

Bayesian Design of Decision Rules for Failure Detection

EDWARD Y. CHOW, Member, IEEE
Schlumberger-Doll Research

ALAN S. WILLISKY, Senior Member, IEEE
Massachusetts Institute of Technology

The formulation of the decision making process of a failure detection algorithm as a Bayes sequential decision problem provides a simple conceptualization of the decision rule design problem. As the optimal Bayes rule is not computable, a methodology that is based on the Bayesian approach and aimed at a reduced computational requirement is developed for designing suboptimal rules. A numerical algorithm is constructed to facilitate the design and performance evaluation of these suboptimal rules. The result of applying this design methodology to an example shows that this approach is potentially a useful one.

I. INTRODUCTION

A failure detection and identification (FDI) process consists of two basic stages: residual generation and decision making. In the first stage, sensor outputs are processed to form residuals that typically have distinct characteristics under normal (no-fail) conditions and under the various possible failures modes. (See [1] for a discussion of the design of residual generation processes.) The function of the second stage is to monitor the residuals and make decisions concerning the occurrence and identity of failure modes. The decision mechanism is based on a compromise among speed of detection, false-alarm rates, and identification accuracy, and it belongs to the extensively studied class of sequential tests or sequential decision rules [2-15]. Most previous works, however, were focused on either the detection of a single type of change (failure) [5-9], or the sequential testing of M hypothesis, which is analogous to the problem of identifying the failure mode given the onset time is known [12-14]. In this paper, we employ the Bayesian approach to the design of decision rules that *directly confront the problem of detecting and distinguishing the various possible failure modes which may occur at unknown times.*

In Section II we describe the Bayes formulation of the FDI decision problem. Although the optimal rule is generally not computable, the structure of the Bayesian approach can be used to derive practical suboptimal rules. The design of suboptimal rules based on the Bayes formulation is discussed in Section III. The approximations and simplifications that are made in order to obtain these rules make systematic use of the important features specific to the problem of dynamic failure detection and consequently allow us to interpret each step in our simplification procedure in terms of its implications for failure detection. In Section IV we report on our experience with this approach to designing decision rules through a numerical example and simulation.

II. THE BAYESIAN APPROACH

In this section we adapt and specialize the standard Bayes sequential decision problem (BSDP) [15] to the problem of failure detection. The BSDP formulation of the FDI problem consists of six elements:

1) Θ : the set of states of nature or failure hypotheses. For simplicity in this development we assume that only single failures may occur. In general an element θ of Θ conveys several pieces of information, namely, the type of failure mode, its time of occurrence, and probably a variable specifying the severity of the failure. For example, if a particular failure mode corresponds to the onset of a sensor bias, the level of this bias could be specified in the corresponding element of θ . In many applications, however, it suffices simply to identify the failure type without estimating its severity. Furthermore,

Manuscript received March 8, 1983.

This work was supported in part by the Office of Naval Research under Contract N00014-77-C-0224 and in part by NASA Ames Research Center under Grant NGL-22-009-124.

Authors' addresses: E. Y. Chow, Schlumberger-Doll Research, P.O. Box 307, Ridgefield, CT 06877; A. S. Willsky, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139.

0018-9251/84/1100-0761 \$1.00 © 1984 IEEE

what is often done to eliminate this nuisance parameter completely is to hypothesize a fixed scale for each failure type corresponding to the smallest deviation from normal behavior that one would like to detect. For example, this approach was used with great success for the detection of aircraft sensor failure in [16]. We will adapt this approach here, and consequently elements of θ are 2-tuples, $\theta = (i, \tau)$, corresponding to the onset of the i th failure mode at time τ . We assume that there are M hypothesized failure modes and also denote by $(0, -)$ that element of Θ corresponding to no failure. Thus,

$$\Theta = \{(i, \tau), i = 1, \dots, M, \tau = 1, 2, \dots\} \cup \{(0, -)\}$$

2) μ : the prior probability mass function (PMF) over the nature set Θ . This PMF represents the a priori information concerning possible failures, i.e., how likely it is for each type of failure to occur, and when is a failure likely to occur. Because this information may not be available or accurate in some cases, the need to specify μ is a drawback of the Bayes approach for such cases. Nevertheless, we will see that it can be regarded as a design parameter in the specification of the Bayes rule.

In general, μ may be arbitrary. Here, we assume the underlying failure process has two properties: 1) the occurrence of each of the M failure modes is independent of the other, and 2) the occurrence of each failure i is a Bernoulli process with (success) parameter ρ_i , a common model for failures in physical components. The independent assumption is also a reasonable one in most applications. It is straightforward to show that

$$\mu(i, \tau) = a(i) \rho(1 - \rho)^{\tau-1} \quad i = 1, \dots, M, \quad \tau = 1, 2, \dots$$

where

$$\rho = 1 - \prod_{j=1}^M (1 - \rho_j)$$

$$a(i) = \rho_i (1 - \rho_i)^{-1} \left[\sum_{j=1}^M \rho_j (1 - \rho_j)^{-1} \right]^{-1}.$$

The parameter ρ may be regarded as the parameter of the combined (Bernoulli) failure process which specifies the statistics of the occurrence of the first failure; $a(i)$ can be interpreted as the marginal probability that the first failure is of type i . Note that the present choice of μ indicates that the arrival of the first failure is memoryless.

3) $D(k)$: The discrete set of terminal decisions available to the decision maker when the residual-monitoring is interrupted at time k in order to make failure identification. An element δ of $D(k)$ may denote the pair (j, t) , i.e., the declaration of a type j failure to have occurred at time $t \leq k$. Alternatively, δ may represent an identification of the j th failure without regard for the failure time, or it may signify the presence of a failure without specifying its type or time, i.e., simply an alarm. Note that the number of terminal decisions specifying failure times grow with k (as there are more times at which a failure could have occurred) while the number of decisions not specifying times will remain the

same. In addition, $D(k)$ does not include the declaration of no-failure, since the residual monitoring is stopped only when a failure appears to have occurred. It is worth pointing out that in some application one may not be interested in estimating failure onset times, there are others in which one is. For example, if a failed sensor has been used for sometime in a closed-loop filter and control law, one may wish to estimate how long the failure has been present in order to compensate for the effect of this erroneous signal. In addition, onset time estimates are critical in other event detection problems such as electrocardiogram analysis [17] and maneuver detection [18, 19].

4) $L(k; \theta, \delta)$: the terminal decision cost function at time k . $L(k; \theta, \delta)$ denotes the penalty for deciding $\delta \in D(k)$ at time k when the true state of nature is θ . It is assumed to be bounded and nonnegative and have the structure:

$$L(k; \theta, \delta) = \begin{cases} L((i, \tau), \delta), & \tau \leq k, \delta \in D(k) \\ L_F, & \tau > k, \delta \in D(k) \end{cases}$$

$$L(k; (0, -), \delta) = L_F$$

where $L((i, \tau), \delta)$ is the underlying cost function for deciding δ when failure (i, τ) has already occurred. Also, L_F denotes the penalty for a false alarm (note that a false alarm corresponds to making a failure declaration before one occurs), and it can be generalized by allowing it to be a function of δ .

