Survey of Model-Based Failure Detection and Isolation in Complex Plants

Janos J. Gertler

ABSTRACT: This paper presents a survey of techniques to detect and isolate failures in complex technological systems, such as sensor biases, actuator malfunctions, leaks, and equipment deterioration. The methods are based on analytical redundancy afforded by a mathematical model of the system. The main components of such techniques are residual generation using the model, signature generation via statistical testing, and signature analysis. Model-structural conditions for failure isolation are introduced together with transformation methods to implement them. Sensitivity and robustness considerations are presented, and a design framework based on model redundancy is proposed.

Introduction

The detection and diagnosis of faults in complex process plants is one of the most important tasks assigned to the computers supervising such plants. The early indication of incipient failures can help avoid major plant breakdowns and catastrophies, ones that could otherwise result in substantial material damage and even human fatalities. Similarly, failure detection and isolation has become a critical issue in the operation of high-performance ships, submarines, airplanes, space vehicles, and structures, where safety, mission satisfaction, and significant material value are at stake. Quite recently, computerized diagnostic systems have been applied to such mass-produced consumer equipment as automobiles and household appliances. By assisting the human operator in assessing the nature and extent of the fault, automatic diagnostic systems may contribute significantly to the fast and proper reaction to failure situations, with such reactions ranging from immediate emergency actions to long-term modification of the maintenance schedule.

An early version of this paper was presented at the IFAC Symposium on Microcomputer Applications to Process Control, Istanbul, Turkey, July 22–25, 1986, and published in Vol. 7 of the *IFAC Proceedings Series*, 1987. Janos J. Gertler is with the School of Information Technology, George Mason University, Fairfax, VA 22030.

From the very beginning of computer process control, most practical systems have contained some form of failure detection and diagnosis. In the majority of these systems, the detection and diagnostic function is rather simple and is based on straight limit checking. The development of computational equipment and techniques has set the scene for the general application of more sophisticated and powerful methods.

In an article written shortly before the Space Shuttle disaster, Cikanek (1986) described the status of the failure detection system for ground test of the Shuttle main engine and outlined the general directions for future development. In the existing system, limit checking with fixed thresholds on single physical variables was the primary approach. A model-based method was used to check the operation of valve actuators, but the information represented by multiple interrelated plant measurements was otherwise not utilized. This paper calls for a system integrating a number of detection and isolation approaches, with a strong emphasis on model-based techniques.

The literature of process fault diagnosis is not very extensive, especially when compared to other areas of control engineering. The first major survey was written by Willsky (1976) followed by a more recent survey by Isermann (1984). Three significant books were published on the subject (Himmelblau, 1978; Pau, 1981; Basseville and Benveniste, 1986).

Nature of the Task

According to generally accepted terminology, failure detection and diagnostics consist of the following tasks (see also Fig. 1):

(1) *Failure detection*, i.e., the indication that something is going wrong in the system.

(2) *Failure isolation*, i.e., the determination of the exact location of the failure.

(3) *Failure identification*, i.e., the determination of the size of the failure.

While any relative weights attached to the preceding three components are obviously subjective, one may venture to say that detection is an absolute must in any practical system and isolation is almost equally important. Failure identification, on the other hand, while undoubtedly helpful, may not be worth the extra effort it requires. (As we will show later, much of the detection and isolation problem can be handled in a Boolean-logical framework. Failure identification, however, seeks numerical estimates and, therefore, does not render itself to such treatment.)

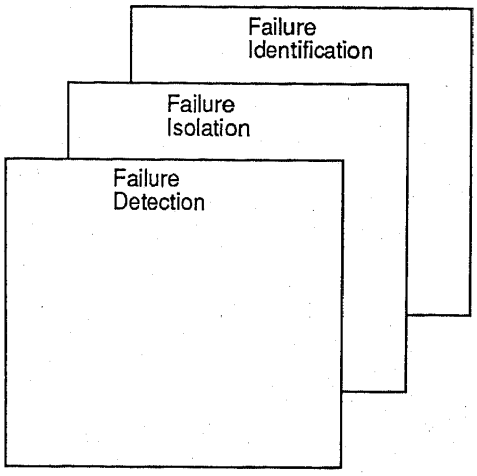

Fig. 1. Three layers of process diagnostics.

