i|j}) \\ &= C_i^{-1}(k, \hat{\theta}_{i|j}) [\tilde{d}_{i|j}(k, \hat{\theta}_{i|j}) + C_{i|j}(k, \hat{\theta}_{i|j} | \theta_t) v]\end{aligned}\quad (3.11)$$

with the expected value

$$E(\hat{v}_{i|j}(k)) = C_i^{-1}(k, \hat{\theta}_{i|j}) C_{i|j}(k, \hat{\theta}_{i|j} | \theta_t) v \quad (3.12)$$

where $\hat{\theta}_{i|j}$ is the estimate of the failure time obtained in the maximization of the likelihood ratios over θ by the i^{th} detector:

$$\hat{\theta}_{i|j}(k) = \arg \max_{\theta \in S} \ell_{i|j}(k, \theta | \theta_t) \quad , \quad S = \{\theta | k-m < \theta < k-N\} \quad (3.13)$$

All these quantities differ from their counterparts for the correct detectors only in the role which the cross-detection information matrix plays. This quantity captures the relevant information concerning the correlation which exists between the modes of the detector and failures. The nature of the response of a GLR detector to the "wrong" type of failure is largely determined by the characteristics of $C_{i|j}(k, \theta, \theta_t)$ as a function of its arguments. Its behavior as θ and θ_t vary can tell us the way the estimate $\hat{\theta}_{i|j}$ will react, if detection takes place. We saw in Chapter 2 how the behavior of δ^2 , with θ changing, reflects intrinsic properties of the mode of a failure. As a function of k , $C_{i|j}$ (and therefore $\delta_{i|j}^2$) indicates the sensitivity of the i^{th} detector to the failure as time progresses.

It is expected that if full advantage is taken of these characteristics of the cross-detection responses, modified decision rules may

be constructed which lessen the difficulties with distinguishability for the detectors. As Chapter 2 showed, detection may be very good assuming that the correct failure mode can be established (still with perfect knowledge of the parameters of the model).

The problem is illustrated in Figure 3.1. It presents the $\delta_{2|2}^2(k, \theta|\theta_t)$ and $\delta_{4|2}^2(k, \theta|\theta_t)^\dagger$ profiles across a detector window $k-40 < \theta < k$, for fixed k . A .1 σ q state step failure is assumed to have occurred at $\theta_t = k-30$. Notice that $\delta_{2|2}^2$, for the state step detector, has a clear peak at $\theta = \theta_t$ while $\delta_{4|2}^2$ does not in the sensor step detector. Also note that for $\theta \neq \theta_t$ the wrong-time, cross-detection δ^2 can actually be greater than the wrong-time correct detector δ^2 . "The physical interpretation of the situation is that a particular failure mode may not look much like another one occurring at the same time, but it may be highly correlated with the other mode started at a different time ...", [12]. As k varies, the δ^2 profiles over the window will vary for both detectors. It is this kind of information about the patterns in the ℓ 's for different detectors which can be used in modifying the GLR system to alleviate cross-detection problems. Some other examples are given in section 3.3.

$^\dagger \delta_{2|2}^2$ and $\delta_{4|2}^2$ are the non-centrality parameters of the state step (type 2) and sensor step (type 4) detectors for a state step failure, as given by the notation introduced in relation to equation (3.4)




Fig. 3.1 Wrong Time Cross Detection Noncentrality Parameter

Following the geometric interpretation of δ^2 in (2.24), $v_{1i|j}^{\prime C}(k, \theta | \theta_t) v_2$ can be thought of as an inner product between the signature of a type i failure at time θ and a type j failure at θ_t .

$$v_{1i|j}^{\prime C}(k, \theta | \theta_t) v_2 = \langle G_i(\cdot, \theta) v_1, G_j(\cdot, \theta_t) v_2 \rangle_{V^{-1}(\cdot)} \quad (3.14)$$

Thus, two failure modes having a large value in (3.14) (e.g., $i=2$ and $j=4$, state and sensor steps) are likely to have distinguishability difficulties. When $v_{1i|j}^{\prime C} v_2$ is small (for example, $i=2$, $j=3$, state step and sensor jump), we expect little or no problem with cross-detection.

Performance probabilities can be defined in much the same way as in the case of correct detection. These now involve the distribution values of $\ell_{i|j}(k, \theta | \theta_t)$ around its expected value: $\delta_{i|j}^2$, plus the dimension of the failure space. The cross-detection probability P_{CD} is the probability that $\ell_{i|j}$ exceeds the threshold for some failure of type j . P_{CD} is a special case -- where $\theta = \theta_t$ -- of the wrong-time, cross-detection probability P_{CW} . Both of these are defined as follows.

$$P_{CD}(k, i, j, \theta, \nu) \triangleq \text{Prob}(\ell_{i|j}(k, \theta) > \epsilon | H_j, \theta, \nu) \quad (3.15)$$

$$P_{CW}(k, i, j, \theta, \theta_t, \nu) \triangleq \text{Prob}(\ell_{i|j}(k, \theta) > \epsilon | H_j, \theta_t, \nu) \quad (3.16)$$

These probabilities are increasing functions of $\delta_{i|j}^2$ as given in (3.10). If P_{CD} is small (e.g. $P_{CD} \leq P_F$) the modes are easily distinguished. If $P_{CD} \geq P_D$, the two modes in question are highly correlated. Since the full GLR selects the most likely failure without restrictions, it may have difficulties with cross-detection. When the failure estimate is

constrained to a specific set of directions (CGLR) or magnitudes (SGLR), these probabilities will be restricted accordingly and the cross-detection problem may be greatly reduced. Thus, incorporating into the GLR formulation reasonable expectations about realistic, possible failures (directions and/or magnitudes) can result in improved overall performance in isolating failures. We will expand on this in section 3.4.

3.3 CROSS-DETECTION SIMULATIONS

3.3.1 The Likelihood Ratios

In order to gain some concrete understanding of the difficulties of the GLR with failure mode distinguishability, we have made a number of cross-detection simulation runs. In addition to the correct GLR detector, other detectors based on different modes were implemented simultaneously. As the failures in different modes were simulated, the output and decisions of all detectors were recorded. Particular emphasis was placed on the distinguishability between state and sensor step failures. These are failure modes whose signatures are, in most cases, highly correlated (and which are modes that model failures of great practical interest).

If the correlation between the signatures of two failure modes is not too pronounced, simple decision rules, such as (3.1), which compare the ℓ 's of the detectors can effectively isolate the correct failure mode. When the cross-detection difficulties are more severe, however, more sophisticated decision rules must be used. These must utilize

the information about the patterns in the ℓ 's of each detector. The availability of the values of the ℓ 's for all the θ 's in the window suggests that window decision rules should be used. A detection decision would then be made conditional on $\ell(k, \theta)$ having a certain shape across the window $(k-M < \theta < k+N)$ for any k . For example, a window decision can declare a failure at time k if K of the ℓ 's in the window exceed the threshold at that time. Then $\hat{\theta}$ can be chosen as that value of θ in the window with the largest ℓ . Other rules can be similarly designed.

Another reason for looking at window decision rules is to take full advantage of the structured behavior of the ℓ 's as a function of θ . It is clear from Chapter two that $\ell(k, \theta)$ -- for fixed k and varying θ -- displays the differences between the signatures of the different modes.

Decision rules that look at $\ell(k, \theta)$ as a function of k are also possible. These rules can check for the growth in time of the likelihood ratios, which differs for the various failure modes. They can be designed to search for continued increase in the ℓ 's -- as for most step failures -- or, for steady states -- as for jump failures (or α sensor steps). The price, however, is that a certain waiting time is necessary before a decision can be made, since the ℓ 's must be monitored over some specified interval of time. Also important is the θ (or θ 's) in the windows for which this test applies. We have seen (in Chapter 2) how different the growth in ℓ can be, for some types of failures, for

two θ 's wide apart in the window.

Let us look at the cross-detection problem between state step and sensor step failures. The generally persistent effects of these modes (Figures 2.4 and 2.5) on the residuals leads to indistinguishability difficulties. It is possible for both kinds of failures to elicit a similar response in the detectors for either mode. Figure 3.2, (a) and (b), are plots of $\ell_{2|4}(k, \theta)$ for simulation runs of 100' sensor step failures in q and α ($\theta_t=5$). Compare these with the response of this same state step detector to state step failures in q and α shown in Figure 2.19, (a) and (b). Notice the similarity in the patterns of the likelihood ratios in Figures 2.19 (a) and 3.2 (a), and their differences in Figures 2.19 (b) and 3.2 (b).

A decision rule for the state step detector which declares a failure only if the ℓ 's across the window at time k have the shape of the correct detector response (Figure 2.19, (a) and (b)), would work as follows if excited by a sensor step. If a sensor step in α (Figure 3.2 (b)) occurs, it would not satisfy the test and the state step detector would not declare a failure. The failure mode is easily distinguished in this case. However, if the failure is a q sensor step, Figure 3.2 (a) indicates that it would be accepted by the state step detector and detection declared. Here the indistinguishability is more severe. Yet, the response of the correct detector (sensor step) to this same failure -- shown in Figure 2.22 (a) -- shows larger values for ℓ




Fig. 3.2(a) Likelihood Ratios: 100% q Sensor Step Failure; State Step Detector




Fig. 3.2(b) Likelihood Ratios: 100% α Sensor Step Failure; State Step Detector

than the incorrect detector does. Hence, a comparison of the ℓ 's for the two detectors can resolve the problem.

Unfortunately, it is not always as simple as indicated. Figure 3.3, (a) and (b), are plots of $\ell_{2|4}(k, \theta)^\dagger$ when the failures are 100' sensor steps in q and α ($\theta_t=5$). The situation remains the same for the α failure: the wrong failure does not meet the requirement of the concave shape of the ℓ 's across the window. But things are different for the q sensor step failure. Figure 2.21 (a) has the response to this failure of the (correct) sensor step detector. A simple comparison of the ℓ 's in both detectors will not necessarily select the correct failure mode. The reason is that the difference between the non-centrality parameters of the correct and incorrect detectors, for a given failure, is proportional to the square of v (see equations (3.10) and (2.26)). Thus, this difference is larger in the case of the 100' sensor step failure.

All this suggests is that information about the structure of the likelihood ratios cannot be used to resolve cross detection problems for all cases (perhaps their behavior with increasing k is also necessary). We have seen that there are different degrees of indistinguishability. In particular, separating state steps from q sensor steps requires even more than testing the shape of the ℓ 's across the window. Other infor-

[†]That is, the ℓ of the state step detector in response to a sensor step failure.




Fig. 3.3(a) Likelihood Ratios: $1\sigma'$ q Sensor Step Failure; State Step Detector




Fig. 3.3(b) Likelihood Ratios: $1\sigma'$ α Sensor Step Failure; State Step Detector

mation must be used.

For example, compare the values of the ℓ 's achieved in Figures 3.2 and 2.19. A state step failure with magnitude of the order of the noise in the dynamics actually results in much larger values than a sensor failure ten times the measurement noise.

Thus, using a priori information as to what are physically meaningful failure sizes, we may set the threshold for the state step detector at a very high level -- allowing for detection of reasonable size state step failures, but avoiding cross-detection of all except very large q-sensor step failures. The point here is that additional information is available which can be used to distinguish modes. In the next section we look at another fact which helps alleviate the cross-detection difficulties.

This discussion can be illustrated with the aid of Table 3.1. It summarizes the shapes of the likelihood ratio profiles over the window of the state and sensor step detectors (i.e., $\ell(k, \theta)$ versus θ for fixed k), for the correct and cross-detection combinations. We have characterized these, for simplicity, as either concave (as in Figure 3.2 (a)) or convex (as in Figure 3.2 (b)). The shape of the ℓ 's in the correct detector/failure combinations are underlined (entries #1, 2, 7 and 8). A decision rule which tests the ℓ 's for consistency with the correct patterns would easily distinguish a failure with a different shape (e.g., between entries #2 and 4 or between #6 and 8). The real distinguishability problem is then reduced to that between a q sensor step

and a q state step (entries # 1 and 3 or 5 and 7).

FAILURE	STATE STEP DETECTOR	SENSOR STEP DETECTOR
state step in q	<u>concave</u> 1	concave 5
state step in α	<u>concave</u> ** 2	concave ** 6
sensor step in q	concave 3	<u>concave</u> 7
sensor step in α	convex ** 4	<u>convex</u> ** 8

TABLE 3.1 Likelihood Ratio Profiles Over Window

The situation for sensor step detectors is illustrated in Figure 3.4 (a) and (b). They are plots of $\ell_{4|2}(k, \theta | \theta_t = 5)$ when the failures are q, (a), and α , (b), state steps. These correspond to entries no. 5 and 6 in Table 3.1. The response of this detector to sensor step failures can be seen in Figure 2.22 (corresponding to entries no. 7 and 8 in Table 3.1). A similar decision rule would have no difficulty distinguishing a sensor step in α from the other failures. Again, it is a sensor step in q which gives rise to indistinguishability.

We conclude that the indistinguishability between failure modes can be reduced in many cases by exploiting the distinctive patterns in the likelihood ratios. Thus the cross-detection problem is narrowed down to the more fundamental causes of difficulty. Other information in the GLR still remains unexamined which can further contribute to

$$l(k, \theta)$$

$$l_{\max} : 5,150$$




Fig. 3.4(a) Likelihood Ratios: 10 q State Step Failure; Sensor Step Detector

$$l(k, \theta)$$

$$l_{\max} : 5,691$$





Fig. 3.4(b) Likelihood Ratios: 10 q State Step Failure; Sensor Step Detector

alleviating these limitations. In the next section we look at the failure estimates.