The cost function $L((i, \tau), \delta)$ generally has some additional structure. For example, a terminal decision cost that indicates the correct failure (and/or onset time) should receive a lower cost than one with the wrong failure (and/or onset time) indication. We further assume that the penalty due to an incorrect identification of the failure time is only dependent on the error in such an identification. That is for $\delta = (j, t)$,

$$L(k; (i, \tau), (j, t)) = L(i, j, (t - \tau)).$$

Note that $L(i, i, (t - \tau))$ corresponds to the penalty for an incorrect time estimate of when the failure type is correctly determined. Again the use and importance of this cost depends upon the application. Finally, if onset time is unimportant, so that δ does not obtain a time specification, we have

$$L((i, \tau), \delta) = L(i, \delta).$$

5) $r(k)$: the m -dimensional residual (observation) sequence. We shall let $p(r(1), \dots, r(k) | i, \tau)$ denote their joint conditional density when (i, τ) is true. Since the residual is affected by the failure in a causal manner, its conditional density has the property

$$p(r(1), \dots, r(k) | i, \tau) = p(r(1), \dots, r(k) | 0, -), \\ i = 1, \dots, M, \quad \tau > k.$$

In this paper, we will assume that the residual is an independent Gaussian sequence with V ($m \times m$ matrix)

as the time-independent covariance function and $g_i(k-\tau)$ as the mean given that the failure (i, τ) has occurred. With the covariance assumed to be the same for all failures, the mean function $g_i(k-\tau)$ characterizes the effect of the failure (i, τ) , and it is henceforth called the signatures of (i, τ) (with $g_i(k-\tau) = 0$ for $i = 0$, or $\tau \geq k$). We have chosen to study this type of residuals because its special structure facilitates the development of insights into the design of decision rules. Such a model arises in the case in which the residuals are generated by a Kalman filter based on normal operation and in which the failures enter additively in the system dynamics or sensor outputs [20]. While this model is not correct if parametric failures are considered (since in this case the correlation structure of the residuals is also affected by the failure), the general concepts we develop for the formulation of a BSDP for failure detection carry over to the parametric case. Furthermore, as reported in [16, 21], an FDI system based on an appropriate additive-failure model can often work very well in detecting parametric failures.

6) $c(k, (i, \tau))$: the delay cost function having the properties:

$$c(k, (i, \tau)) = \begin{cases} c(i, k-\tau) > 0, & \tau < k \\ 0, & \tau \geq k \end{cases}$$

$$c(i, k_1 - \tau) > c(i, k_2 - \tau), \quad k_1 > k_2 > \tau.$$

After a failure has occurred at time τ , there is a penalty for delaying the terminal decision until time $k \geq \tau$ with the penalty an increasing function of the delay $(k - \tau)$. In the absence of a failure, no penalty is imposed on residual sampling. In this paper we will consider a delay cost function that is linear in the delay, i.e., $c(i, k - \tau) = c(i)(k - \tau)$, where $c(i)$ is a positive function of the failure type i , and may be used to provide different delay penalty for different types of failures.

A sequential decision rule naturally consists of two parts: a stopping rule (sampling plan) and a terminal decision rule. The stopping rule is essentially a detection rule as its purpose is to determine whether monitoring should be interrupted in order to identify a failure. The terminal decision rule then performs the subsequent identification. The stopping rule denoted by $\Phi = (\phi(0), \phi(1; r(1)), \dots, \phi(k; r(1), \dots, r(k)), \dots)$ is a sequence of functions of the observed residual samples, with $\phi(k; r(1), \dots, r(k)) = 1$ or 0. When $\phi(k; r(1), \dots, r(k)) = 1$ (0), residual monitoring or sampling is interrupted (continued) after the k th residual sample $r(k)$ is observed. Alternatively, the stopping rule may be defined by another sequence of functions $\Psi = (\psi(0), \psi(1; r(1)), \dots, \psi(k; r(1), \dots, r(k)), \dots)$, where $\psi(k; r(1), \dots, r(k)) = 1$ indicates that residual monitoring has not been interrupted up to and including time $(k - 1)$ but will be interrupted when residual samples $r(1), \dots, r(k)$ are observed [15]. The functions Φ and Ψ are related to each other in the following way:

$$\psi(k; r(1), \dots, r(k)) = \phi(k; r(1), \dots, r(k)) \times \prod_{s=0}^{k-1} [1 - \phi(r(1), \dots, r(s))], \quad k \geq 1$$

with $\psi(0) = \phi(0)$.

The terminal decision rule is a sequence of functions, $D = (d(0), d(1; r(1)), \dots, d(k; r(1), \dots, r(k)), \dots)$. The function $d(k; r(1), \dots, r(k))$ maps the residual samples $r(1), \dots, r(k)$ into the terminal decision set $D(k)$ and represents the decision rule used to arrive at a failure identification if sampling is interrupted at time k .

If (i, τ) is the true state of nature and if the sequential decision rule (Φ, D) is used, then the total expected cost, i.e., the expectation of the sum of the delay and terminal decision costs is

$$U[(i, \tau), (\Phi, D)] = \sum_{k=0}^{\infty} E_{i, \tau} \{ \psi(k; r(1), \dots, r(k)) [c(k, (i, \tau)) + L(k; (i, \tau), d(k; r(1), \dots, r(k)))] \}$$

where $E_{i, \tau}$ denotes the expectation given that (i, τ) is true. The Bayes sequential decision rule (BSDR) with respect to μ is defined to be the sequential decision rule (Φ^*, D^*) that minimizes the sequential Bayes risk $U_s(\Phi, D)$ which is given by

$$U_s(\Phi, D) = E \{ U[(i, \tau), (\Phi, D)] \}$$

$$= \sum_{i=1}^M \sum_{\tau=1}^{\infty} \mu(i, \tau) U[(i, \tau), (\Phi, D)].$$

Now we discuss an interpretation of the sequential Bayes risk for the FDI problem. Let us define the following notation:

$$P_F(\tau) = \sum_{k=1}^{\tau-1} E_{0, \tau} \{ \psi(k; r(1), \dots, r(k)) \},$$

$$D = \bigcup_{k=0}^{\infty} D(k),$$

$$\tilde{S}(k, \delta) = \{ \{r(1), \dots, r(k)\} : \psi(k; r(1), \dots, r(k)) = 1, d(k; r(1), \dots, r(k)) = \delta \}, \quad \delta \in D,$$

$$\Pr\{\tilde{S}(k, \delta) | i, \tau\} = \int_{\tilde{S}(k, \delta)} p(r(1), \dots, r(k) | i, \tau) \times dr(1), \dots, dr(k),$$

$$\bar{t}(i, \tau) = \sum_{k=\tau}^{\infty} [(k - \tau) (1 - P_F(\tau))^{-1} \times E_{i, \tau} \{ \psi(k; r(1), \dots, r(k)) \}],$$

$$P((i, \tau), \delta) = (1 - P_F(\tau))^{-1} \sum_{k=\tau}^{\infty} \Pr\{\tilde{S}(k, \delta) | i, \tau\},$$

where $P_F(\tau)$ is the probability of stopping to declare a failure before the failure occurs at τ , i.e., the probability of false alarm when a failure occurs at time τ or later. D is the set of terminal decisions for all times. $\tilde{S}(k, \delta)$ is the

region in the sample space of the first k residuals where the sequential rule (Φ, D) yields the terminal decision δ . Clearly, the $\hat{S}(k, \delta)$'s are disjoint sets with respect to both k and δ . The expressions $\bar{t}(i, \tau)$ and $P((i, \tau), \delta)$ are, respectively, the conditional expected delay and the conditional probability of declaring δ , given a type i failure has occurred at time τ and no false alarm has been signalled before this time. $P((i, \tau), \delta)$ is called the generalized cross-detection probability. Using these quantities the sequential Bayes risk can be written as

$$U_s(\Phi, D) = \sum_{i=1}^M \sum_{\tau=1}^{\infty} \mu(i, \tau) \{L_F P_F(\tau) + (1 - P_F(\tau)) [c(i) \bar{t}(i, \tau) + \sum_{\delta \in D} L((i, \tau), \delta) P((i, \tau), \delta)]\}. \quad (1)$$

Equation (1) indicates that the sequential Bayes risk is a weighted combination of the conditional false-alarm probability, expected delay to decision, and cross-detection probabilities, and the optimal sequential rule (Φ^*, D^*) minimizes such a combination. From this vantage point, the cost functions (L and c) and the prior distribution (μ) act as the weighting coefficients and hence serve as a basis for specifying the tradeoff relationships among the various performance issues. The advantage of this approach is that only the total expected cost instead of every individual performance issue needs to be considered explicitly in designing a sequential rule. The drawback, however, lies in the need to choose a set of appropriate cost functions (and the prior distribution) when the physical problem does not have a natural set, as it doesn't in general. In this case, the Bayes approach is most useful with the cost functions and the prior distribution considered as design parameters that may be adjusted to obtain an acceptable design.

The optimal terminal decision rule D^* can be easily shown to be a sequence of fixed-sample-size tests [15]. The determination of the optimal stopping rule Φ^* is a dynamic programming problem [22]. The immense storage and computation required makes Φ^* impossible to compute, and suboptimal rules must be used.

Despite the impractical nature of its solution, the BSDP provides a useful framework for designing suboptimal decision rules for the FDI problem because of its inherent characteristic of explicitly weighing the tradeoffs between detection speed and accuracy (in terms of its cost structure). A sequential decision rule specifies a set of sequential decision regions $\hat{S}(k, \delta)$, and the decision regions corresponding to the BSDR yield the minimum risk. From this vantage point, the design of a suboptimal rule can be viewed as the problem of choosing a set of decision regions that would yield a reasonably small risk. This is the essence of the approach to suboptimal rule design that we take in this paper and describe next.

III. DESIGN OF SUBOPTIMAL RULES

A. Suboptimal Rules Based on the BSDR

The Sliding Window Approximation. The immense computation associated with the BSDR is partly due to the increasing number of possible failure times that must be considered as time progresses. The remedy for this problem is the use of a sliding window to limit the number of failure onset times to be considered at each time. The assumption made under the sliding window approximation is that essentially all failures can be detected within W time steps after they have occurred, or that if a failure is not detected within this time it will not be detected in the future. Here, the window size W is a design parameter, and it should be chosen long enough so that detection and identification of failures are possible, but short enough so that implementation is feasible [22].

The sliding window rule (Φ^W, D^W) divides the sample space of the sliding window of residuals $r(k - W + 1), \dots, r(k)$, or equivalently, the space of vectors of posterior probabilities, likelihood ratios, or log-likelihood ratios of the sliding window of failure hypotheses into disjoint time-independent *sequential decision regions* S_0, S_1, \dots, S_N . Here, $N = M$ if no failure time indication is involved in the terminal decision, while $N = MW$ if a failure time estimate is also required. Because the residuals are assumed to be Gaussian variables with variances that do not depend on the hypothesis, it is easy to check that an equivalent set of sufficient statistics is given by [20, 23]

$$\Lambda(k) = [\Lambda'_0(k), \dots, \Lambda'_{W-1}(k)]'$$

where, for $0 \leq \sigma \leq W - 1$,

$$\Lambda_\sigma(k) = [\Lambda(k; 1, \sigma), \dots, \Lambda(k; M, \sigma)]'$$

$$\Lambda(k; i, \sigma) = \sum_{s=0}^{\sigma} g'_i(s) V^{-1} r(k - \sigma + s).$$

Here σ indexes the possible failure onset times measured relative to the present time k (i.e., σ corresponds to a failure onset at time $k - \sigma$). The quantities $\Lambda(k; i, \sigma)$ differ only by an unimportant constant from the log-likelihood ratios for each hypothesis versus the no-fail hypothesis. The sliding window decision procedure operates as follows. At each time $k \geq W$, we form the decision statistics $\Lambda(k)$ from the window of residual samples. If $\Lambda(k) \in S_i$, for $i = 1, \dots, N$, we stop sampling to declare δ_i ; otherwise, $\Lambda(k) \in S_0$ and we proceed without making any immediate decision. The Bayes design problem is to determine a set of regions $S_0^*, S_1^*, \dots, S_N^*$ that minimizes the corresponding sequential risk $U_s^W(\{S_i\})$ (the expression for which we will describe shortly). This represents a functional minimization problem that is generally very difficult to solve. A simplification of this problem is to constrain the decision regions to take on special shapes $S_i(f)$ that are parameterized by a fixed dimensional vector f of design variables. A typical choice for these parametrically

specified regions might be in terms of the relative ordering of the sizes of the $L(k; i, \tau)$ and a set of threshold levels which correspond to the components of f (see (3) below). While such a constrained structure will lead to a suboptimal solution, the difference between the performance resulting from using the best constrained solution and that achieved by the optimal will be small if the constrained structure is chosen carefully. Furthermore, it is our contention that this performance difference will typically be mostly an artifact of the idealized problem formulation rather than a reality. That is, the unconstrained problem seeks to find the best boundaries between decision regions, while the constrained problem fixes the boundary shapes (e.g., straight lines of polygonal boundaries). Given that the residual statistical model used to define the problem is subjected to error, the extra drop of performance resulting from being able to "fine tune" the boundary shapes will generally be dwarfed in the uncertainty arising from modeling errors.

In the remainder of this paper we focus our attention on a special set of parameterized sequential decision regions, because they are simple and they serve well to illustrate that the Bayes formulation can be exploited, in a systematic fashion, to obtain simple suboptimal rules that are capable of delivering good performance. These decision regions are:

$$\begin{aligned} S(j, t) &= \{ \Lambda(k); \Lambda(k; j, t) > f(j, t), \\ &\quad \epsilon^{-1}(j, t) [\Lambda(k; j, t) - f(j, t)] \\ &\quad > \epsilon^{-1}(i, s) [\Lambda(k; i, s) - f(i, s)], \\ &\quad (i, s) \neq (j, t) \} \end{aligned} \quad (3a)$$

$$\begin{aligned} S(0, -) &= \{ \Lambda(k); \Lambda(k; i, s) < f(i, s), \\ &\quad i = 1, \dots, M, s = 0, \dots, W - 1 \} \end{aligned} \quad (3b)$$

where $S(j, t)$ is the stop-to-declare- $(j, k - t)$ region and $S(0, -)$ is the continue region. See Fig. 1 for a pictorial representation of the structure of (3) in the case where




Fig. 1. Sequential decision regions in two dimensions.

there are only two failures hypothesized, failure $(j, k - t)$ and $(i, k - s)$. Generally, the ϵ 's may be regarded as design parameters, but here, $\epsilon(j, t)$ is simply taken to be the standard deviation of $\Lambda(k; j, t)$.

To evaluate $U_s^W(f)$, the Bayes risk due to the use of (3), we need to determine the set of probabilities, $\Pr[\Lambda(k) \in S(j, t), \Lambda(k-1) \in S(0, -), \dots, \Lambda(W) \in S(0, -) | i, \tau]$, $k \geq W, j = 1, \dots, M, t = 0, \dots, W - 1$, which, indeed, is the goal of many research efforts in so-

called level-crossing problems [24]. As it stands, each of the probabilities is an integral of a kMW -dimensional Gaussian density over the compound region $S(0, -) \times \dots \times S(0, -) \times S(j, t)$, which for large kMW becomes extremely unwieldy and difficult to evaluate. A variety of approximations and bounds [25-28] have been developed for the evaluation of quantities such as this. We have not investigated the utility of any of these for our problem but rather have developed a systematic approach which is particularly appropriate for the dynamic FDI problem and which greatly simplifies the required calculations.

As a first step in this process, we reduce the dimension of the decision statistic $\Lambda(k)$ from MW to M . Specifically, we will base our decision process solely on the values of the log-likelihood ratios for each of the M failures modes assuming an onset time precisely at $\sigma = W - 1$, i.e., the beginning of the window. Since we are not estimating failure time in this case, the terminal decision to be made is simply the identification of the failure modes. The rationale behind this simplification has several aspects. First, in many applications, such as the aircraft sensor FDI problem [16] and the detection of freeway incidents [21], where the failure time need not be explicitly identified, the failure time resolution power provided by the full window of decision statistics is not needed. Furthermore, even if failure onset time information is desired, resolution of this time within a block of length W may often be sufficient. If not, one can imagine a two-level decision-making structure in which one first determines the failure type (using the procedure to be described) and then estimates the onset time. Note that this overall system will have decidedly lower complicity than one based on simultaneous detection, identification, and onset time estimation. In assuming the utility of the approach just described one must make sure that the resulting decision algorithm does not have a significantly elevated probability of incorrectly identifying the failure type. That is, if a failure of type i occurs at a time before the end of the window and if a detection occurs, one would want the subsequent identification to also be i with high probability. Determining whether this is the case can be done completely in terms of the failure signatures [29]. We can expect good performance if cross-correlation among signatures for failures of the same type at different times are significantly higher than the cross-correlations of signatures corresponding to different failure types. We note that this is often the case in practice, and in fact an often used goal for the residual generation process is that of producing signatures which are orthogonal or which at least lie in trivially overlapping subspaces [1, 22].