3.3.2 The Failure Estimates

In addition to the likelihood ratios, the estimates of the failure have something important to say about cross-detection. We have seen that on the basis of the ℓ 's -- using all the available information in the detectors -- the distinguishability between different failure modes can be enhanced. Since the estimates of the failure (in the event of detection) are readily accessible, they offer a convenient source of information. The quality and behavior of $\hat{v}(k)$ was seen in Chapter 2 to reflect some of the characteristics of the signatures of the different modes. In the case of an incorrect detector one might also expect to see the estimates provide information indicative of the mode of a failure.

In the simulations of the correct and incorrect detectors, the failure estimates were recorded whenever detection took place. Figures 3.5 and 3.6 show the behavior of the failure estimates for the state and sensor step detectors when 1σ state step failures were simulated in q and α . The correct estimate $\hat{v}_{2|2}(k)$ slowly converges to the vicinity of the true failure in the phase plane. The incorrect estimate $\hat{v}_{4|2}(k)$, however, behaves quite differently. First of all, it converges to a point which indicates contributions to the failure from from q and α while θ_t remains in the window. A decision rule which



STATE STEP DETECTOR



SENSOR STEP DETECTOR

Fig. 3.5 Failure Estimates for 10σ State Step





Fig. 3.6 Failure Estimates for 10 α State Step

only accepts failures that are physically meaningful -- e.g., failures only in q or in α -- by monitoring the convergence of \hat{v} , can be a substantial aid in identifying the correct failure mode. Notice that what such a rule amounts to is CGLR.

Also of interest is the reaction of \hat{v} when θ_t is no longer in the window. In these simulations ($M=50$, $N=0$, $\theta_t=5$) this occurs at $k=55$. Notice that while $\hat{v}_{2|2}(k)$ changes slowly after θ_t is out of the window, $\hat{v}_{4|2}(k)$ goes through a sharp reversal in direction as soon as that time is reached. Also, the correct estimates tend to vary less before the end of the window is reached. This information about the incorrect estimate can be used if one is willing to wait the time necessary for θ_t to leave the window ($M-N+1$ time steps) before making a decision based on a particular $\ell(k, \theta)$ for a fixed θ and varying k .

The kind of behavior shown here by the correct and incorrect failure estimates shows clearly that this information can further contribute towards identifying the true failure mode. Together with modifications which make use of the patterns in the ℓ 's, the cross-detection problem can be systematically studied, ameliorated, and possibly, eliminated.

3.4 DISTINGUISHABILITY WITH CGLR

The observation made in the last section, concerning the behavior of the incorrect failure estimate, merits a closer look. In Figures 3.5 and 3.6 it was seen how for the full GLR the incorrect estimates

$\hat{v}_{i|j}$ differ from the correct detector estimates (for a sensor step detector and state step failure). The failures simulated were orthogonal: failure in either q or α . Yet, the incorrect estimates approached -- while θ_t remained in the window -- an estimate indicating substantial components in both q and α . The full GLR detector chooses the estimate \hat{v} which is most likely to have resulted in the maximum (over the window) likelihood ratio, assuming the particular failure mode hypothesized by the detector in fact occurred. It is free to choose \hat{v} as any point in the failure space of appropriate dimension. This raises the following question: how much better would the constrained GLR (CGLR) or the simplified GLR (SGLR) perform with respect to distinguishability problems? The assumption that these formulations will, in fact, do at least as well as the full GLR is discussed below. The comments in this section refer to the CGLR and it is expected that many of them carry over to SGLR as well.

The CGLR (and SGLR) formulation was discussed briefly in Chapter

1. We saw that it consists of computing, for each failure mode, the likelihood ratios and failure estimates along specific directions: $\hat{v} = \hat{\beta} \hat{f}_m$, where \hat{f}_m is one out of a set of possible failure directions and $\hat{\beta}$ is an estimate of the magnitude of the failure along \hat{f}_m (i.e., a scaling factor). We note that when a failure occurs in the assumed direction, \hat{f}_m , the δ^2 in the CGLR corresponding to that direction should be the same as for the full GLR (In CGLR a $\delta^2(m)$ is associated with

every direction f_m). This is true assuming that the correct mode is known. When $v = \beta f_m$ in the full GLR, the expressions for δ^2 and $E(\hat{v})$ coincide with those for CGLR.

In most applications the physically meaningful failures are not likely to be "free" -- i.e., with v totally unknown. For example, a failure in a particular sensor may be independent of the other measurements. Restricting the search for failures to a set of specific directions should reduce the uncertainty in identifying the actual failures. In a cross detection situation between modes with considerable distinguishability problems, the potential gains in using CGLR are very appealing. Suppose, in reference to the results in Figures 3.5 and 3.6, that the incorrect detector had been constrained to pure q and pure α failures. In all probability, then the incorrect $\ell_{i|j}$'s along those directions would be appreciably smaller than the ℓ 's for the correct detector. The reason for this is the following. For the full GLR the incorrect detector estimate $\hat{v}_{i|j}$ -- the most likely type i failure which accounts for the $\ell_{i|j}$ achieved -- required a mixed q and α failure. Constrained to orthogonal failure directions, the likelihood ratios would have to be lower in value, and hence they would be smaller than the ℓ 's in the correct detector, which should not change too much (as they tend toward predominately q or α failure estimates).

Thus, a comparison between the capabilities of full GLR and CGLR should be investigated in order to gain a better understanding of the tradeoffs vis a vis cross-detection. Since the dominant failure modes

and directions in any application are given to us, the design limitations of the GLR are very much context-dependent. Hence, further work on this question should concentrate on developing a systematic treatment of these limitations and tradeoffs between full GLR and CGLR (and, by extension, SGLR).

The likelihood ratios and failure estimates in CGLR are computed only for failures in certain directions which are specified beforehand (see Chapter 1, section 1.2). Thus, for every failure mode the corresponding CGLR detector searches along these directions and selects the one in which a failure is most likely to have occurred. A scheme in which the outputs of all detectors are compared for each failure direction can then be studied by means of analysis and simulation. For example, one can evaluate δ^2 for the state step and sensor step detectors in the q and α directions in CGLR and compare them to δ^2 in the full GLR for those same detectors. The wrong-time, cross-detection version of these quantities then tell us what can be gained by using CGLR as far as failure mode distinguishability is concerned.

The general expression for the wrong-time, correct detector δ^2 and failure estimate in CGLR are

$$\ell(k, \theta | \theta_t, m) = \frac{b^2(k, \theta | \theta_t, m)}{a(k, \theta | m)} \quad (3.17)$$

$$\hat{\beta}(k) = \frac{b(k, \hat{\theta} | \theta_t, m)}{a(k, \theta | m)} \quad (3.18)$$

where

$$b(k, \theta | \theta_t, m) = f_m' d(k, \theta | \theta_t) \quad (3.19)$$

$$a(k, \theta | m) = f_m' C(k, \theta) f_m \quad (3.20)$$

and f_m is the m^{th} failure direction (i.e., $v = \beta f_m$). The actual failure leads to a residual sequence

$$\gamma(k) = \tilde{\gamma}(k) + \beta G(k, \theta_t) f_m \quad (3.21)$$

Corresponding to (3.17) and (3.18), we have

$$\begin{aligned} \delta^2(m) &= \frac{\beta^2 (f_m' C(k, \theta | \theta_t) f_m)^2}{f_m' C(k, \theta) f_m} \\ &= \frac{\beta^2 a^2(k, \theta | \theta_t)}{a(k, \theta)} \end{aligned} \quad (3.22)$$

and

$$E(\hat{\beta}) = \frac{\beta a(k, \hat{\theta} | \theta_t)}{a(k, \hat{\theta})} \quad (3.23)$$

In a cross-detection situation, for a fixed direction, the expressions for these quantities become

$$\delta_{i|j}^2(m) = \frac{\beta^2 (f_{m_i|j}' C_{i|j}(k, \theta | \theta_t) f_m)^2}{f_{m_i}' C_i(k, \theta) f_m} \quad (3.24)$$

and

$$E(\hat{\beta}) = \frac{\beta \mathbf{f}'_{m i} \mathbf{C}_i |_{j} (k, \theta | \theta_t) \mathbf{f}_m}{\mathbf{f}'_{m i} \mathbf{C}_i (k, \theta) \mathbf{f}_m} \quad (3.25)$$

Comparing these quantities with those given by (3.17) and (3.18) for \mathbf{f}_m in the q and α directions will tell us how well the CGLR handles distinguishability between failure modes with correlated signatures. Comparing them to the corresponding ones in full GLR then indicates the relative advantages in cross-detection between the two GLR formulations. A similar kind of analysis can be performed with respect to simplified GLR to see how further restriction to include the magnitude of v affects distinguishability. Finally, we note that (3.24) and (3.25) can be expanded to include the case where, in addition, the failure is in a direction \mathbf{f}_n other than the one hypothesized in the detector.

3.5 CROSS-DETECTION: DISCUSSION

This concludes our brief look into the difficulties with the GLR technique in distinguishing between several failure modes. We have seen some of the fundamental limitations of this approach. The nature of this indistinguishability -- similar behavior of the residuals under different failure conditions -- is built into the GLR technique since it searches for specific developments in the residuals.

There are several ways to deal with these difficulties in determining which is the true mode of a failure. The main consideration common to all is the realization of the fact that the principal design

limitations with the GLR are very much context-dependent. Therefore, efforts to improve detection must take advantage of the way in which these dependencies are displayed in the various quantities computed in the detectors. In one approach, full use should be made of all available information concerning the ℓ 's and the estimates. The patterned behavior of these quantities for the correct detector can be used effectively to provide increased distinguishability. The design of "smart" decision rules, which check for this distinctive kind of behavior consistent with the hypothesized failure mode, can eliminate or restrict many of the limitations of this method.

Another approach, complementary to the use of modified decision rules, is the search for the most applicable GLR formulation for the problem at hand. The flexibility offered by full, constrained and simplified GLR allows us to provide a better match to our knowledge about the system and failures. The assumptions about the failures which can be made on physical considerations motivate the various restrictions on ν made by SGLR, CGLR and full GLR. If only a particular failure can occur, or if its origin must necessarily be related to specific points in the systems, such information should be acknowledged explicitly in selecting the detector formulation. Further work with the GLR technique should concentrate on development of a methodology with which to select the best design, for a given situation, in a systematic fashion.

CHAPTER 4

SENSITIVITY TO MODELLING ERRORS

4.1 DETECTION UNDER MISMATCHED CONDITIONS

In this chapter we consider the sensitivity of the performance of the GLR detectors to modeling errors. Up to now we have assumed that the model of the system used to compute the Kalman filter gains and detector matrices is exact: all parameters are perfectly known. In previous chapters we have examined the detection performance of the GLR under this assumption.

The concern over the issue of modeling errors is motivated by two important considerations:

- . one can never really measure the model parameters exactly, and, in fact, the model is often a vast mathematical simplification of the physical system
- . even if one could, the true system parameters are likely to have some time-varying behavior

With respect to the first of these, one would like to know the relationship between quality of detection and the accuracy to which the system parameters are known. Perhaps of more importance, however, is the second observation. In many applications the true system drifts slowly away from the dynamical system specified by the model. Or, the true system is only approximated by the model for certain regions in the state space. For example, the F-8C model used here consists of a sequence of linear, time-invariant systems corresponding to different flight conditions along the flight trajectory. A more exact representation of the aircraft would be a nonlinear dynamic system [16].

In order to develop some intuition into the possible implications of modeling errors for reliability in detection via the GLR, some analysis and simulations were carried out. These results are only preliminary, and much more work remains to be done. However, we feel that some of the basic issues have been identified and that it is possible to deal with the problem of parameter errors in a number of ways.

Finally, we view the subject of this chapter as a further look at the basic limitations of the GLR technique. Some of the difficulties in distinguishing between failure modes were demonstrated in Chapter 3. In fact, it will be seen that an analogous situation exists between certain other failure modes -- hard-over sensor failures in particular -- and the kind of modeling errors considered here. Therefore, on the one hand we learn about the effect of these errors on performance and on the other we can view this as providing insight into failures (Chapter 1, section 1.2) using detectors based on the simpler models -- such as the step detectors.

Reference to the term 'mismatch' in this chapter denotes the situation where some or all of the model parameters on which the detectors and the filter are based are in error.

4.2 THE FILTER MATCHED TO THE DETECTOR

4.2.1 Description of Mismatch

In approaching the sensitivity of GLR performance to modeling errors, an analytical treatment becomes intractable very quickly. We

therefore rely substantially on the use of simulations. The simulation program MDSP ([13] and Chapter 2, section 2.3.1) easily accomodates various forms of mismatches. The detectors and filter can each be computed based on the same model of the system as the true one, or, on another one which can be specified. For our purposes reduced-order models for the F-8C aircraft were used, for flight conditions 10 and 12 (Chapter 1, section 1.2.1 and [16]). These correspond to flight at Mach 0.4 and Mach 0.8, respectively, with the same altitude and weather conditions as flight condition 11 (20,000 feet and cumulus clouds).

The most significant differences between these three models are:

- the $H_{\alpha\alpha}$ term changes approximately by a factor of 2 from each flight condition to the next (due to marked changes in dynamic pressure)
- the Φ matrix changes so that the period of the aircraft oscillations increases by almost 50% from condition 12 to condition 11, and again from condition 11 to condition 10.