A decision rule of the type just described consists of sequential decision regions that are similar to (3) but are only defined in terms of the M components $\Lambda(k; i, W - 1)$, $i = 1, \dots, M$:

$$\begin{aligned} \Lambda_{W-1}(k) &= [\Lambda(k; 1, W - 1), \Lambda(k; 2, W - 1), \dots, \\ &\quad \Lambda(k; M, W - 1)]' \end{aligned} \quad (4a)$$

$$S_j = \{\Lambda_{W-1}(k): \Lambda(k;j, W-1) > f_j, \\ \epsilon^{-1}(j, W-1) [\Lambda(k;j, W-1) - f_j] \\ > \epsilon^{-1}(i, W-1) [\Lambda(k;i, W-1) - f_i], \\ j \neq i\} \quad (4b)$$

$$S_0 = \{\Lambda_{W-1}(k): \Lambda(k;j, W-1) \leq f_j, \\ j = 1, \dots, M\} \quad (4c)$$

$$U_s^W(f) = L_F \sum_{i=1}^M \sum_{\tau=W+1}^{\infty} [\mu(i, \tau) \\ \times \sum_{k=W}^{\tau-1} \sum_{j=1}^M \Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1)|0, -\} \\ + \sum_{i=1}^M \sum_{\tau=1}^{\infty} \mu(i, \tau) \sum_{k=\max\{W, \tau\}}^{\infty} \sum_{j=1}^M [c(i)(k-\tau) \\ + L(i, j)] \Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1)|i, \tau\}]$$

where S_0 is the event defined below:

$$S_0(k) = \{\Lambda_{W-1}(k) \in S_0, \dots, \Lambda_{W-1}(W) \in S_0\}.$$

The first term in the expression for $U_s^W(f)$ represents the portion of the risk due to false alarms. The key expression here is $\Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1)|0, -\}$, which is the probability that no detections have been made before time k but that an identification for a type j failure is made at time k , given that no failure has occurred. The remaining portion of $U_s^W(f)$ represents that part of the risk corresponding to detection delay and the possibility of incorrect identification. Here the key quantity is the probability $\Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1)|i, \tau\}$, which is the probability that a detection is first made at time k and that the failure is identified as being of type j given that a type i failure occurred at time $\tau \leq k$. The calculation of these probabilities is specified by the following recursions:

$$\Pr\{\Lambda_{W-1}(k+1) \in S_j | S_0(k), i, \tau\} \\ = 1 \int_{S_0} p(\Lambda_{W-1}(k) | S_0(k-1), i, \tau) d\Lambda_{W-1}(k) \\ \times \int_{S_0} [p(\Lambda_{W-1}(k+1) | \Lambda_{W-1}(k), S_0(k-1), i, \tau) \\ \times p(\Lambda_{W-1}(k) | S_0(k-1), i, \tau)] d\Lambda_{W-1}(k), \quad k > W \quad (5)$$

$$\Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1) | i, \tau\} \\ = \Pr\{S_0(k-1) | 0, -\} \int_{S_j} p(\Lambda_{W-1}(k) | S_0(k-1), i, \tau) \\ \times d\Lambda_{W-1}(k), \quad j = 1, \dots, M \quad (6)$$

with

$$\Pr\{\Lambda_{W-1}(W) \in S_j | i, \tau\} \\ = \int_{S_j} p(\Lambda_{W-1}(W) | i, \tau) d\Lambda_{W-1}(W). \quad (7)$$

Note that in essence what we are calculating in (5)–(7) are several different level crossing probabilities, and as we have just shown, it is these calculations that are the central elements to be determined in evaluating the performance of an hypothesized detection rule. For M small, numerical integration of (5)–(7) becomes manageable assuming that the required integrands are available.

Unfortunately, the transition density, $p(\Lambda_{W-1}(k+1) | \Lambda_{W-1}(k), S_0(k-1), i, \tau)$, required in (5), is difficult to calculate, because $\Lambda_{W-1}(k)$ is not a Markov process. In order to facilitate the computation of these probabilities, we use an approximation for this transition density obtained by developing an approximate Markovian model for the evolution of $\Lambda_{W-1}(k)$. A simple, but quite useful approximation is an M -dimensional Gauss-Markov process $l(k)$ that is defined by

$$l(k+1) = A l(k) + \zeta(k+1) \quad (8a)$$

$$\text{cov}\{\zeta(k)\zeta'(t)\} = \Gamma \delta_{k,t} \quad (8b)$$

where A is an $M \times M$ constant matrix, and ζ is a white Gaussian sequence (with covariance equal to the $(M \times M)$ matrix Γ) uncorrelated with $l(k)$. The conditional mean of $\zeta(k)$ will be specified shortly. The reason for choosing this model is twofold. First, just as $\Lambda_{W-1}(k)$, $l(k)$ is Gaussian. Second, $l(k)$ is Markov so that its transition density can be readily determined. In order to have the evolution of $l(k)$ match that of $\Lambda_{W-1}(k)$ as closely as possible, we choose the matrices A and Γ and the conditional mean $E_{i,\tau}\{\zeta(k)\}$ of $\zeta(k)$ under the hypothesis (i, τ) so that

$$E_{i,\tau}\{l(k)\} = E_{i,\tau}\{\Lambda_{W-1}(k)\} \quad (9a)$$

$$E_{0,-}\{l(k)l'(k)\} = E_{0,-}\{\Lambda_{W-1}(k)\Lambda'_{W-1}(k)\} \quad (9b)$$

$$E_{0,-}\{l(k)l'(k+1)\} = E_{0,-}\{\Lambda_{W-1}(k)\Lambda'_{W-1}(k+1)\}. \quad (9c)$$

That is, we have matched the marginal density and the one-step cross-covariance of $l(k)$ to those of $\Lambda_{W-1}(k)$. A straightforward calculation shows that (8)–(10) uniquely specify

$$A = \Sigma_1' \Sigma_0^{-1} \quad (10a)$$

$$\Gamma = \Sigma_0 - \Sigma_1' \Sigma_0^{-1} \Sigma_1 \quad (10b)$$

$$E_{i,\tau}\{\zeta(k+1)\} = E_{i,\tau}\{\Lambda_{W-1}(k+1)\} \\ - A E_{i,\tau}\{\Lambda_{W-1}(k)\} \quad (10c)$$

where

$$\Sigma_0 = E_{0,-}\{\Lambda_{W-1}(k)\Lambda'_{W-1}(k)\} \\ = \sum_{t=0}^{W-1} G_t V^{-1} G_t'$$

$$\begin{aligned}\Sigma_1 &= E_{0,-}\{\Lambda_{W-1}(k)\Lambda'_{W-1}(k+1)\} \\ &= \sum_{t=0}^{W-2} G_{t+1} V^{-1} G'_t \\ E_{i,\tau}\{\Lambda_{W-1}(k)\} &= \begin{cases} 0, & \tau > k \\ \sum_{t=0}^{k-\tau} G_{t+k_0} V^{-1} g_i(t), & k_0 = k - W + 1 - \tau \leq 0 \\ \sum_{t=0}^{W-1} G_t V^{-1} g_i(t+k_0), & k_0 = k - W + 1 - \tau > 0 \end{cases} \\ G_t &= [g_1(t), \dots, g_M(t)]'. \end{aligned}$$

Clearly, Σ_0^{-1} exists if the failure signatures $[g'_i(0), \dots, g'_i(W-1)]$, $i = 1, \dots, M$, are linearly independent. This condition is equivalent to the statement that there is sufficient information in a window of length W to distinguish among all of the M possible failure modes, assuming that if one of these failures has occurred, it did so in the beginning of the window. A sufficient condition for A to be stable, i.e., the magnitude of all its eigenvalues are less than unity, and Γ be invertible is that either G_0 or G_{W-1} is of rank M . (See the appendix for a discussion of the necessary and sufficient conditions for the invertibility of Γ and the stability of A .)

As an alternative to the model specification just given it is possible to choose other Markov approximations for $\Lambda_{W-1}(k)$. For example, one could match the n -step cross-covariance ($1 < n < W$) instead of matching the one-step cross-covariance as in (10). The suitability of a criterion for choosing the matrices A and Γ , such as (9) and (10), depends directly on the failure signatures under consideration and may be examined as an issue separate from the decision rule design problem. Also, a higher order Markov process may be used to approximate $\Lambda_{W-1}(k)$. However, the increase in the computational complexity may negate the benefits of the improved approximation. Finally, we emphasize that the statistic $l(k)$, as we have described it here, is *not* an observable quantity. That is, it cannot be computed from the residuals. Rather, $l(k)$ is an artificial process introduced in order to obtain approximations for the calculation of the statistics of $\Lambda_{W-1}(k)$. Later in this section we will describe a suboptimal test statistic to replace $\Lambda_{W-1}(k)$ which is computable from the residuals and which is also Markov.

Using the model we have developed for $l(k)$ we can approximate the required probabilities by substituting $l(k)$ for $\Lambda_{W-1}(k)$ in the calculations. That is,

$$\begin{aligned}\Pr\{\Lambda_{W-1}(k) \in S_j, S_0(k-1)|i, \tau\} \\ \approx \Pr\{l(k) \in S_j, S_0(k-1)|i, \tau\}, \\ j = 0, 1, \dots, M, k > W\end{aligned}$$

and

$$\begin{aligned}\Pr\{l(k) \in S_j, S_0(k-1)|i, \tau\} &= \Pr\{S_0(k-1)|i, \tau\} \\ &\times \int_{S_j} p(l(k)|S_0(k-1), i, \tau) dl(k). \quad (11)\end{aligned}$$

Assuming Γ^{-1} exists, we have

$$\begin{aligned}p(l(k+1)|S_0(k), i, \tau) \\ = [\int_{S_0} p(l(k)|S_0(k-1), i, \tau) dl(k)]^{-1} \\ \times \int_{S_0} [p(\zeta(k+1) = l(k+1) \\ - Al(k))|i, \tau) p(l(k)|S_0(k-1), i, \tau) dl(k)] \quad (12)\end{aligned}$$

where $p(\zeta(k)|i, \tau)$ is the Gaussian density of $\zeta(k)$ under the failure (i, τ) . The key simplification that results from using the Markovian approximation is

$$\begin{aligned}p(l(k+1)|l(k), S_0(k-1), i, \tau) &= p(l(k+1)|l(k), i, \tau) \\ &= p(\zeta(k+1) = l(k+1) \\ &\quad - Al(k)|i, \tau).\end{aligned}$$

Because of this, the integrands in (12) are readily obtained (the first comes from the previous step of the recursion) and thus the integrals in (12) can be calculated more easily.

In the event that Γ is not invertible, the density for $\zeta(k)$ is degenerate and (12) is more difficult to evaluate. As discussed in the appendix, the invertibility of Γ is related to the distinguishability of the M failure modes. Consequently, in any well-posed failure detection problem, W will be chosen so that the invertibility of Γ is assured.

Nonwindow Sequential Decision Rules. Here we describe another simple decision rule that has the same decision region as the simplified sliding window rule (4), but the vector z of M -dimensional statistics is obtained differently as follows:

$$z(k+1) = \tilde{A}z(k) + \tilde{B}r(k+1) \quad (13)$$

where \tilde{A} is a constant stable $M \times M$ matrix, and \tilde{B} is a $M \times m$ constant matrix of rank M . Unlike the Markov model $l(k)$ that approximates $\Lambda_{W-1}(k)$, $z(k)$ is a realizable Markov process driven by the residuals. The advantages of using $z(k)$ as the decision statistics are: 1) less storage is required, because residual samples need not be stored as necessary in the sliding window scheme, and 2) since

$z(k)$ is Markov, the required probability integrals are of the forms (11) and (12) so that the same integration algorithm can be directly applied. Of course, $z(k)$ is a suboptimal decision statistic. One could, if desired, use a higher-order model for $z(k)$ so that it more nearly equals $\Lambda_{w-1}(k)$, but the added computational complexity may negate the advantages.

In order to form the statistics $z(k)$, we need to choose the matrices \tilde{A} and \tilde{B} . When the failure signatures under consideration are constant biases and $M \leq m$, \tilde{B} can simply be set to equal $G_0 V^{-1}$ (provided G_0 is of rank M), and \tilde{A} can be chosen to be αI , where $0 < \alpha < 1$. Then, the term $\tilde{B}r$ in (13) provides the correlation of the residuals with the signatures as in (2), while the time constant α^{-1} characterizes the memory span of $z(k)$ just as W characterizes that of the sliding window statistics.

More generally, if we consider the case where failure signatures are not constant biases, rank $(G_0) < M$, or $m < M$, the choice of \tilde{A} may still be handled in the same way as in the previous case, but the selection of \tilde{B} is more involved. With some insights into the nature of the signatures, a reasonable choice of \tilde{B} can often be made in order to have distinct components of $z(k)$ respond primarily to the corresponding failure. To illustrate how this may be accomplished, we will consider an example with two failure modes ($M = 2$) and an m -dimensional residual vector. Let

$$g_1(k - \tau) = \beta_1$$

$$g_2(k - \tau) = \beta_2(k - \tau + 1).$$

That is, g_1 is a constant bias, and g_2 is a ramp. If β_1 and β_2 are orthogonal a simple choice of \tilde{B} is available:

$$\tilde{B} = \begin{bmatrix} \beta_1' \\ \beta_2' \end{bmatrix}.$$

This choice may often be acceptable even when $\beta_1' \beta_2 \neq 0$. It is clearly not of any use when β_1 and β_2 are multiples of the same vector β , or when they are scalars (corresponding to $m = 1$), as the rank of \tilde{B} is less than 2. In these cases we can consider processing groups of residuals. For example, suppose we batch process every two residual samples together, i.e., we use the augmented residual sequence $\tilde{r}(k) = [r'(2\tilde{k} - 1), r'(2\tilde{k})]'$, $\tilde{k} = 1, 2, \dots$. In this case we can set \tilde{B} to be

$$\tilde{B} = \begin{bmatrix} \beta' & \beta' \\ \beta' & 2\beta' \end{bmatrix}.$$

Thus, this \tilde{B} is of dimension $M \times 2m$ and has rank M ($= 2$). The first and second rows of \tilde{B} captures the constant bias and ramp nature g_1 and g_2 , respectively. The use of the modified residual $\tilde{r}(k)$ in this case causes no adverse effect, since it only lengthens slightly the interval between times when terminal decisions can be made. Clearly one can consider further augmentation and batch processing of the residuals, and in general the logical choice of \tilde{B} is one in which each row of \tilde{B} contain in sequence the initial values of the corresponding failure

signature. In this case the mean values of $z(k)$ will exactly equal that of $\Lambda_{w-1}(k)$ for a number of time steps following a failure equal to the level of augmentation used. The utility of this approach clearly depends on the temporal structure of the failure signatures. For problems where the signatures vary drastically as a function of the elapsed time and the distinguishability among failures depends essentially on these variations, the effectiveness of using $z(k)$ diminishes. In such cases the sliding window decision rule should provide better performance, although it should be noted that in this case one would typically have to use a comparatively long window in order to obtain an adequate degree of distinguishability.