We note that these are fairly large modeling errors and therefore we are considering extreme, but meaningful, situations. Table 4.1 summarizes the parameters of the model and Kalman filter corresponding to flight condition 12 (the parameters for flight condition 11 are given in section 1.2.2). All the results reported in this chapter involve a mismatch between flight conditions 11 and 12.

The period of oscillation of these two models are approximately 92 (3 sec.) and 67 (2 sec) time steps respectively. The $H_{\alpha\alpha}$ term, which doubles from flight condition 11 to condition 12, depends on the dynamic pressure and thus the velocity. We use flight condition 12

$$\Phi = \begin{bmatrix} 0.97501 & -0.273 \\ 0.0303 & 0.95852 \end{bmatrix}$$

$$G = \begin{bmatrix} 0.027354 & 0.0 \\ 0.0041515 & 0.0002776 \end{bmatrix}$$

$$H = \begin{bmatrix} 1.0 & 0.0 \\ 0.0 & 31.276 \end{bmatrix}$$

$$D = \begin{bmatrix} 0.00873 & 0.0 \\ 0.0 & 0.06 \end{bmatrix}$$

$$\lambda_i(\Phi) = \begin{cases} 0.966765 + j(0.090575) & i=1 \\ 0.966765 - j(0.090575) & i=2 \end{cases}$$

$$K = \begin{bmatrix} 6.6267 \times 10^{-1} & 5.7414 \times 10^{-2} \\ 8.6715 \times 10^{-2} & 1.2029 \times 10^{-2} \end{bmatrix}$$

$$P = \begin{bmatrix} 7.9284 \times 10^{-4} & 1.2081 \times 10^{-4} \\ 1.2081 \times 10^{-4} & 1.9015 \times 10^{-5} \end{bmatrix}$$

$$\Theta = \Phi(I - \bar{K}H) = \begin{bmatrix} 0.305227 & 1.580514 \\ 0.093339 & 0.652316 \end{bmatrix}$$

$$\lambda_i(\Theta) = \begin{cases} 0.057296 & i=1 \\ 0.900246 & i=2 \end{cases}$$

Table 4.1 System and Filter Parameters for Aircraft Model at Flight Condition 12

rather than 10 for the real system because 12 has a shorter period. Hence, the effects that depend on this are likely to happen faster.

4.2.2 Complete Mismatch

We first look at the case where the detector and the filter are calculated based on the same flight condition which differs from the true system (on which the measurements are made). This is a realistic situation since the filter gains and detector matrices must be either generated on-line or kept in storage. However, it is an idealization of the actual situation since the true system is likely to be gradually changing. It is impossible to obtain them for every change in the actual system. So we only assume that the system model is "near enough" to the real system to be meaningful. The case where the filter gains, but not the detectors, can be matched to the true system is discussed briefly in section 4.4.

In order to understand how detection performance is affected by the modeling errors we look at the way γ , ℓ , $\hat{\theta}$ and \hat{v} change. We can then infer the probable rates of false alarms, detection delays, etc. We begin with the behavior of the residuals under mismatch. Since the Kalman filter is not matched to the system, these innovations need no longer be a white noise process in the absence of failures (although whiteness in the residuals is not a sufficient condition to declare a model exact, as is pointed out in [17]).

Figure 4.1 gives plots of the residuals for simulation runs where the true system is at flight condition 12 while the filter and




Fig. 4.1(a) Residuals Under Mismatch:
No Failure; $q(0)=0, \alpha(0)=5^\circ$




Fig. 4.1(b) Mean Value of Residuals in
Fig. 4.1(a)

detector are computed for flight condition 11. Figure 4.1 (a) shows the actual residuals generated, and in 4.1 (b) the same residuals are shown for a deterministic run (all noises equal zero for all k). Both are the residuals of the unfailed system. The latter indicates how the mean value of the residuals responds to the mismatch. Both of these correspond to the system started at an initial condition with $q(0)=0$ and $\alpha(0)=5^\circ$, the reason for which will be discussed shortly. By deleting the noises we have isolated the pure effect of the mismatch on the residuals. The relationship between this oscillatory behavior in γ and the system dynamics is explored later on in this chapter.

This kind of behavior in the residuals suggests how the GLR detectors will be sensitive to modeling errors. To the extent that the residuals over some interval of time resemble the failure signatures for some mode, false alarms will increase. The likelihood ratios in the state step and sensor step detectors for these residuals are plotted in Figures 4.2 (a) and 4.2 (b). Notice, first of all, the large values of the ℓ 's in both detectors. Thus, false alarms will surely occur (assuming the use of an instantaneous decision rule). Furthermore, notice the peculiar nature of the patterns in the ℓ 's, involving oscillatory behavior in both k and θ . The latter fact should not surprise us, given the oscillations in the residuals and that $\ell(k,\theta)$ is a quadratic function of $\gamma(k)$.

Consider the ℓ 's in Figure 4.2(a) as they develop with increasing k . At first only the first peak is in the window and a $\hat{\theta}$ will correspond




Fig. 4.2(a) Likelihood Ratios Under Mismatch:
No Failure; State Step Detector




Fig. 4.2(b) Likelihood Ratios Under Mismatch:
No Failure; Sensor Step Detector

to its largest value. Later, as the first peak dies out and a new one appears for later θ 's, $\hat{\theta}$ undergoes an abrupt shift. As we know, \hat{v} also changes sharply for such behavior in $\hat{\theta}$. The correlation between the ℓ 's for θ 's wide apart can be quite small. In Figure 4.3(a) we have plotted $\hat{\theta}(k)$ corresponding to the likelihood ratios, shown in Figure 4.2(a), for the state step detector under mismatch. The estimate $\hat{\theta}$ for the sensor step detector, under similar conditions, is shown in Fig. 4.3(b). In the first case, the estimate of the declared failure undergoes a large discontinuity. This increase equals the distance in θ between the two peaks observed in the ℓ 's in Fig. 4.2(a). The estimate $\hat{\theta}$ in the sensor step detector behaves quite differently, as do the ℓ 's shown in Fig. 4.2(b).

Two comments can be made based on these observations. In the case of the state step detector, Figures 4.2(a) and 4.3(a) suggest that sudden, large shifts in $\hat{\theta}$, corresponding to times within the window, are indicative of a mismatch situation. Furthermore, the size of the shift $\Delta\hat{\theta}$ is a function of the period in the ℓ 's and thus, is directly related to the oscillations in the residuals. The importance of this fact should become clear in the discussion which follows. On the other hand, in the case of the sensor step detector (Figures 4.2(b) and 4.3(b)) $\hat{\theta}(k)$ is indistinguishable from the estimated time of an actual failure, as it consistently selects the same value for $\hat{\theta}(k)$. Hence, other information must be used in this case to distinguish between a real failure and simply a condition of mismatch. This information




Fig. 4.3(a) Estimated Failure Time Under Mismatch: No Failure; State Step Detector




Fig. 4.3(b) Estimated Failure Time Under Mismatch: No Failure; Sensor Step Detector

must come from the failure estimates $\hat{v}(k)$ and/or from the patterns of the likelihood ratios (from this detector and from the state step detector).

The distinctive shapes of the ℓ 's in Fig. 4.2 indicates that the kind of decision rules suggested in the last chapter (in relation to cross-detection) might prove useful here also. For example, under mismatch the ℓ 's at times exhibit more than one peak across the window simultaneously. This was never seen before in detection of failures of any kind under matched conditions. Similar comments apply to decision rules which look at the behavior of the ℓ 's with time.

Let us consider the nature of the changes in the residuals in response to the mismatch. Suppose that the real system can be described by

$$x(k+1) = \Phi_1 x(k) + \Gamma_1 w(k) \quad (4.1)$$

$$z(k) = H_1 x(k) + v(k) \quad (4.2)$$

where $w(k)$ and $v(k)$ are the same Gaussian white noise sequences given in Chapter 1.

If the Kalman filter is based on a system model with Φ , H and Γ instead of Φ_1, H_1 and Γ_1 , the estimates in the filter are given by

$$\begin{aligned} \hat{x}(k+1) &= \Phi(I-KH) \hat{x}(k) + \Phi K z(k) \\ &= \Phi(I-KH) \hat{x}(k) + \Phi K H_1 \hat{x}(k) + \Phi K v(k) \\ &= \Phi(I-KH) \hat{x}(k) + \Phi K H x(k) + \Phi K \Delta H x(k) + \Phi K v(k) \end{aligned} \quad (4.3)$$

where

$$\hat{x}(k) = \hat{x}(k|k-1) \quad (4.4)$$

and $z(k)$ in (4.3) comes from (4.2). If we now let

$$\Phi_1 = \Phi + \Delta\Phi \quad (4.5)$$

$$H_1 = H + \Delta H \quad (4.6)$$

$$\Gamma_1 = \Gamma + \Delta\Gamma \quad (4.7)$$

we can express the residual in the filter:

$$\begin{aligned} \gamma(k) &= z(k) - \hat{Hx}(k) \\ &= (H + \Delta H)x(k) - H\hat{x}(k) + v(k) \end{aligned} \quad (4.8)$$

Letting $e(k)$ denote the estimation error

$$e(k) = x(k) - \hat{x}(k) \quad (4.9)$$

then

$$\begin{aligned} e(k+1) &= \Phi(I-KH)e(k) + (\Gamma+\Delta\Gamma)w(k) + \Delta\Phi x(k) \\ &\quad - \Phi Kv(k) - \Phi K\Delta Hx(k) \end{aligned} \quad (4.10)$$

and

$$\gamma(k) = He(k) + \Delta Hx(k) + v(k) \quad (4.11)$$

Thus, we can think of the residuals as the output of an augmented system with state vector

$$\begin{bmatrix} e(k+1) \\ x(k+1) \end{bmatrix} = \begin{bmatrix} \Phi(I-KH) & \Delta\Phi - \Phi K\Delta H \\ 0 & \Phi + \Delta\Phi \end{bmatrix} \begin{bmatrix} e(k) \\ x(k) \end{bmatrix} + \begin{bmatrix} \Gamma w(k) - \Phi Kv(k) \\ \Gamma w(k) \end{bmatrix} \quad (4.12)$$

$$\gamma(k) = [H \quad \Delta H] \begin{bmatrix} e(k) \\ x(k) \end{bmatrix} + v(k) \quad (4.13)$$

Assuming that $\Phi(I-KH)$ is a stable matrix with real eigenvalues, $e(k)$ and $\gamma(k)$ oscillate when forced by $x(k)$, and therefore at a frequency determined by $\Phi + \Delta\Phi = \Phi_1$, i.e., the real system. If the eigenvalues of $\Phi(I-KH)$ are not real, then the previous statement still holds, but the behavior of $e(k)$ in equation (4.10) will show mixed oscillatory modes.

Returning to our example, $\Phi(I-KH)$ for flight condition 11 is a stable matrix with real eigenvalues as shown by equation (1.65). The residuals in Figure 4.1 in fact oscillate with a period close to that of Φ in flight condition 12. The initial condition in α was chosen to provide the maximum possible effect, as ΔH in equation (4.11) is of the form:

$$\Delta H = \begin{bmatrix} 0 & 0 \\ 0 & \Delta h_{\alpha\alpha} \end{bmatrix} \quad (4.14)$$

It should be pointed out that the value of that initial condition ($q(0)=0$, $\alpha(0)=5^\circ$) is quite large: $\alpha(0)=5^\circ$ has a value twenty times the standard deviation of the noise in the α dynamics. Therefore, it is expected that the effects of mismatch will be less apparent for a more subdued state trajectory.

Finally, in Figure 4.4(a) and 4.4(b), we present the ℓ 's under mismatch between the same two flight conditions -- and for no failures -- when the initial condition is one in q : $q(0) = 26^\circ/\text{sec}$ (20σ) and




Fig. 4.4(a) Likelihood Ratios Under Mismatch:
No Failure; State Step Detector




Fig. 4.4(b) Likelihood Ratios Under Mismatch:
No Failure; Sensor Step Detector

$\alpha(0) = 0$. Although there are some differences in the magnitude of the ℓ 's achieved and, for the sensor step detector, in their shape, all our earlier observations apply equally well here. The interaction between q and α through the dynamics results in similar responses in the detectors, as the pitch rate q is integrated into angle-of-attack α .

In the next section we consider the case when failures occur while a condition of mismatch prevails.

4.2.3 Failure Detection Under Complete Mismatch

In the last section we saw how the sensitivity of the GLR detectors to modeling errors can result in large values of the likelihood ratios, and consequently in false alarms. The response of the detectors to the mismatch was seen to have very definite characteristics, as well as to be quite dependent on the state of the system.

In order to determine how detection might be affected when a failure does occur under these conditions, a number of simulation runs were made for different failures and initial conditions. Figure 4.5(a) shows the ℓ 's in the state detector in response to a 1σ state step failure in q . The initial condition was $q(0)=0$, $\alpha(0)=5^\circ$ and $\theta_t=5$ was the failure time. The ℓ 's in Figure 4.5(b) correspond to the same kind of failure and mismatch, but for $q(0)=0$ and $\alpha(0)=-5^\circ$. The unfailed, mismatched response of the same detector, for comparison, was given in Figure 4.2(a). While the oscillatory behavior of the ℓ 's due to the mismatch are clearly present, it is worth noticing that they (the ℓ 's) are actually growing




Fig. 4.5(a) Likelihood Ratios Under Mismatch:
1σ q State Step Failure; State
State Detector




Fig. 4.5(b) Likelihood Ratios Under Mismatch:
1σ q State Step Failure; State
Step Detector

with time. This is in contrast to the unfailed case where they are expected to decrease with time, in line with the decaying tendency shown by the residuals in Figure 4.1. The reason for the decrease in the residuals is that these oscillations depend on $x(k)$ (see equations (4.10) and (4.11)). Since the system is stable, the effects of the initial conditions diminish with time, in the absence of any forcing input. Notice that this "free" response of the system is mainly determined by the initial state (the noise disturbances are the other factor).