B. Risk Evaluation

An algorithm based on one-dimensional Gaussian quadrature formulas [30, 31] has been developed to compute the probability integrals of (11) and (12) for the case $M = 2$. (It can be extended to higher dimension with an increase in computation.) The details of this quadrature algorithm are described in [22]. Its accuracy has been assessed via comparison with Monte Carlo simulations (see the numerical example in Section IV). With this algorithm we can evaluate the performance probabilities and risks associated with the suboptimal decision rules described above.

In the absence of a failure, the conditional density for $l(k)$ (12) has been observed in numerous examples to essentially reach a steady state¹ at some finite time $T > W$. Assuming this is the case, we have for $k \geq \tau \geq T$,

$$\Pr\{l(k) \in S_j | S_0(k-1), 0, -\} = \bar{b}_j \quad (14)$$

$$\Pr\{l(k) \in S_j, l(k-1) \in S_0, \dots, l(\tau)$$

$$\in S_0 | S_0(\tau-1), i, \tau\} = b_j(k - \tau | i). \quad (15)$$

That is, once steady state is reached, only the elapsed time since failure is important. Generally, failures occur infrequently, and decision rules with low false-alarm probability are employed. Thus, it is reasonable to assume 1) $\rho \ll 1$, i.e., $(1 - \rho)^{-1} \approx 1$, and 2) $\Pr\{S_0(T) | 0, -\} \approx 1$. The sequential risk associated with (4) for $M = 2$ can be approximated by

$$U_s^W(f) = P_F L_F + (1 - P_F) \sum_{i=1}^2 a(i) \sum_{j=1}^2 \sum_{t=0}^{\infty} [c(i) + L(i, j)] b_j(t | i) \quad (16)$$

where

$$P_F = \frac{(1 - \rho)(1 - \bar{b}_0)}{1 - \bar{b}_0(1 - \rho)}$$

P_F is the unconditional false-alarm probability, i.e., the probability of one false alarm over all time.

¹Unfortunately, we have not been able to prove such convergence behavior using elementary techniques. More advanced function-theoretic methods may be necessary.

Next, we seek to replace the infinite sum over t in (16) by the finite sum up to $t = Q$ plus a term approximating the remainder of the infinite sum. Suppose we have been sampling for Q steps since a failure occurred. Define

$$P_t(j|i) = \Pr\{I(t) \in S_j | S_0(t-1), i, 0\}, \quad j = 0, 1, 2.$$

If we stop computing the probabilities after Q , we may approximate

$$P_t(j|i) \approx P_Q(j|i), \quad j = 0, 1, 2, \quad t > Q. \quad (17)$$

That is we assume that after a detection delay of Q steps the conditional probability of detection at any time given no detection at any previous time reaches a constant steady-state value. This is the same as assuming that beyond Q steps of delay, the additional detection delay is exponentially distributed. This assumption is reasonable for constant failure signatures or signatures that reach steady state. While the assumption may not be valid for signatures which continue to vary, the effect of this approximation is generally quite small, since for all i and j one typically can choose Q so that

$$b_j(t|i) \approx 0, \quad t > Q.$$

That is, for each failure mode the probability of a detection delay greater than Q steps is negligible.

Substituting (17) in (16), we obtain

$$U_s^w(f) = P_F L_F + (1 - P_F) \sum_{i=1}^2 a(i) [c(i) \bar{t}_i + \sum_{j=1}^2 L(i, j) P(i, j)] \quad (18)$$

where

$$\bar{t}_i = \sum_{j=1}^2 \sum_{t=0}^Q t b_j(t|i) + Q b_0(Q|i) + \frac{1}{1 - P_Q(0|i)} \quad (19)$$

$$P(i, j) = \sum_{t=0}^Q b_j(t|i) + b_0(Q|i) \frac{P_Q(j|i)}{1 - P_Q(0|i)}. \quad (20)$$

Here, \bar{t}_i is the conditional expected delay to decision, given that a type i failure has occurred, and $P(j, i)$ is the conditional probability of declaring a type j failure, given that failure i has occurred. From the assumption that $\Pr\{S_0(T)|0, -\} \approx 1$ and the steady-state condition (14), it can be shown that the mean time between false alarms is simply $(1 - \bar{b}_0)^{-1}$. Now all the probabilities in (18)–(20) can be computed by using our quadrature algorithm. Note that the risk expression (18) consists only of finite sums and it can be evaluated with a reasonable amount of computational effort. With such an approximation of the sequential risk, we are able to consider the problem of determining the decision regions (i.e., the thresholds f_j^*) that minimize the risk.

It should be noted that we could consider choosing a set of thresholds that minimizes a weighted combination of certain detection probabilities ($P(i, j)$), the expected

delay (\bar{t}_i), and the mean time between false alarms $(1 - \bar{b}_0)^{-1}$. Although such an objective function will not result in a Bayesian design in general, it is a valid design criterion that may be useful for some applications.

C. Risk Minimization

The risk minimization has two features that deserve special attention. First, the sequential risk is not a simple function of the threshold f , and its derivatives with respect to f are not readily available. Second, calculating the risk is a computationally intensive task. Therefore, the minimum-seeking procedure to be used must require new function evaluations, and it must not require derivatives. For these reasons we chose to use the sequence-of-quadratic-programs (SQP) algorithm studied by Winfield [32] to solve this problem, because it does not need any derivative information and it appears to require fewer function evaluations than other well-known algorithms [32]. Furthermore, the SQP is simple, and it has quadratic convergence. Very briefly, the algorithm consists of the following. At each step of the iteration, a quadratic surface is fitted to the risk function locally using the preceding guesses at the optimal value of f and the corresponding risk function evaluations. The resulting quadratic model is minimized over a constrained region (hence the name SQP). The risk function is evaluated at this minimum and is used in the surface fitting of the next iteration. The details of the application of SQP to risk minimization are reported in [22].

IV. NUMERICAL EXAMPLE

Now we discuss an application of the suboptimal rule design methodology developed in this paper. We consider the detection of two possible failure modes (without identifying the failure time). The residual is a two-dimensional vector, and the vector failure signatures, $g_i(t)$, $i = 1, 2$, as functions of the elapse time t are shown in Table I. The signature of the first failure is simply a constant vector. The first component of $g_2(t)$ is a constant, while the second component is a ramp. We

TABLE I
Failure Signatures

$g_1(t) = \begin{bmatrix} 1 \\ 0.5 \end{bmatrix}, \quad g_2(t) = \begin{bmatrix} 0.5 \\ 0.25 + 0.25t \end{bmatrix}, \quad V = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$		
---	--	--

have chosen to examine these types of signatures because they are simple and describe a large variety of failure signatures that are commonly seen in practice. For simplicity, we have chosen V , the covariance of r , to be the identity matrix.

Both a simplified sliding window rule (that uses Λ_{w-1}) and a rule using the Markov statistic z were examined. The parameters associated with Λ_{w-1} , l , and z are shown in Table II, and the cost functions and the prior probability are shown in Table III. To facilitate

TABLE II
Parameters for Λ_{W-1} , l , and z

$\Sigma = \begin{bmatrix} 10 & 8.5 \\ 8.5 & 14.75 \end{bmatrix}$	$A = \begin{bmatrix} 0.826 & 0.058 \\ 0.116 & 0.837 \end{bmatrix}$	$\Gamma = \begin{bmatrix} 2.32 & 2.01 \\ 2.01 & 4.58 \end{bmatrix}$
$\bar{A} = \begin{bmatrix} 0.875 & 0 \\ 0 & 0.875 \end{bmatrix}$	$\bar{B} = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 2 \end{bmatrix}$	

TABLE III
Cost Functions and Prior Probability

$c_1 = c_2 = 1$
$L(1,2) = L(2,1) = 10, \quad L(1,1) = L(2,2) = 0, \quad L_F = 9$
$T = 8, \quad Q = 8, \quad \rho = 0.0002$
$\mu(i, \tau) = 0.5 \rho (1 - \rho)^{\tau-1}, \quad i = 1, 2$

discussion, we introduce the following terminology. We refer to a Monte Carlo simulation of the sliding window rule as SW, a simulation of the rule using the Markov statistic z as Markov Implementation (MI), and a simulation of the nonimplementable decision process




Fig. 2. $b_0(t|0)$: SW, MA, and Q20.