Figure 4.5(b) shows the λ 's, also from the state step detector, for the same failure when the system is started at a different value of the state. While the same comments apply, note the larger values of the λ 's and the apparent change in the oscillations. Figures 4.6(a) and 4.6(b) show the λ 's in the sensor step detector in response to a $10'$ q sensor step failure, with the same two initial conditions. Notice that observations can be made analogous to those for the state step detector. One difference is that for the initial condition with $\alpha(0)=-5^\circ$, the likelihood ratios are actually smaller than in the unfailed case shown in Figure 4.2(b). The effect of the q failure on the value of α is opposite that of the negative initial conditions. Both tend to cancel out somewhat the effect on the residuals.

These simulations indicate that detection is still possible under mismatch, although not without difficulties. The increasing, although oscillatory, nature of the λ 's for failures under these conditions imply




Fig. 4.6(a) Likelihood Ratio Under Mismatch: 1σ q Sensor Step Failure; Sensor Step Detector




Fig. 4.6(b) Likelihood Ratios Under Mismatch: 1σ q Sensor Step Failure; Sensor Step Detector

continued detectability of the failure. By raising the threshold to a very high value we can insure that only failures can lead to detection. The large thresholds necessary to guarantee this will lead to increased delays before a failure can be detected, or decreased probability of detection. However, we have seen in Chapter 2 that step failures, state steps in particular, of moderate size result in extremely high values of the λ 's very quickly. Therefore the real problem is that modeling errors may limit the smallest failures we can hope to detect.

The dependence of the λ 's on the values of the state of the system is a factor which is unique to the condition of mismatch; it was not seen before in matched detection analysis and results. As with the likelihood ratio, the estimates of the failure and failure time $\hat{\nu}$ and $\hat{\theta}$, also depend on the state.

Since some of the failures enter into the state of the system -- state jump and state step failures -- the behavior of the outputs of the detectors for those failures is more sensitive to the state trajectory than the detectors for the sensor failures.

Figures 4.7 and 4.8 display the likelihood ratios under the mismatch between flight conditions 11 and 12 for α step failures ($\theta_t=5$) in the state and sensors. Figures 4.7(a) and 4.8(a) show the λ 's in both the state step and sensor step detectors for an α initial condition of 5° . The corresponding detector response to the same failures for an initial condition of $\alpha=-5^\circ$ is shown in Figures 4.7(b) and 4.8(b). As before, the differences in the state lead to changes in the λ 's. However,




Fig. 4.7(a) Likelihood Ratios Under Mismatch: 1σ α State Step Failure; State Step Detector




Fig. 4.7(b) Likelihood Ratios Under Mismatch: 1σ α State Step Failure; State Step Detector




Fig. 4.8(a) Likelihood Ratios Under Mismatch: 10° α Sensor Step Failure; Sensor Step Detector




Fig. 4.8(b) Likelihood Ratios Under Mismatch: 10° α Sensor Step Failure; Sensor Step Detector

compare the resulting changes in the ℓ 's due to the different state of the system in Figures 4.5 and 4.7 to those changes observed in Figures 4.6 and 4.8. We can see that

- the way in which the ℓ 's respond to the differences in the state is not the same for both detectors: in one case they increase and they decrease in the other
- the ℓ 's in the state step detector change by as much as 100% depending on the state, for the same failure; the ℓ 's in the sensor step detector only differ by about 25% (the value of the state can tend to cancel out part of the failure when it is a state failure -- but not when it is a sensor failure)

Knowledge of these characteristics of the detector responses in the presence of modeling errors can be used to advantage, once understood in more detail.

Just as in the case of cross-detection (Chapter 3), the best way to deal with these difficulties is the full use of the behavior of the ℓ 's, and the estimates $\hat{v}(k)$ and $\hat{\theta}(k)$. The maximum likelihood interpretation of these quantities for the various detector hypotheses can tell us something about the activity of the residuals, and, hence, also about the behavior of the real system. It may be possible to use all this information in an integrated manner in connection with a scheme for system identification.

Finally, in Figures 4.9, (a) and (b), we show the estimates of the failure time for the state step detector. The failures are a 1σ q state step (Figure 4.9 (a)) and a 1σ α state step (Figure 4.9 (b)) and $\theta_t = 5$ for both. They are given the initial condition in α of 5° and also -5° .




Fig. 4.9(a) Estimated Failure Time for a 10 q State Step: State Step Detector




Fig. 4.9(b) Estimated Failure Time for a 10 α State Step: State Step Detector

We notice that the kind of failure -- failure in q or in α -- can enhance or diminish the differences between the resulting $\hat{\theta}(k)$'s for two different values of the initial state. Compare these plots to the unfailed response of the same detector for an initial condition in α of 5° , shown in Figure 4.3(a).

First of all, notice that the failure in α (Fig. 4.9(b)) overrides the effect of the mismatch to some extent and maintains a consistent $\hat{\theta}$, for both state trajectories. The failure in q , on the other hand (Fig. 4.9(a)), shows erratic estimates not too different from the unfailed estimate in Figure 4.3(a). It must be kept in mind, however, that these are small failures and thus the problem is expected to be less severe for larger failures. Also of interest is the way that the sequence $\hat{\theta}(k)$ differs, for the same failure and detector, for two separate state paths. This is particularly evident in Figure 4.9(a). A more detailed analysis of this state dependence might prove fruitful. It may tell us if we may corroborate a detection decision by deliberate perturbation of the actual flight trajectory.

Although the data presented here is limited to a few situations of special interest, we feel that our observations can be generalized. A more complete study of this sensitivity to modeling inaccuracies should be rewarding. In the next section we look at another possible approach to improving detection under mismatch.

4.3 PARTIAL MISMATCH: ΔH

4.3.1 Compensation for ΔH

It was pointed out in section 4.2.1 that one of the major parameter errors in the mismatch between flight conditions corresponds the $H_{\alpha\alpha}$ term of the observation matrix. A number of simulations were made where the only mismatch between the system model and the actual system consisted of the difference in the H. The detector and filter were based on flight condition 11. The real system was the same, except that the H was that of flight condition 12.

Figures 4.10(a) and 4.10(b) show these ℓ 's under the same conditions as in Figure 4.2: no failure and an initial condition in α of 5° . The great similarity to the previous likelihood ratios, under the complete mismatch between the flight conditions, makes it clear that it is in fact H which contributes, in our case, to the main effect. Thus, if one could correct for this parameter error, one would expect to be able to mitigate greatly the sensitivity to the mismatch.

It was mentioned earlier that the $H_{\alpha\alpha}$ term, which is the only one contributing toward ΔH (see equation (4.13)), is a function of the dynamic pressure. Hence, if one can determine the actual H by some means -- such as measuring dynamic pressure directly -- one may be able to compensate for ΔH . Figures 4.11 and 4.12 display the ℓ 's in the state step and sensor step detectors for a complete mismatch where ΔH has been compensated for (i.e., the actual $\Delta H=0$). There is no failure in either




Fig. 4.10(a) Likelihood Ratios Under H Mismatch: No Failure; State Step Detector




Fig. 4.10(b) Likelihood Ratios Under H Mismatch: No Failure; Sensor Step Detector




Fig. 4.11(a) Likelihood Ratios for Partially Compensated Mismatch: No Failure; State Step Detector




Fig. 4.11(b) Likelihood Ratios for Partially Compensated Mismatch: No Failure; Sensor Step Detector




Fig. 4.12(a) Likelihood Ratios for Partially Compensated Mismatch: No Failure; State Step Detector




Fig. 4.12(b) Likelihood Ratios for Partially Compensated Mismatch: No Failure, Sensor Step Detector

case. The ℓ 's in Figure 4.11 correspond to the system started at an α (20σ) initial condition and for the ℓ 's in Figure 4.12 the system started with a 20σ q initial condition ($\alpha=0$). Note that in both cases the values of the ℓ 's are significantly smaller than in the cases where there was no such compensation (Figures 4.2 and 4.4). Selecting a high threshold can eliminate the false alarms without significantly diminishing the capability for detecting failures.

If we recognize that our example consists of relatively large modeling errors and initial conditions, these results suggest that in many cases it may be possible to isolate the main source of the difficulties and to eliminate them.

4.3.2 Approximate Analysis of Complete Mismatch

We mentioned earlier that analytical treatment of the sensitivity of GLR performance to modeling errors is in general intractable. To consider changes in all parameters of the system and their effects on the filter residuals results in too large a burden which overrides any gain in clarity or understanding of the situation. However, some analysis may be possible in those cases where adequate approximations can be made. In this section we present some expressions obtained, for the aircraft example, by recognizing the dominant sources of the mismatch.

In section 4.3.1 it was pointed out that for our example the dominant parameter error was ΔH . This term accounts for the major trends in the behavior of the residuals under mismatch. Figure 4.10 shows that

ΔH alone produces a response in the ℓ 's of the GLR detectors which is qualitatively similar to that shown in Figure 4.2 for the complete mismatch. In Appendix B we have calculated the change in the residual due to the presence of the ΔH between the true system and the model in the filter (which is the same as for the detector).

The residuals in the Kalman filter, expressed in terms of the residuals for the matched, unfailed case ($\tilde{\gamma}(k)$), are then given by

$$\gamma(k) = \tilde{\gamma}(k) + \sum_{\ell=0}^k \Lambda(k, \ell) x(\ell) \quad (4.15)$$

where

$$\Lambda(k, k) = \Delta H \quad (4.16)$$

$$\Lambda(k, \ell) = -H\Theta^{k-\ell-1}\Phi K\Delta H, \quad \ell < k \quad (4.17)$$

and

$$\Theta = \Phi(I-KH) \quad (4.18)$$

Θ is recognized to be the transition matrix of the Kalman filter.

Note that this is another expression for the residuals in (4.13) and that the dependence on the past and present values of the state is made explicit. The quantities which go into $\Lambda(k, \ell)$ are all known and thus it is not hard to compute Λ for increasing values of $\tau=k-\ell$. If knowledge of ΔH can be obtained, or if we can estimate it, it may be possible to filter out a large part of the non-white component of the residuals in (4.15) by subtracting out values computed with an approximation with a small number of terms in that convolution. This pre-filtering would provide the new residuals to be used by the detectors:

$$\gamma_{\text{new}}^{(k)} = \gamma_{\text{old}}^{(k)} - \sum_{\ell=k-T}^k \tilde{\Lambda}(k, \ell) \hat{x}(\ell|k) \quad (4.19)$$

where $\tilde{\Lambda}$ and T are to be optimized. We must use the estimate $\hat{x}(k|k)$ since we do not have x .

Otherwise, the effect of the residuals in (4.15) on the likelihood ratios of the GLR detectors (for ΔH) can be calculated, as in Appendix B, to be:

$$\begin{aligned} E(\lambda(k, \theta)) = & \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) V(j) \right] \\ & + 2 \text{tr} \left[\sum_{h=\theta}^k \sum_{j=\theta}^k \sum_{\ell=0}^j \sum_{i=0}^{\min(\ell-1, h-1)} L(k, \theta; h, j) \Lambda(j, \ell) \Phi^{\ell-i-1} \Gamma' \Gamma (H \Phi^{h-i-1})' \right] \\ & + \text{tr} \left[\sum_{j=\theta}^k \sum_{m=\theta}^k \sum_{\ell=0}^j \sum_{n=0}^m L(k, \theta; j, m) \Lambda(m, n) \Phi^n x_0 x_0' \left(\Lambda(j, \ell) \Phi^\ell \right)' \right] \\ & + \text{tr} \left[\sum_{j=\theta}^k \sum_{m=\theta}^k \sum_{\ell=0}^j \sum_{n=0}^m L(k, \theta; j, m) \Lambda(m, n) \left(\sum_{s=0}^{\min(\ell-1, n-1)} \Phi^{n-s-1} \Gamma' \Gamma \cdot \right. \right. \\ & \left. \left. \cdot \left[\Lambda(j, \ell) \Phi^{\ell-s-1} \right] \right) \right] \end{aligned} \quad (4.20)$$

where

$$V(j) = HP(j)H' + R \quad (4.21)$$

$$P(j) = E(e(j)e(j)') \quad (4.22)$$

$$L(k, \theta; h, j) = V^{-1}(h)G(h, \theta)C^{-1}(k, \theta)G'(j, \theta)V^{-1}(j) \quad (4.23)$$

$e(k)$, θ and Λ are given by equations (4.9) and (4.15)-(4.18) and $x_0 = x(0)$.

Thus we find that even for the idealized mismatch with only ΔH ,

the analysis of the likelihood ratios is not a trivial exercise. Hence, our heavy reliance on simulations is justified as the more immediate and useful approach. Based on the observed patterns in the behavior of the various detector outputs, 'smart' decision rules must be designed and tested for improved performance. Otherwise some form of compensation must be implemented either by the pre-filtering suggested above or by means of the suggestion in section 4.3.1 (correcting for the dominant parameter by directly measuring it, whenever possible).