Fig. 3. $b_0(t|1)$: SW, MA, and Q20.




Fig. 4. $b_0(t|2)$: SW, MA, and Q20.




Fig. 5. $\sum_{s=0}^t b_1(s|1)$: SW, MA, and Q20.

using the approximation l as Markov Approximation (MA). (All simulations are based on 10 000 trajectories.) The notation Q20 refers to the results of applying the quadrature algorithm to calculate the various performance indices of the sliding window rule while using l to approximate Λ_{W-1} (12).

The results of SW, MA, and Q20 for the thresholds [8.85, 12.05] are shown in Figs. 2-6 (see (15) for the definition of notation). The quadrature results Q20 are very close to those obtained by Monte Carlo simulations for MA, indicating the excellent accuracy of the quadrature algorithm. In comparing SW with MA, it is evident that the Markov approximation slightly underestimates the false-alarm rate of the sliding window rule (SW). However, the response of the Markov approximation to failure is very close to that of the sliding window rule. In the present example, Λ_{W-1} is a seventh order process, while its approximation l is only first order. In view of this fact we can conclude that l provides a very reasonable and useful approximation of Λ_{W-1} .

The successive choices of thresholds by SQP for the sliding window rule are plotted in Fig. 7. Note that we have not carried the SQP algorithm so far that the successive choices of thresholds are, say, within 0.001 of




Fig. 6. $\sum_{s=0}^t b_2(s|2)$: SW, MA, and Q20.




Fig. 7. Thresholds chosen by SQP. Circles indicate initial guesses. Dots indicate thresholds determined by SQP.

each other. This is because near the optimum the expected risk is relatively insensitive to small changes in f . This implies that fine-scale optimization is not generally worthwhile. This conclusion is supported by the fact that the residual signature models used in designing failure detection systems are typically idealizations, and thus minor improvements in Bayes risk are generally an artifact of the mathematical formulation. Furthermore, it should be remembered that the use of the Bayes formulation is simply for the purpose of providing a




Fig. 8. $b_0(t|0)$: SW and MI.




Fig. 9. $b_0(t|1)$: SW and MI.




Fig. 10. $b_0(t|2)$: SW and MI.




Fig. 11. $\sum_{s=0}^t b_1(s|1)$: SW and MI.




Fig. 12. $\sum_{s=0}^t b_2(s|2)$: SW and MI.

mechanism for determining high-performance decision rules, and thus the precise optimization of the Bayes risk is not the central issue. In fact, the cost parameters L , c , μ , and W should be used as design parameters. In the event that the optimal thresholds resulting from a particular choice of Bayes risk do not provide the desired detection performance, the design parameters may be adjusted and the SQP may be repeated to get a new design. A practical alternative method is to make use of the list of performance indices (e.g., $P(i,j)$) that are generated in the risk calculation, and choose a pair of thresholds that yields the desired performance tradeoff.

The performance of the decision rule using Λ_{W-1} and z as determined by SQP are shown in Figs. 8–12. (The thresholds for Λ_{W-1} are [8.85, 12.05] and those for z are [6.28, 11.69].) We note that MI has a higher false-alarm rate than SW. The speeds of detection for the two rules are similar. While MI has a slightly higher type 1 correct detection probability ($\sum_{s=0}^t b_1(s|1)$) than SW, SW has a consistently higher type 2 correct detection probability ($\sum_{s=0}^t b_2(s|2)$) than MI. By raising the thresholds of the rule using z appropriately, we can decrease the false-alarm rate of MI down to that of SW with an increase in detection delay and slightly improved correct detection probability for the type 2 failure. Thus the sliding window rule is slightly superior to the rule using z in the sense that when both are designed to yield a comparable false-alarm rate, the latter will have longer detection delays and a slightly lower correct detection probability for a type 2 failure. In view of the fact that a decision rule using z is much simpler to implement, it is worthy of being considered as an alternative to the sliding window rule.

In summary, this example illustrates the utility of our approach. The quadrature algorithm has been shown to be accurate and useful, and the Markov approximation of Λ_{W-1} by l is a valid one. The simplicity and usefulness of the SQP algorithm has also been demonstrated. Finally, the Markov decision statistic z has been shown to be a worthy alternative to the sliding window statistic Λ_{W-1} .

V. CONCLUSION

A computationally feasible methodology based on the Bayesian approach has been developed for designing suboptimal sequential decision rules for FDI. This methodology was applied to a numerical example, and the results indicate that it is a potentially useful design approach.

APPENDIX

Theorem. Consider the Gauss-Markov process $l(k)$ specified by (8)–(10) in Section IIIA. Given $E\{\Lambda_{W-1}(k)\Lambda'_{W-1}(k)\} = \Sigma_0 > 0$, A has at least one unity eigenvalue and Γ is semipositive definite if and only if there exist M -vectors $\alpha \neq 0$ and $\beta \neq 0$ such that

$$\alpha' G_i = \beta' G_{i+1}, \quad i = 0, \dots, W-2 \quad (\text{A1}) \quad \alpha' \Lambda_{W-1}(k+1) = \beta' \Lambda_{W-1}(k). \quad (\text{A4})$$

$$\alpha' G_{W-1} = 0$$

$$(\text{A2}) \quad \text{Recall}$$

and

$$\beta' G_0 = 0.$$

Proof. Let

$$\begin{aligned} \tilde{\Sigma} &= E \left\{ \begin{bmatrix} \Lambda_{W-1}(k) \\ \Lambda_{W-1}(k+1) \end{bmatrix} [\Lambda'_{W-1}(k), \Lambda'_{W-1}(k+1)] \right\} \\ &= \begin{bmatrix} \Sigma_0 & \Sigma_1 \\ \Sigma_1' & \Sigma_0 \end{bmatrix}. \end{aligned}$$

Using the transformation T

$$T = \begin{bmatrix} I & 0 \\ -\Sigma_1' \Sigma_0^{-1} & I \end{bmatrix}$$

we obtain

$$T \tilde{\Sigma} T' = \begin{bmatrix} \Sigma_0 & 0 \\ 0 & \Gamma \end{bmatrix}.$$

Since T is full rank and $\Sigma_0 > 0$, $\tilde{\Sigma}$ and Γ are semipositive definite if and only if there are non-zero M -vectors α and β such that

$$(\text{A3}) \quad \Lambda_{W-1}(k) = \sum_{s=0}^{W-1} G_s V^{-1} r(k-W+1+s).$$

Therefore, (A4) is equivalent to the conditions (A1)–(A3). From (8) we obtain

$$\Sigma_0 = A \Sigma_0 A' + \Gamma.$$

It follows that A only has eigenvalues of magnitudes less than or equal to unity, and it has at least one unity eigenvalue if and only if Γ is semipositive definite. Q.E.D.

Suppose all signatures vanish for elapse times greater than $W-1$, i.e., $g_i(t) = 0$, for $t > W-1$, and $i = 1, \dots, M$. Then, (A1)–(A3) are equivalent to the condition that it is not possible to distinguish between a failure occurring at a certain time and failures occurring one time step earlier or later. Moreover, (A1)–(A3) indicate that only a special class of failure signatures would satisfy this indistinguishability condition for all values of W . Generally, it is possible to choose a sufficiently large W so that this situation is avoided.

- [1] Chow, E.Y., and Willsky, A.S. (1984)
Analytical redundancy and the design of robust failure detection systems.
IEEE Transactions on Automatic Control, vol. AC-29 (July 1984), 603–614.
- [2] Wald, A. (1947)
Sequential Analysis.
New York: Wiley, 1947.
- [3] Blackwell, D., and Girshick, M.A. (1951)
Theory of Games and Statistical Decisions.
New York: Wiley, 1951.
- [4] Chernoff, H. (1972)
Sequential analysis and optimal design.
In: SIAM Regional Conference Series in Applied Mathematics, vol. 8, Philadelphia, 1972.
- [5] Page, E.S. (1954)
Continuous inspection schemes.
Biometrika, 41 (1954), 100–114.
- [6] Gagliardi, R.M., and Reed, I.S. (1965)
On the sequential detection of emerging targets.
IEEE Transactions on Information Theory, IT-11 (Apr. 1965), 260–263.
- [7] Chien, T.T., and Adams, M.B. (1976)
A sequential failure detection technique and its application.
IEEE Transactions on Automatic Control, AC-21 (Oct. 1976), 750–757.
- [8] Basseville, M. (1981)
Edge detection using sequential methods for change in level—part II: sequential detection of change in mean.
IEEE Transactions on Acoustics, Speech, and Signal Processing, ASSP-29 (Feb. 1981), 32–50.
- [9] Shiryaev, A.N. (1963)
On optimum methods in quickest detection problems.
In *Theory of Probability and Applications*, Vol. 8, 1963, pp. 22–46.
- [10] Shiryaev, A.N. (1963)
On the detection of disorder in a manufacturing process I.
In *Theory of Probability and Applications*, Vol. 8, 1963, pp. 247–265.
- [11] Shiryaev, A.N. (1963)
On the detection of disorder in a manufacturing process II.
In *Theory of Probability and Applications*, Vol. 8, 1963, pp. 402–413.
- [12] Roberts, R.A., and Mullis, C.T. (1970)
A Bayes sequential test of M hypotheses.
IEEE Transactions on Information Theory, IT-16 (Jan. 1970), 91–94.
- [13] El-Sawy, A.H., and Vandelinde, V.D. (1979)
Robust sequential detection of signals in noise.
IEEE Transactions on Information Theory, IT-25 (May 1979), 346–353.
- [14] Fleisher, S., and Shwedyk, E. (1980)
A sequential multiple hypotheses test for the unknown parameters of a Gaussian distribution.
IEEE Transactions on Information Theory, IT-26 (Mar. 1980), 255–259.
- [15] Ferguson, T.S. (1967)
Mathematical Statistics, A Decision Theoretic Approach.
New York: Academic Press, 1967.
- [16] Deckert, J.C., Desai, M.N., Deyst, J.J., and Willsky, A.S. (1977)
F-8 DFBW sensor failure identification using analytic redundancy.
IEEE Transactions on Automatic Control, AC-22 (Oct. 1977), 795–803.
- [17] Gustafson, D.E., Willsky, A.S., Wang, J.Y., Lancaster, M.C., and Triebwasser, J.H. (1978)
ECG/VCG rhythm diagnosis using statistical signal analysis—II. Identification of transient rhythms.
IEEE Transactions on Biomedical Engineering, BME-25 (July, 1978), 353–361.
- [18] McAulay, R.J., and Denlinger, E. (1973)
A decision-directed adaptive tracker.
IEEE Transactions on Aerospace and Electronic Systems, AES-9 (Mar. 1973), 229–236.
- [19] Korn, J., Gully, S.W., and Willsky, A.S. (1982)
Application of the generalized likelihood ratio algorithm to maneuver detection and estimation.
In *Proceedings of the American Control Conference*, 1982.
- [20] Willsky, A.S., and Jones, H.L. (1976)
A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems.
IEEE Transactions on Automatic Control, AC-21 (Feb. 1976), 108–112.
- [21] Willsky, A.S., Chow, E.Y., Gershwin, S.B., Greene, C.S., Houpt, P.K., and Kurkjian, A.L. (1980)
Dynamic model-based techniques for the detection of incidents on freeways.
IEEE Transactions on Automatic Control, AC-25 (June 1980), 347–360.
- [22] Chow, E.Y. (1980)
A failure detection system design methodology.
Sc.D. Thesis, Laboratory for Information and Decision Systems, M.I.T., Cambridge, MA, 1980.
- [23] Van Trees, H.L. (1968)
Detection, Estimation, and Modulation Theory, Vol. I.
New York: Wiley, 1968.
- [24] Blake, I.F., and Lindsey, W.C. (1973)
Level-crossing problems for random processes.
IEEE Transactions on Information Theory, IT-19 (May 1973), 295–315.
- [25] Gallagher, R.G., and Helstrom, C.W. (1969)
A bound on the probability that a Gaussian process exceeds a given function.
IEEE Transactions on Information Theory, IT-15 (Jan. 1969), 163–166.
- [26] Bhati, D.H. (1973)
Approximations to the distribution of sample size for sequential tests, I. Tests of simple hypotheses.
Biometrika, 46 (1973), 130–138.
- [27] Ghosh, B.K. (1969)
Moments of the distribution of sample size in a SPRT.
American Statistical Association Journal, 64 (1969), 1560–1574.
- [28] Walker, B.K. (1980)
A semi-Markov approach to quantifying fault-tolerant system performance.
Sc.D. Thesis, Dept. of Aeronautics and Astronautics, M.I.T., Cambridge, MA, 1980.
- [29] Chow, E.Y. (1976)
Analytical studies of the generalized likelihood ratio technique for failure detection.
S.M. Thesis, Department of Electrical Engineering and Computer Science, M.I.T., Cambridge, MA, 1976.
- [30] Davis, P.J., and Rabinowitz, P. (1967)
Numerical Integration.
Waltham, MA: Blaisdell, 1967.
- [31] Stroud, A.H., and Secrest, D. (1966)
Gaussian Quadrature Formulas.
Englewood Cliffs, NJ: Prentice Hall, 1966.
- [32] Winfield, D.H. (1967)
Function minimization without derivatives by a sequence of quadratic programming problems.
Technical Report No. 537, Division of Engineering and Applied Physics, Harvard University, Cambridge, MA, Aug. 1967.



Edward Y. Chow (S'77—M'80) received the S.B., S.M., and Sc.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, Mass., in 1974, 1976, and 1981, respectively.

In 1976 he worked at the Charles Stark Draper Laboratory, Cambridge, Mass. Since September 1980 he has been with Schlumberger-Doll Research, Ridgefield, CT, working on the application of estimation and signal processing theory to well-logging problems.

Dr. Chow is a member of Eta Kappa Nu and Tau Beta Pi.



Alan S. Willsky (S'70—M'73—SM'82) received both the S.B. degree and the Ph.D. degree from the Massachusetts Institute of Technology in 1969 and 1973, respectively. From 1969 through 1973 he held a Fannie and John Hertz Foundation Fellowship. He joined the M.I.T. faculty in 1973 and his present position is Professor of Electrical Engineering. From 1974 to 1981 he served as Assistant Director of the M.I.T. Laboratory for Information and Decision Systems. He is also a founder and member of the board of directors of Alphatech, Inc. In 1975 he received the Donald P. Eckman Award from the American Automatic Control Council. He has held visiting positions at Imperial College, London, and L'Université de Paris-Sud. He is Editor of the M.I.T. Press series on signal processing, optimization, and control, was program chairman for the 17th IEEE Conference on Decision and Control, is an associate editor of several journals including the *IEEE Transactions on Automatic Control*, is a member of the Board of Governors of the IEEE Control Systems Society, and was program chairman for the 1981 Bilateral Seminar on Control Systems held in the People's Republic of China. Also he gave the opening plenary lecture at the 20th IEEE Conference on Decision and Control. He is the author of the research monograph *Digital Signal Processing and Control and Estimation Theory* and is coauthor of the undergraduate text *Signals and Systems*. He was awarded the 1979 Alfred Noble Prize by the ASCE and the 1980 Browder J. Thompson Memorial Prize Award by the IEEE for a paper excerpted from his monograph. His present research interests are in problems involving abrupt changes in signals and systems, multidimensional estimation, decision-directed signal processing, and the asymptotic analysis of control and estimation systems.