In the next section we look briefly at another kind of mismatch where the analysis becomes somewhat simpler.

4.4 THE FILTER MATCHED TO THE SYSTEM

So far in this chapter, all the discussion on the limitations of the GLR technique arising from inaccuracies in the modeling of the system of interest has been related to what we call complete mismatch (section 4.2.2). The same model is used to calculate the detector matrices and the gains in the Kalman filter.

Another possibility is having the filter gains matched to the dynamics of the real system and the detector quantities based on the system model. This becomes feasible if the GLR detection system operates simultaneously with an identification scheme -- e.g., the multiple model adaptive control method (MMAC) [18], which is capable of choosing a Kalman filter for the correct flight condition. The detectors for a given flight condition consists of a sequence of matrices $(G(r), C(r): r=0,1, \dots, M-N+1)$

for the interval in the window. Therefore it is desirable to have detectors for a limited number of flight conditions.

Consider the system given by equations (4.1) and (4.2). Suppose, again, that the only difference between this system and its model in the detector is that the latter assumes a value H_d while the actual observation matrix is

$$H_s = H_d + \Delta H \quad (4.24)$$

The subject $(\cdot)_s$ denotes quantities corresponding to the real system while $(\cdot)_d$ denotes quantities based on the model of the detector.

The Kalman filter is matched to the system -- i.e., it is based on H_s .

In the absence of failures, the filter residuals will form a white noise process. However, let us express $\gamma(k)$ in terms of the model in the detector:

$$\begin{aligned} \gamma_s(k) &= z_s(k) - H_s \hat{x}(k) \\ &= (H_d + \Delta H)x(k) - (H_d + \Delta H)\hat{x}(k) + v(k) \\ &= H_d e(k) + v(k) + \Delta H e(k) \\ &= \tilde{\gamma}_d(k) + \Delta H e(k) \\ &= \tilde{\gamma}_s(k) \end{aligned} \quad (4.25)$$

Or,

$$\tilde{\gamma}_s(k) = \tilde{\gamma}_d(k) + \Delta H e(k) \quad (4.26)$$

where $\hat{x}(k)$ and $e(k)$ --- the predicted state estimate and estimation error -- are defined in equations (4.4) and (4.9).

If we compute the expected value of the likelihood ratios in the GLR detectors, based on H_d , for these residuals corresponding to H_s , we obtain (in the absence of failures):

$$E(\tilde{\ell}_{d|s}(k, \theta)) = \text{tr} \sum_{j=\theta}^k L(k, \theta; j) V_s(j) \quad (4.27)$$

$$\text{where } L(k, \theta; j) = L(k, \theta; j, j) \quad (4.28)$$

$$\begin{aligned} V_s &= H_s P H_s' + R \\ &= (H_d + \Delta H) P (H_d + \Delta H)' + R \\ &= H_d P H_d' + R + H_d P \Delta H' + \Delta H P H_d' + \Delta H P \Delta H' \\ &= V_d + \Delta V \end{aligned} \quad (4.29)$$

$$\Delta V = H_d P \Delta H' + \Delta H P H_d' + \Delta H P \Delta H' \quad (4.30)$$

$L(k, \theta; j, j)$ is defined in (4.23), P in (4.22), and V_d is as in (4.21).

The expression $\tilde{\ell}_{d|s}$ refers to the unfailed likelihood ratio as described above and is derived in Appendix C. Thus, we have

$$\begin{aligned} E(\tilde{\ell}_{d|s}(k, \theta)) &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j) V_s(j) \right] = \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j) (V_d(j) + \Delta V(j)) \right] \\ &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j) V_d(j) \right] + \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j) \Delta V(j) \right] \\ &= E(\tilde{\ell}_d(k, \theta)) + \Delta = n + \Delta \end{aligned} \quad (4.31)$$

The last equality follows from the result in Chapter 2 for the matched detection case. Hence, the difference in the ℓ 's in the detectors, due to the error in their model value for H , leads to an increase in the value of the unfailed likelihood ratios. To the extent that ΔH is the most significant modeling error,

$$\begin{aligned} \Delta &= E(\tilde{\ell}_{d|s}(k, \theta)) - E(\tilde{\ell}_d(k, \theta)) \\ &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j) \Delta V(j) \right] \end{aligned} \quad (4.32)$$

predicts the changes needed in the threshold in order to maintain the same detection probabilities (P_d and P_F) as in the case when the detector is perfectly matched to the system. Notice the ΔV , and thus Δ , can be precomputed.

In contrast to the complete mismatch situation in the previous part of this chapter, the approximate analysis for this other kind of detector mismatch is feasible. An analysis of the failure detection performance of the detectors under these conditions can be developed much like in the case of cross-detection in Chapter 3. In Appendix C the following expressions are also derived:

$$E(d_{d|s}(k; \theta)) = E\left(\sum_{j=\theta}^k G'_d(j, \theta) V_d^{-1}(j) \gamma_s(j)\right) = C_{d|s}(k, \theta) v \quad (4.33)$$

$$\delta_{d|s}^2(k, \theta) = v' C'_{d|s}(k, \theta) C_d^{-1}(k, \theta) C_{d|s}(k, \theta) \quad (4.34)$$

$$E(\hat{v}_{d|s}(k)) = C_d^{-1}(k, \hat{\theta}_{d|s}(k)) C_{d|s}(k, \hat{\theta}_{d|s}(k)) v \quad (4.35)$$

and

$$\hat{\theta}_{d|s}(k, \theta) = \arg \max_{\theta} \ell_{d|s}(k, \theta) \quad (4.36)$$

where

$$C_{d|s}(k, \theta) = \sum_{j=\theta}^k G_d'(j, \theta) v_d^{-1}(j) G_s(j, \theta) \quad (4.37)$$

and G_d is the signature based on $H=H_d$. Notice the similarity in form to the same quantities in cross-detection given by equations (3.6), (3.10), (3.12), (3.13) and (3.5), in that order.

Thus, if one can isolate the parameter error which contributes the most to the detector sensitivity, a thorough, although approximate, analysis of the effects on detection is possible. In conjunction with an adaptive estimation control system such as MMAC [18], this may be an attractive approach to dealing with the sensitivity to modeling errors.

4.5 DISCUSSION OF MODELING ERRORS AND THE GLR

In this chapter we have examined some of the implications of inaccuracies in the model parameters used to compute the Kalman filter and detector matrices. This completes our discussion of the limitations inherent in the GLR techniques, which we began in Chapter 3 by looking at distinguishability between the failure modes.

We have looked at the complete mismatch situation -- i.e., where the filter and detector are computed for the same model -- and found that

indeed detection is very sensitive to such errors. However, in view of the large parameter errors assumed, and in the manner in which performance is affected, it is felt that the degradation in detection which may result can be dealt with effectively (i.e., compensating for the modeling errors). It was seen that the detector outputs -- the ℓ 's, $\hat{\theta}(k)$ and $\hat{v}(k)$ -- are affected by the mismatch in a way which is characteristic of the true system. Thus, "smart" decision rules can be constructed which minimize the possibility of false alarms, much as in the case of cross-detection. In addition, some of the information in these outputs may be useful for an identification system.

The correlations in the residuals induced by the parameter errors, even in the absence of failures, are very distinctive. Although analysis of this kind of mismatch is not tractable, it may be possible to isolate the dominant source of error by means of simulations. The results presented with our example indicate that some form of compensation for such errors can reduce the difficulties significantly. In any event, it was seen how parameter errors lead to residual signatures much like those of actual failures. In fact, some modeling errors are equivalent to some of the more complex failure modes mentioned in Chapter 1 -- e.g., an error ΔH is seen to be the same as the hard-over sensor failure. Unlike the simpler failures considered throughout this thesis, the likelihood ratios now depend on the state of the system. This is one of the reasons for the distinctive manner in which the GLR detectors react, corresponding to the actual dynamics associated with the state trajectory.

Viewed in this light, we see that it is possible to detect, and perhaps identify, the more complex failures with detectors based on elementary additive failure models.

Finally, another kind of mismatch was briefly considered where analysis and compensation seems tractable. If in conjunction with an identification technique the filter gains can be made to correspond to the true system, the residuals may remain nearly uncorrelated. We examined an example where the dominant parameter error was assumed to be the only mismatch. The unfailed response of the detectors can then be corrected by a simple change in the threshold. In the event of failure, the characteristics of the detector outputs can be studied analytically, much as in the cross-detection problem considered in the last chapter.

CHAPTER 5

CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

5.1 CONCLUSIONS

We have examined in detail the performance of the generalized likelihood ratio (GLR) technique for failure detection. In particular, we have looked at the full GLR formulation (in which no information about the failure vector v is available) and have applied it to a reduced-order model of the longitudinal dynamics of the F-8C aircraft. A number of failure models have been introduced and the performance of the GLR detectors corresponding to four basic failure modes has been discussed at length.

A qualitative examination of the performance of the full GLR for this application has been made. It has been directed toward an evaluation of the capability of this technique to extract information about the failure, and to the degradation which results from parameter uncertainties. Extensive use of simulations has complemented the analysis of this method.

The failure signatures and GLR detectors for some simple failure modes -- state jumps and steps, sensor jumps and steps -- have been analyzed and discussed in detail. The performance of the full GLR, as measured by such indices as

- . false alarm rates
- . detection delays
- . ability to distinguish among the various failure modes
- . sensitivity to modeling errors

indicates great sensitivity to failures. In most cases (including very small failures) the thresholds can be chosen to guarantee fast detection with few, if any, false alarms. This is true as long as the modeling is accurate. Nevertheless, even under these conditions there are cross-detection problems, i.e., difficulties in selecting the true mode of a failure from among failure modes with correlated signatures. The problem exists if one simply looks at the values of the likelihood ratios for single times (as opposed to their values over intervals of time).

The simulations verify the analysis under matched conditions (exact modeling) and suggest ways to deal with the difficulties. The distinctive patterns in the likelihood ratios indicate that it is possible to develop more sophisticated, or "smart", decision rules which make full use of the available information. These rules -- which remain to be formulated -- should be able to improve detection performance. The basis for these rules is to look for specific behavior of the likelihood ratios, and of the estimates of the failure and its time of occurrence, which is characteristic of each mode. The joint detection, isolation, and estimation of failures can greatly improve overall detection performance (i.e., lower false alarm rates, etc.). When looked at as simultaneous tasks, rather than sequential operations -- i.e., first detect, then isolate the failure mode and finally estimate the failure -- better use is made of the available information.

These comments apply whether distinguishability between failure modes or the effects of parameter errors is the main concern. The

development of such modified decision rules is related to the use of other GLR formulations. It was seen that where full GLR can have serious distinguishability difficulties, the CGLR (and possibly the SGLR) can be of great utility.

Our experience with the GLR method up to now suggests that it is a useful and reliable technique for the detection of failures, or abrupt system changes. The performance of the GLR method can be studied analytically [12]: the analysis becomes simpler for the more restrictive formulation of SGLR. In addition, we have seen that GLR can be successfully examined by way of simulations. Using both approaches, the GLR offers ample flexibility in the range of implementations possible. This allows us to match the technique to the available information about the system and the failures of interest.

5.2 SUGGESTIONS FOR FURTHER WORK

Based on the results obtained, it is felt that future efforts should concentrate on the following immediate issues:

- The study of the correlations between the $\ell_i(k, \theta)$ for the different failure modes. New analytical techniques and performance measures are needed with which to develop and evaluate the suggested "smart" decision rules.
- The development of modifications to the signatures in order to increase the distinguishability of the failure modes. The concepts of "orthogonal" signatures on the one hand (to minimize cross-detection effects), and "universal" signatures

(signatures which can detect all or groups of the failures of interest) on the other should be explored further. Computational and storage savings are issues which can also be dealt with by means of modifications in the detectors.

- Further study of the sensitivity to parameter errors and of the more complex failure models is needed. Their interconnections are important to understand and can contribute to the analysis of "smart" decision rules.
- Work should continue with CGLR and SGLR. It may be possible to obtain improved failure mode distinguishability and reduced sensitivity to parameter errors by looking for failures in specific directions (and/or magnitudes) only.

On a more distant horizon, work should continue toward:

- An integrated study of the options offered by the various GLR formulations, as part of an overall design methodology for maximum utilization of the characteristics of each system.
- The study of tradeoffs between the use of these techniques which rely on analytical redundancy and simpler detection systems relying on hardware redundancy. These complexity-performance tradeoffs should be examined in the light of the GLR and compared to other methods for failure detection.
- Finally, an evaluation of the performance of such failure detection systems within an overall scheme of self-organizing control systems which automatically restructure themselves as compensation to failures in the system. The criteria with which to critically consider the performance of failure detection systems can be different for closed-loop operation. For example, the "price" of a false alarm or delayed detection

can be quite different (from an open-loop situation) if these decisions are used in a feedback control system.

APPENDIX A

THE FAILURE SIGNATURES FOR FAILURE MODELS 1-4

The failure signature matrices $G_i(k, \theta)$ are presented here for the following failure modes: state jumps, state steps, sensor jumps and sensor steps (failure models 1-4 in Chapter 1). These expressions are derived by Chow ([12], Chapter 2).

For the linear dynamic system

$$x(k+1) = \Phi(k)x(k) + w(k) \quad (A.1)$$

$$z(k) = H(k)x(k) + v(k) \quad (A.2)$$

and Kalman-Bucy filter based on the no-failure hypothesis,

$$\hat{x}(k+1|k) = \Phi(k)\hat{x}(k|k) \quad (A.3)$$

$$\hat{x}(k|k) = \hat{x}(k|k-1) + K(k)\gamma(k) \quad (A.4)$$

with $K(k)$, the optimal gains (see (1-18)-(1-21), the measurement residuals are

$$\gamma(k) = \tilde{\gamma}(k) \quad (A.5)$$

where $\tilde{\gamma}(k)$ is a white noise process. In the event of a failure v at time θ in one of the above modes we can write, by linearity,

$$\gamma(k) = \tilde{\gamma}(k) + s(k, v, \theta) \quad (A.6)$$

$$s(k, v, \theta) = G(k, \theta)v \quad (A.7)$$

where $G(k, \theta)$ is non-zero for $k \geq \theta$ and satisfy the following recursive equations. The matrix $F(k, \theta)$ used below is defined by

$$\hat{x}(k|k) = \hat{x}_1(k|k) + \hat{x}_2(k|k) \quad (A.8)$$

$$\hat{x}_2(k|k) = F(k, \theta)v \quad (A.9)$$

where $\hat{x}_1(k|k)$ is the 'unfailed' estimate and $\hat{x}_2(k|k)$ is due solely to the failure. $\bar{\Theta}(k, k-1)$ is the filter transition matrix for the updated estimated in (A.4).

$$\bar{\Theta}(k, k-1) = \bar{\Theta}(k-1) = [I - K(k)H(k)]\Phi(k-1) \quad (A.10)$$

$$\bar{\Theta}(k, j) = \bar{\Theta}(k-1)\bar{\Theta}(k-2) \dots \bar{\Theta}(j) \quad (A.11)$$

$$\Phi(k, j) = \Phi(k)\Phi(k-1) \dots \Phi(j) \quad (A.12)$$

The signatures are:

State Jump Failures (Mode 1)

$$F_1(k, \theta) = \begin{cases} 0 & k < \theta \\ \sum_{j=\theta}^k \bar{\Theta}(k, j)K(j)H(j)\Phi(j, \theta) & k \geq \theta \end{cases} \quad (A.13)$$

$$G_1(k, \theta) = \begin{cases} 0 & k < \theta \\ H(k) [\Phi(k, \theta) - \Phi(k, k-1)F_1(k-1, \theta)] & k \geq \theta \end{cases} \quad (A.14)$$

State Step Failure (Mode 2)

$$F_2(k, \theta) = \begin{cases} 0 & k < \theta \\ \sum_{i=\theta}^k \sum_{j=i}^k \bar{\Theta}(k, j)K(j)H(j)\Phi(j, i) & k \geq \theta \end{cases} \quad (A.15)$$

$$G_2(k, \theta) = \begin{cases} 0 & k < \theta \\ H(k) \left[\sum_{j=\theta}^k \Phi(k, j) - \Phi(k, k-1) F_2(k-1, \theta) \right] & k \geq \theta \end{cases} \quad (\text{A.16})$$

Sensor Jump Failure (mode 3)

$$F_3(k, \theta) = \begin{cases} 0 & k < \theta \\ \bar{\Theta}(k, \theta) K(\theta) & k \geq \theta \end{cases} \quad (\text{A.17})$$

$$G_3(k, \theta) = \begin{cases} 0 & k < \theta \\ I & k = \theta \\ -H(k) \Phi(k, k-1) F_3(k-1, \theta) & k > \theta \end{cases} \quad (\text{A.18})$$

Sensor Step Failure (mode 4)

$$F_4(k, \theta) = \begin{cases} 0 & k < \theta \\ \sum_{j=\theta}^k \bar{\Theta}(k, j) K(j) & k \geq \theta \end{cases} \quad (\text{A.19})$$

$$G_4(k, \theta) = \begin{cases} 0 & k < \theta \\ I & k = \theta \\ I - H(k) \Phi(k, k-1) F_4(k-1, \theta) & k > \theta \end{cases} \quad (\text{A.20})$$

APPENDIX B

Approximate Analysis of the Likelihood Ratios: Partial Mismatch (in H) with the Filter Matched to the GLR Detectors

I. The Residuals

Consider the following dynamical system:

$$x(k+1) = \Phi x(k) + \Gamma w(k) \quad (B.1)$$

$$z(k) = Hx(k) + v(k) \quad (B.2)$$

with $w(k)$ and $v(k)$ uncorrelated, white noise sequences with zero mean value and covariance matrices:

$$E(w(i)w'(j)) = I\delta_{ij} \quad (B.3)$$

$$E(v(i)v'(j)) = I\delta_{ij} \quad (B.4)$$

The corresponding residuals in the optimal Kalman-Bucy filter for this system, with $z(k)$ as input to the filter, are given by

$$\begin{aligned} \tilde{\gamma}(k) &= z(k) - \hat{z}(k|k-1) \\ &= z(k) - H\hat{x}(k|k-1) \end{aligned} \quad (B.5)$$

where $\hat{x}(k|k-1)$ is the (one-step) prediction of the state which is the optimal estimate of $x(k)$ based on the values of $z(i)$ up to and including $z(k-1)$. When the filter has reached steady-state, it evolves with the following dynamics:

$$\hat{x}(k+1|k) = \Phi\hat{x}(k|k-1) + \Phi K[z(k) - H\hat{x}(k|k-1)] \quad (B.6)$$

where K is the steady-state optimal gain (see equations (1.15)-(1.17)).

Rearranging terms we obtain the equivalent formulation:

$$\hat{x}(k+1|k) = [\Phi - \Phi KH] \hat{x}(k|k-1) + \Phi Kz(k) \quad (B.7)$$

or, letting

$$\Theta = \Phi(I - KH), \quad (B.8)$$

the filter transition matrix, we obtain,

$$\hat{x}(k+1|k) = \Theta \hat{x}(k|k-1) + \Phi Kz(k) \quad (B.9)$$

which is the steady-state equation for the optimal estimate.

The (non-recursive) solutions of (B.1) and (B.9) express $x(k)$ and $\hat{x}(k|k-1)$ as follows:

$$x(k) = \Phi^k x_0 + \sum_{\ell=0}^{k-1} \Phi^{k-\ell-1} \Gamma w(\ell) \quad (B.10)$$

$$\hat{x}(k|k-1) = \Theta^k \hat{x}_0 + \sum_{\ell=0}^{k-1} \Theta^{k-\ell-1} \Phi Kz(\ell) \quad (B.11)$$

where

$$x_0 = x(0), \quad \hat{x}(0) = \hat{x}(0|-1) \quad (B.12)$$

If, instead of the observations given by (B.2), the actual input to the filter is

$$\begin{aligned} z_1(k) &= (H + \Delta H) x(k) + v(k) \\ &= Hx(k) + v(k) + \Delta Hx(k) \\ &= z(k) + \Delta Hx(k), \end{aligned} \quad (B.13)$$

then and (B.11) changes to

$$\begin{aligned} \hat{x}_1(k|k-1) &= \Theta^k \hat{x}_0 + \sum_{\ell=0}^{k-1} \Theta^{k-\ell-1} \Phi K[z(\ell) + \Delta Hx(\ell)] \\ &= \hat{x}(k|k-1) + \sum_{\ell=0}^{k-1} \Theta^{k-\ell-1} \Phi K \Delta Hx(\ell) \end{aligned} \quad (B.14)$$

The actual residuals in the Kalman Bucy filter can now be expressed (using (B.5), (B.13) and (B.14)) as

$$\begin{aligned}
 \gamma_1(k) &= z_1(k) - H\hat{x}_1(k|k-1) \\
 &= z(k) - H\hat{x}(k|k-1) + \Delta Hx(k) - H \sum_{\ell=0}^{k-1} \Theta^{k-\ell-1} \Phi_K \Delta Hx(\ell) \\
 &= \tilde{\gamma}(k) + \Delta Hx(k) - H \sum_{\ell=0}^{k-1} \Theta^{k-\ell-1} \Phi_K \Delta Hx(\ell)
 \end{aligned} \tag{B.15}$$

or, more generally,

$$\gamma_1(k) = \tilde{\gamma}(k) + \sum_{\ell=0}^k \Lambda(k, \ell) x(\ell) \tag{B.16}$$

with

$$\Lambda(k, \ell) = \begin{cases} \Delta H & , \quad k=\ell \\ -H\Theta^{k-\ell-1} \Phi_K \Delta H & , \quad k>\ell \end{cases} \tag{B.17}$$

II. The Log-Likelihood Ratios

The detector equation giving the log-likelihood ratios, we recall from Chapter 1,

$$\ell(k, \theta) = d'(k, \theta) C^{-1}(k, \theta) d(k, \theta) \tag{B.18}$$

$$d(k, \theta) = \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \gamma(j) \tag{B.19}$$

$$C(k, \theta) = \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) G(j, \theta) \tag{B.20}$$

When the GLR detectors are computed assuming $z(k)$ as in (B.2) but, instead, we have $z_1(k)$ as in (B.13), we can compute the effect or change in $E(\ell(k, \theta))$.

Substituting $\gamma_1(j)$ from (B.16) for $\gamma(j)$ in (B.19),

$$\begin{aligned}
 d(k, \theta) &= \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) [\tilde{\gamma}(j) + \sum_{\ell=0}^j \Lambda(j, \ell) x(\ell)] \\
 &= \sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \tilde{\gamma}(j) + \sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \\
 &= \tilde{d}(k, \theta) + \sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \tag{B.21}
 \end{aligned}$$

where $\tilde{d}(k, \theta)$ denotes the quantity that would be obtained were everything matched (i.e., if $\Delta H = 0$).

From (B.18), the log-likelihood ratios can be evaluated with the use of (B.21):

$$\begin{aligned}
 \ell(k, \theta) &= \left[\tilde{d}(k, \theta) + \sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \right] C^{-1}(k, \theta) \cdot \\
 &\quad \cdot \left[\tilde{d}(k, \theta) + \sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \right] \\
 &= \tilde{\ell}(k, \theta) + 2\tilde{d}'(k, \theta) C^{-1}(k, \theta) \left[\sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \right] \\
 &\quad + \left[\sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \right]' C^{-1}(k, \theta) \cdot \\
 &\quad \cdot \left[\sum_{j=\theta}^k \sum_{\ell=0}^j G'(j, \theta) V^{-1}(j) \Lambda(j, \ell) x(\ell) \right] \tag{B.22}
 \end{aligned}$$

where $\tilde{\ell}(k, \theta)$ denotes the log-likelihood ratios corresponding to $\tilde{d}(k, \theta)$.

In order to evaluate the expected value of $\ell(k, \theta)$, $E(\ell(k, \theta))$, we proceed with one term at a time. Let us express (B.22) as

$$E(\ell(k, \theta)) = E_1 + 2E_2 + E_3 \quad (\text{B.23})$$

then, taking the terms in order:

$$\begin{aligned} E_1 &= E(\tilde{\ell}(k, \theta)) \\ &= E(\tilde{d}'(k, \theta) C^{-1}(k, \theta) \tilde{d}(k, \theta)) \\ &= E\left(\left[\sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \tilde{\gamma}(j) \right]' C^{-1}(k, \theta) \left[\sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \tilde{\gamma}(j) \right] \right) \\ &= E\left(\left[\sum_{j=\theta}^k \tilde{\gamma}'(j) V^{-1}(j) G(j, \theta) \right] C^{-1}(k, \theta) \left[\sum_{j=\theta}^k G'(j, \theta) V^{-1}(j) \tilde{\gamma}(j) \right] \right) \\ &= E\left(\sum_{i=\theta}^k \sum_{j=\theta}^k \tilde{\gamma}'(i) L(k, \theta; i, j) \tilde{\gamma}(j) \right) \end{aligned} \quad (\text{B.24})$$

where we define

$$L(k, \theta; i, j) \triangleq V^{-1}(i) G(i, \theta) C^{-1}(k, \theta) G'(j, \theta) V^{-1}(j) \quad (\text{B.25})$$

for simplicity.

Continuing with (B.24), applying the expectation to each term in the sum,

$$\begin{aligned} E_1 &= E(\tilde{\ell}(k, \theta)) \\ &= \sum_{i=\theta}^k \sum_{j=\theta}^k E(\tilde{\gamma}'(i) L(k, \theta; i, j) \tilde{\gamma}(j)) \\ &= \sum_{i=\theta}^k \sum_{j=\theta}^k \text{tr}[L(k, \theta; i, j) E(\tilde{\gamma}(i) \tilde{\gamma}'(j))] \end{aligned} \quad (\text{B.26})$$

where we have used the property of the expectation operator by which

$$E(\mathbf{x}'\mathbf{A}\mathbf{y}) = \text{tr}[\mathbf{A}E(\mathbf{y}\mathbf{x}')] \quad (\text{B.27})$$

Since the residuals in the Kalman-Bucy filter are uncorrelated, we have

$$E(\tilde{\gamma}(i)\tilde{\gamma}'(j)) = \begin{cases} 0 & i \neq j \\ \mathbf{v}(j) & i = j \end{cases} \quad (\text{B.28})$$

and consequently,

$$\begin{aligned} E_1 &= \sum_{j=\theta}^k \text{tr} L(k, \theta; j, j) \mathbf{v}(j) \\ &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) \mathbf{v}(j) \right] \end{aligned} \quad (\text{B.29})$$

Proceeding to the second term in (B.22):

$$\begin{aligned} \text{ii) } E_2 &= E \left(\tilde{\mathbf{d}}'(k, \theta) \mathbf{C}^{-1}(k, \theta) \left[\sum_{j=\theta}^k \sum_{\ell=0}^j \mathbf{G}'(j, \theta) \mathbf{v}^{-1}(j) \Lambda(j, \ell) \mathbf{x}(\ell) \right] \right) \\ &= E \left(\left[\sum_{h=\theta}^k \tilde{\gamma}'(h) \mathbf{v}^{-1}(h) \mathbf{G}(h, \theta) \right] \mathbf{C}^{-1}(k, \theta) \left[\sum_{j=\theta}^k \sum_{\ell=0}^j \mathbf{G}'(j, \theta) \mathbf{v}^{-1}(j) \Lambda(j, \ell) \mathbf{x}(\ell) \right] \right) \\ &= E \left(\sum_{h=\theta}^k \sum_{j=\theta}^k \sum_{\ell=0}^j \tilde{\gamma}'(h) L(k, \theta; h, j) \Lambda(j, \ell) \mathbf{x}(\ell) \right) \end{aligned} \quad (\text{B.30})$$

where use was made of (B.25). Continuing,

$$E_2 = \sum_{h=\theta}^k \sum_{j=\theta}^k \sum_{\ell=0}^j \text{tr} [L(k, \theta; h, j) \Lambda(j, \ell) E(\mathbf{x}(\ell) \tilde{\gamma}'(h))] \quad (\text{B.31})$$

where, using (B.10),

$$\begin{aligned}
 E(\mathbf{x}(\ell)\tilde{\gamma}'(h)) &= E(\phi^\ell \mathbf{x}_0 \tilde{\gamma}'(h)) + E\left(\sum_{i=0}^{\ell-1} \phi^{\ell-i-1} \Gamma \mathbf{w}(i) \tilde{\gamma}'(h)\right) \\
 &= \sum_{i=0}^{\ell-1} \phi^{\ell-i-1} \Gamma E(\mathbf{w}(i) \tilde{\gamma}'(h))
 \end{aligned} \tag{B.32}$$

by linearity of the expectation, and assuming that \mathbf{x}_0 and $\tilde{\gamma}$ are uncorrelated (since \mathbf{x}_0 is assumed uncorrelated with both \mathbf{w} and \mathbf{v} in (B.1) and (B.2)).

In order to obtain the value of $E(\mathbf{w}(i)\tilde{\gamma}'(h))$ we need the following expressions:

$$\begin{aligned}
 \tilde{\gamma}(k) &= \mathbf{z}(k) - \hat{\mathbf{H}}\mathbf{x}(k|k-1) \\
 &= \mathbf{H}\mathbf{x}(k) + \mathbf{v}(k) - \hat{\mathbf{H}}\hat{\mathbf{x}}(k|k-1) \\
 &= \mathbf{H}\mathbf{e}(k|k-1) + \mathbf{v}(k)
 \end{aligned} \tag{B.33}$$

$$\mathbf{e}(k|k-1) = \mathbf{x}(k) - \hat{\mathbf{x}}(k|k-1) \tag{B.34}$$

and

$$\mathbf{e}(k|k-1) = \theta^k \mathbf{e}_0 + \sum_{r=0}^{k-1} \theta^{k-r-1} [\Gamma \mathbf{w}(r) - \Phi \mathbf{K} \mathbf{v}(r)] \tag{B.35}$$

Equation (B.35) can be derived from equations (B.1), (B.2), (B.7) and (B.34). Then, by (B.33),

$$\begin{aligned}
 E(\mathbf{w}(i)\tilde{\gamma}'(h)) &= E(\mathbf{w}(i)\mathbf{e}'(h|h-1))\mathbf{H}' + E(\mathbf{w}(i)\mathbf{v}'(h)) \\
 &= E(\mathbf{w}(i)\mathbf{e}'(h|h-1))\mathbf{H}'
 \end{aligned} \tag{B.36}$$

since the noises ω and v are uncorrelated for all times. Therefore,

$$\begin{aligned}
 E(w(i)\tilde{\gamma}(h)) &= E(w(i)e_0') [\Theta^h]' H' \\
 &+ E(w(i) \sum_{r=0}^{h-1} w'(r) \Gamma' (\Theta^{h-r-1})') H' \\
 &+ E(w(i) \sum_{r=0}^{h-1} v'(r) K' \Phi' (\Theta^{h-r-1})') H' \\
 &= \sum_{r=0}^{h-1} E(w(i)w'(r)) \Gamma' (\Theta^{h-r-1})' H' \\
 &= \begin{cases} [H\Theta^{h-i-1}\Gamma]' & r=i, i \leq h-1 \\ 0 & \text{otherwise} \end{cases} \quad (B.37)
 \end{aligned}$$

due to the fact that: 1) ω and e_0 are uncorrelated, 2) equation (B.3).

Consequently, (B.32) becomes

$$E(x(\ell)\tilde{\gamma}'(h)) = \sum_{i=0}^{\ell-1} \Phi^{\ell-i-1} \Gamma' (\Theta^{h-i-1})' H' \quad (B.38)$$

and E_2 , from (B.31), is given by

$$\begin{aligned}
 E_2 &= \sum_{h=\theta}^k \sum_{j=\theta}^k \sum_{\ell=0}^j \text{tr}[L(k,\theta;h,i)\Lambda(j,\ell)] \sum_{i=0}^{\min(\ell-1,h-1)} \Phi^{\ell-i-1} \Gamma' (\Theta^{h-i-1})' H' \\
 &= \text{tr} \sum_{h=\theta}^k \sum_{j=\theta}^k \sum_{\ell=0}^j \sum_{i=0}^{\min(\ell-1,h-1)} L(k,\theta;h,j)\Lambda(j,\ell) \Phi^{\ell-i-1} \Gamma' (\Theta^{h-i-1})' H' \quad (B.39)
 \end{aligned}$$

where the limits on the sums are constrained to values such that the expression makes sense. Finally we get to the third term, E_3 , in (B.23).

iii) We will need the cross-correlation function for $x(k)$. For $x(k)$ as given in (B.1), with $X_0 = E(x_0 x_0')$,

$$\begin{aligned}
 E(x(n)x'(\ell)) &= E\left(\left[\Phi^n x_0 + \sum_{s=0}^{n-1} \Phi^{n-s-1} \Gamma_w(s)\right] \left[\Phi^\ell x_0 + \sum_{t=0}^{\ell-1} \Phi^{\ell-t-1} \Gamma_w(t)\right]\right) \\
 &= \Phi^n X_0 (\Phi^\ell)' + \sum_{s=0}^{\min(n-1, \ell-1)} \Phi^{n-s-1} \Gamma \Gamma' (\Phi^{\ell-s-1})', \quad (B.40)
 \end{aligned}$$

There, using (B.22), E_3 becomes

$$\begin{aligned}
 E_3 &= E\left(\sum_{j=\theta}^k \sum_{m=\theta}^k \sum_{\ell=0}^j \sum_{n=0}^m x'(\ell) \Lambda'(j, \ell) L(k, \theta; j, m) \Lambda(m, n) x(n)\right) \\
 &= \text{tr} \left[\sum_{j=\theta}^k \sum_{m=\theta}^k \sum_{\ell=0}^j \sum_{n=0}^m L(k, \theta; j, m) \Lambda(m, n) E(x(n)x'(\ell)) \Lambda'(j, \ell) \right] \\
 &= \text{tr} \left[\sum_{j=\theta}^k \sum_{m=\theta}^k \sum_{\ell=0}^j \sum_{n=0}^m L(k, \theta; j, m) \Lambda(m, n) \cdot \right. \\
 &\quad \left. \cdot \left(\Phi^n X_0 (\Phi^\ell)' + \sum_{s=0}^{\min(\ell-1, n-1)} \Phi^{n-s-1} \Gamma \Gamma' (\Phi^{\ell-s-1})' \right) \Lambda'(j, \ell) \right] \quad (B.41)
 \end{aligned}$$

We now have a full expression for $E(\ell(k, \theta))$ for a mismatch in H , as given by (B.22) and (B.23). E_1 , E_2 and E_3 are given by (B.29), (B.39) and (B.41), respectively.

APPENDIX C

APPROXIMATE ANALYSIS: PARTIAL MISMATCH (IN H) WITH THE FILTER MATCHED TO THE SYSTEM

I. Without Failures

Suppose that for a system represented by

$$x(k+1) = x(k) + w(k) \quad (C.1)$$

$$z_d(k) = H_d x(k) + v(k) \quad (C.2)$$

GLR detectors are computed for some failure mode. However, let us assume that the actual system observations really correspond to

$$\begin{aligned} z_s(k) &= H_s x(k) + v(k) \\ &= (H_d + \Delta H) x(k) + v(k) \\ &= z(k) + \Delta H x(k) \end{aligned} \quad (C.3)$$

with H_s also being the value used in the Kalman-Bucy filter equations for the estimates and optimal gains.

The residuals in the filter form an uncorrelated, white noise process with zero mean and with covariance

$$\begin{aligned} \text{Cov}(\tilde{\gamma}(k)) &= E(\tilde{\gamma}(k)\tilde{\gamma}'(k)) \\ &= H_s P(k|k-1)H_s^T + R \\ &= V_s(k) \end{aligned} \quad (C.4)$$

where $P(k|k-1)$ is the predicted covariance of the estimate error, and R , the covariance matrix of $v(k)$. We have

$$\begin{aligned}
 \tilde{\gamma}_s(k) &= z_s(k) - H_s \hat{x}(k|k-1) \\
 &= (H_d + \Delta H)x(k) - (H_d + \Delta H)x(k) \\
 &= H_d(x(k) - \hat{x}(k|k-1)) + \Delta H(x(k) - \hat{x}(k|k-1)) + v(k) \quad (C.5)
 \end{aligned}$$

or,

$$\tilde{\gamma}_s(k) = \tilde{\gamma}_d(k) + \Delta H(x(k) - \hat{x}(k|k-1)) \quad (C.6)$$

where $\tilde{\gamma}_d$ denotes

$$\begin{aligned}
 \tilde{\gamma}_d(k) &= H_d(x(k) - \hat{x}(k|k-1)) + v(k) \\
 &= z_d(k) - H_d \hat{x}(k|k-1) \quad (C.7)
 \end{aligned}$$

which would be their value if $\Delta H=0$.

We will now see the effect on the GLR detectors. The subscript d/s will indicate quantities in a detector (which is based on H_d) when the residuals for its input are $\gamma_s(k)$. First, in the absence of failures (from (B.18) - (B.20)),

$$\begin{aligned}
 \tilde{d}_{d/s}(k, \theta) &= \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) \tilde{\gamma}_s(j) \\
 &= \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) [\tilde{\gamma}_d(j) + \Delta H(x(j) - \hat{x}(j|j-1))] \\
 &= \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) \tilde{\gamma}_d(j) + \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) \Delta H[x(j) - \hat{x}(j|j-1)] \\
 &= \tilde{d}_d(k, \theta) + \Delta d(k, \theta) \quad (C.8)
 \end{aligned}$$

where we let

$$\Delta d(k, \theta) = \sum_{j=\theta}^k G_d'(j, \theta) v_d^{-1}(j) \Delta \gamma(j) \quad (C.9)$$

$$\begin{aligned} \Delta \gamma(j) &= \Delta H[x(j) - \hat{x}(j|j-1)] \\ &= \Delta H e(j|j-1) \end{aligned} \quad (C.10)$$

Then the log-likelihood ratios become

$$\begin{aligned} \tilde{\ell}_{d|s}(k, \theta) &= \tilde{d}_{d|s}'(k, \theta) C_d^{-1}(k, \theta) \tilde{d}_{d|s}(k, \theta) \\ &= [\tilde{d}_d(k, \theta) + \Delta d(k, \theta)]' C_d^{-1}(k, \theta) [\tilde{d}_d(k, \theta) + \Delta d(k, \theta)] \\ &= \tilde{d}_d'(k, \theta) C_d^{-1}(k, \theta) \tilde{d}_d(k, \theta) + \tilde{d}_d'(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta) \\ &\quad + \Delta d'(k, \theta) C_d^{-1}(k, \theta) \tilde{d}_d(k, \theta) + \Delta d'(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta) \end{aligned} \quad (C.11)$$

or

$$\tilde{\ell}_{d|s}(k, \theta) = \tilde{\ell}_d(k, \theta) + \Delta \ell(k, \theta) + \Delta \ell'(k, \theta) + \Delta^2 \ell(k, \theta) \quad (C.12)$$

with

$$\Delta \ell(k, \theta) = \tilde{d}_d'(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta) \quad (C.13)$$

$$\Delta^2 \ell(k, \theta) = \Delta d'(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta) \quad (C.14)$$

C_d is the detector information matrix and \tilde{d}_d , $\tilde{\ell}_d$ are \tilde{d} and $\tilde{\ell}$ for $\Delta H=0$.

If we take expectations in (C.12),

$$E(\tilde{\ell}_{d|s}(k, \theta)) = E_1 + E_2 + E_2' + E_3 \quad (C.15)$$

where

$$\begin{aligned}
 E_1 &= E(\tilde{\ell}_d(k, \theta)) \\
 &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) V_d(j) \right] \tag{C.16}
 \end{aligned}$$

$$E_2 = E(\Delta \ell(k, \theta)) \tag{C.17}$$

$$E_3 = E(\Delta^2 \ell(k, \theta)) \tag{C.18}$$

The expression for L in (C.16) was given in Appendix B, (B.25).

If we proceed to evaluate E_2 , using (C.13):

$$\begin{aligned}
 E_2 &= E(\tilde{d}'_d(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta)) \\
 &= E \left(\left[\sum_{i=\theta}^k G'_d(i, \theta) V_d^{-1}(i) \tilde{\gamma}_d(i) \right]' C_d^{-1}(k, \theta) \left[\sum_{j=\theta}^k G'_d(j, \theta) V_d^{-1}(j) \Delta \gamma(j) \right] \right) \\
 &= E \left(\sum_{i=\theta}^k \sum_{j=\theta}^k \tilde{\gamma}'_d(i) V_d^{-1}(i) G_d(i, \theta) C_d^{-1}(k, \theta) G'_d(j, \theta) V_d^{-1}(j) \Delta \gamma(j) \right) \\
 &= E \left(\sum_{i=\theta}^k \sum_{j=\theta}^k \tilde{\gamma}'_d(i) L(k, \theta; i, j) \Delta \gamma(j) \right) \tag{C.19}
 \end{aligned}$$

With the use of (C.10) we obtain, by linearity,

$$\begin{aligned}
 E_2 &= \sum_{i=\theta}^k \sum_{j=\theta}^k E(\tilde{\gamma}'_d(i) L(k, \theta; i, j) \Delta H[x(j) - \hat{x}(j|j-1)]) \\
 &= \sum_{i=\theta}^k \sum_{j=\theta}^k \text{tr} [L(k, \theta; i, j) \Delta H E([x(j) - \hat{x}(j|j-1)] \tilde{\gamma}'_d(i))] \\
 &= \text{tr} \sum_{i=\theta}^k \sum_{j=\theta}^k L(k, \theta; i, j) \Delta H E(e(j|j-1) \tilde{\gamma}'_d(j)) \tag{C.20}
 \end{aligned}$$

where

$$\begin{aligned}
 E(e(j|j-1)\tilde{\gamma}'_d(i)) &= E(e(j|j-1)[H'_d e(i|i-1) + v(i)]') \\
 &= E(e(j|j-1)e'(i|i-1))H'_d + E(e(j|j-1)v'(i)) \\
 &= \begin{cases} P(j|j-1)H'_d & i=j \\ 0 & i \neq j \end{cases} \quad (C.21)
 \end{aligned}$$

We have made use of the facts that the optimal error is not correlated with itself shifted in time, nor with $v(i)$ for any time.

Consequently, the expression for E_2 becomes

$$E_2 = \text{tr} \left[\sum_{j=0}^k L(k, \theta; j, j) \Delta H P(j|j-1) H'_d \right] \quad (C.22)$$

and similarly,

$$E'_2 = \text{tr} \left[\sum_{j=0}^k L(k, \theta; j, j) H_d P(j|j-1) \Delta H' \right] \quad (C.23)$$

Finally, E_3 is given by

$$\begin{aligned}
 E_3 &= E(\Delta d'(k, \theta) C_d^{-1}(k, \theta) \Delta d(k, \theta)) \\
 &= E \left(\left[\sum_{i=0}^k \Delta \gamma'(i) v_d^{-1}(i) G_d(i, \theta) \right] C_d^{-1}(k, \theta) \left[\sum_{j=0}^k G'_d(j, \theta) v_d^{-1}(j) \Delta \gamma(j) \right] \right) \\
 &= \sum_{i=0}^k \sum_{j=0}^k E(e'(i|i-1) \Delta H' L(k, \theta; i, j) \Delta H e(j|j-1)) \\
 &= \sum_{i=0}^k \sum_{j=0}^k \text{tr} [L(k, \theta; i, j) \Delta H E(e(j|j-1)e'(i|i-1) \Delta H')] \\
 &= \text{tr} \left[\sum_{j=0}^k L(k, \theta; j, j) \Delta H P(j|j-1) \Delta H' \right] \quad (C.24)
 \end{aligned}$$

where we have used (C.9), (C.10), and since $E(e(j|j-1)e'(i|i-1)) = 0$ for $i \neq j$.

Collecting E_1 , E_2 , E_2' and E_3 for $E(\tilde{\ell}_{d|s})$ in (C.15), we obtain (with (C.16), (C.22)-(C.24)),

$$E(\tilde{\ell}_{d|s}(k, \theta)) = \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) (V_d(j) + \Delta H P(j|j-1) H_d' + H_d P(j|j-1) \Delta H' + \Delta H P(j|j-1) \Delta H') \right] \quad (C.25)$$

and by substituting the expression for $V_d(j)$

$$V_d(j) = H_d P(j|j-1) H_d' + R \quad (C.26)$$

and regrouping terms in (C.25),

$$\begin{aligned} E(\tilde{\ell}_{d|s}(k, \theta)) &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) (R + [H_d + \Delta H] P(j|j-1) [H_d + \Delta H]') \right] \\ &= \text{tr} \sum_{j=\theta}^k L(k, \theta; j, j) (H_s P(j|j-1) H_s' + R) \\ &= \text{tr} \left[\sum_{j=\theta}^k L(k, \theta; j, j) V_s(j) \right] \end{aligned} \quad (C.27)$$

by use of (C.4).

II. With Failures

If a failure v takes place at time θ the residuals change and can be represented as

$$\gamma_s(k) = \tilde{\gamma}_s(k) + G_s(k, \theta)v \quad (C.28)$$

where $\tilde{\gamma}_s$ are the unfailed residuals in (C.5) when the input to the filter is z_s (in (C.3)).

In the detectors we get

$$\begin{aligned} d_{d|s}(k, \theta) &= \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) \gamma_s(j) \\ &= \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) [\tilde{\gamma}_s(j) + G_s(j, \theta)v] \\ &= \tilde{d}_{d|s}(k, \theta) + C_{d|s}(k, \theta)v \end{aligned} \quad (C.29)$$

where

$$C_{d|s}(k, \theta) = \sum_{j=\theta}^k G'_d(j, \theta) v_d^{-1}(j) G_s(j, \theta) \quad (C.30)$$

and $\tilde{d}_{d|s}$ was given in (C.8), by which we can see that

$$E(d_{d|s}(k, \theta)) = C_{d|s}(k, \theta)v \quad (C.31)$$

The log-likelihood ratios for the failure v now becomes:

$$\begin{aligned} \ell_{d|s}(k, \theta) &= d'_{d|s}(k, \theta) C_d^{-1}(k, \theta) d_{d|s}(k, \theta) \\ &= [\tilde{d}_{d|s}(k, \theta) + C_{d|s}(k, \theta)v]' C_d^{-1}(k, \theta) [\tilde{d}_{d|s}(k, \theta) + C_{d|s}(k, \theta)v] \\ &= \tilde{d}'_{d|s}(k, \theta) C_d^{-1}(k, \theta) \tilde{d}_{d|s}(k, \theta) + 2v' C'_{d|s}(k, \theta) C_d^{-1}(k, \theta) \tilde{d}_{d|s}(k, \theta) \\ &\quad + v' C'_{d|s}(k, \theta) C_d^{-1}(k, \theta) C_{d|s}(k, \theta)v \end{aligned} \quad (C.32)$$

Recognizing the first term as $\tilde{\ell}_{d|s}(k, \theta)$ in (C.11), with its expectation in (C.27), and noticing that $\tilde{d}_{d|s}$ has zero mean value (see (C.8)-(C.10)) we can write,

$$E(\ell_{d|s}(k, \theta)) = E(\tilde{\ell}_{d|s}(k, \theta)) + v' C_{d|s}'(k, \theta) C_d^{-1}(k, \theta) C_{d|s}(k, \theta) v \quad (C.33)$$

In analogy to the matched case, we call the non-centrality parameter

$$\delta_{d|s}^2(k, \theta) = v' C_{d|s}'(k, \theta) C_d^{-1}(k, \theta) C_{d|s}(k, \theta) v \quad (C.34)$$

This expression is similar to the cross-detection δ^2 and can be rearranged

$$\begin{aligned} \delta_{d|s}^2(k, \theta) &= v' C_{d|s}'(k, \theta) C_d^{-1}(k, \theta) \cdot I \cdot C_{d|s}(k, \theta) v \\ &= [v' C_{d|s}'(k, \theta) C_d^{-1}(k, \theta)] C_d^{-1}(k, \theta) [C_d^{-1}(k, \theta) C_{d|s}(k, \theta) v] \\ &= v_{d|s}' C_d(k, \theta) v_{d|s} \end{aligned} \quad (C.35)$$

where

$$v_{d|s} = C_d^{-1}(k, \theta) C_{d|s}(k, \theta) v = v_{d|s}(k, \theta) \quad (C.36)$$

is a transformed failure vector.

Finally, the failure estimates can be seen to be

$$\hat{\theta}_{d|s}(k) = \arg \max_{\theta} \ell_{d|s}(k, \theta) \quad (C.37)$$

and

$$\hat{v}_{d|s}(k) = C_d^{-1}(k, \hat{\theta}_{d|s}(k)) d_{d|s}(k, \hat{\theta}_{d|s}(k)) \quad (C.38)$$

Evaluating its expectation using (C.31) and (C.36),

$$\begin{aligned} E(\hat{v}_{d|s}(k)) &= C_d^{-1}(k, \hat{\theta}_{d|s}(k)) E(d_{d|s}(k, \hat{\theta}_{d|s}(k))) \\ &= C_d^{-1}(k, \hat{\theta}_{d|s}(k)) C_{d|s}(k, \hat{\theta}_{d|s}(k)) v \\ &= v_{d|s}(k, \theta) \end{aligned} \quad (C.39)$$

REFERENCES

- [1] Jazwinski, A.H., Stochastic Processes and Filtering Theory, Academic Press, New York, 1970.
- [2] Tsypkin, Ya. Z., "Adaptation, Training and Self-Organization in Automatic Systems", Automation and Remote Control, Vol. 27, No. 1, Jan. 1966, pp. 16-51.
- [3] Jazwinski, A.H., "Adaptive Filtering", Automatica, Vol. 5, 1969, pp. 475-485.
- [4] Willsky, A.S., "A Survey of Design Methods for Failure Detection in Dynamic Systems", Automatica, November, 1976.
- [5] Willsky, A.S. and H.L. Jones, "A Generalized Likelihood Ratio Approach to State Estimation in Linear Systems Subject to Abrupt Changes", Proc. of the 1974 IEEE Conference on Decision and Control, Phoenix, Arizona, November 1974.
- [6] Willsky, A.S. and H.L. Jones, "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems", IEEE Trans. on Automatic Control, AC-21, No. 1, Feb. 1976, pp. 108-112.
- [7] Beard, R.V., Failure Accomodation in Linear Systems Through Self-Reorganization, Rept. MVT-71-1, Man Vehicle Laboratory, Cambridge, Massachusetts, February 1971.
- [8] Jones, H.L., Failure Detection in Linear Systems, Ph.D. Thesis, Dept. of Aeronautics and Astronautics, M.I.T., September 1973.
- [9] Chien, Tze-Thong, "An Adaptive Technique for a Redundant Sensor Navigation System", Sc.D. Thesis, Dept. of Aeronautics and Astronautics, M.I.T., January 1972.
- [10] Lumel'skii, V.Ya., "Algorithm for Detecting the Time of Change of Properties of a Random Process", Automation and Remote Control, Vol. 33, No. 10, Part 1, October 1972, pp. 1620-1625.
- [11] Deckert, J.C., M.N. Desai, J.J. Deyst, and A.S. Willsky, "Dual Redundant Sensor FDI Techniques Applied to the NASA DFBW Aircraft", IEEE Trans. on Automatic Control, to appear.
- [12] Chow, E.Y., Analytical Studies of the Generalized Likelihood Ratio Technique for Failure Detection, S.M. Thesis, Dept. of Elec. Eng. and Comp. Sci., M.I.T., February 1976.

- [13] Bueno, R., E.Y. Chow, S. Gershwin and A.S. Willsky, "Research Status Report to NASA Langley Research Center: A Dual-Mode Generalized Likelihood Ratio Approach to Self-Reorganizing Digital Flight Control System Design", M.I.T. Electronic Systems Laboratory, Cambridge, Mass., Paper No. P-633, November 1975.
- [14] Schweppe, F.C., Uncertain Dynamic Systems, Prentice-Hall, New Jersey, 1973.
- [15] Van Trees, H.L., Detection, Estimation and Modulation Theory, Part I: Detection, Estimation and Linear Modulation Theory, John Wiley and Sons, Inc., New York, 1971.
- [16] Athans, M. and K.-P. Dunn, "Linearized Equations for the Continuous Time LQG Problem for the F-8 Aircraft Longitudinal Dynamics", Interim Report #1, ESL-IR-549, E.S.L. M.I.T. April 1, 1974.
- [17] Boozer, D.D. and W.L. McDaniel, Jr., "On Innovation Sequence Testing of the Kalman Filter" IEEE Trans. on Automatic Control, AC-17, No. 1, February 1972, pp. 158-160.
- [18] Athans, M., K.-P. Dunn, C.S. Greene, W.H. Lee, N.R. Sandell, I. Segall and A.S. Willsky, "The Stochastic Control of the F-8C Aircraft Using the Multiple Model Adaptive Control (MMAC) Method, Proc. 1975 IEEE Conf. on Decision and Control, Houston, Texas, December, 1975.