Note that failure detection and isolation are intimately related to the *filtering* of plant measurements: in certain schemes, filtering is used to facilitate isolation. At the same time, filtering performed to reduce noise effects and isolation are, in a sense, opposite activities: while isolation aims at pinpointing irregularities, filtering tends to hide (smear) them. In process engineering practice, it is customary to filter the measurements under constraints so that they satisfy energy and material balance equations or other plant models; this approach is often called *data reconciliation* (see, e.g., Gertler and Almasy, 1973; Mah et al., 1976; Gertler, 1979; Stanley and Mah, 1981). Reconciliation and isolation may use very similar techniques. However, isolation should precede reconciliation and the latter should be applied only after any isolated failure has been removed (at least computationally) from the system.

0272-1708/88/1200-0003 $01.00 © 1988 IEEE

Nature of Failures

We will use the words *failure* and *fault* as synonyms (although, strictly speaking, the term failure suggests complete breakdown, while a fault may connote something tolerable).

Almost any classification of failures (faults) reflects, explicitly or implicitly, the framework in which the detection/isolation problem is posed. Our classification corresponds to a model-based framework and is general enough to serve as a basis for further discussion. We will consider three classes of failures:

(1) *Additive measurement faults.* These are discrepancies between the measured and true values of plant output or input variables. Such faults best describe sensor biases. They can also be used to describe actuator malfunctions; now the discrepancy is between the intended (computed) control and its realization by the actuator.

(2) *Additive process faults.* These are disturbances (unmeasured inputs) acting on the plant, which are normally zero and cause a shift in the plant outputs independent of the measured inputs. Such faults best describe plant leaks, loads, etc.

(3) *Multiplicative process faults.* These are changes (abrupt or gradual) of the plant parameters. Such faults best describe the deterioration of plant equipment, such as partial or total loss of power, surface contamination, etc.

Additive faults (whether measurement or process) are distinguished from noise. Noise is considered random with zero mean. Additive faults are considered deterministic (e.g., a constant bias or drift) or semideterministic (jumps occurring at random intervals with random amplitudes).

Note that the preceding fault categories, although convenient from an analytical point of view, may not describe some practical failure situations in a natural way. A complete sensor failure (zero output) would have to be described either as a variable bias (equal to the true value) or as a multiplicative fault (some parameter changing to zero). Many disturbances are not completely additive: the intensity of a leak, for example, may depend on a pressure which is a process variable. An accurate description of this situation would require nonlinear modeling. Furthermore, an actuator failure might be more naturally described as an additive or multiplicative process fault, depending on the nature of the failure. However, the "additive measurement fault" representation of actuator failures leads to more uniform analytical treatment.

Approaches to Failure Detection and Isolation

The approaches to the problem of failure detection and isolation fall into two major groups:

- methods that do not make use of a plant model; and
- methods that do make use of a plant model.

The rest of this paper will be devoted to techniques that rely on a mathematical model of the plant. The model-free methods are only briefly reviewed.

(1) *Limit checking.* Plant measurements are compared to preset limits; exceeding a limit indicates a failure situation. In many systems, there are two levels of limits: the first level serves for prewarning only, while the second level triggers emergency action.

(2) Installation of *special sensors.* These may be limit sensors basically performing limit checking in hardware (e.g., limit temperature or pressure) or ones measuring some special variables (e.g., sound, vibration, elongation).

(3) Installation of *multiple sensors* (physical redundancy). This measure is aimed especially at detecting and isolating sensor failures. Measurements of the same variable from different sensors are compared. Any serious discrepancy is an indication of the failure of at least one sensor. The measurement that is likely to be correct may be selected in a voting system.

(4) *Frequency analysis* of plant measurements. Some plant measurements have a typical frequency spectrum under normal operating conditions; any deviation from this is an indication of abnormality. Certain types of failure may even have a characteristic signature in the spectrum that can be used for failure isolation.

(5) *Expert system* approach. The usual expert system approach is orthogonal to the preceding methods in that it is aimed at evaluating the symptoms obtained by the detection hardware or software. The system usually consists of a combination of logical rules of the [IF symptom AND symptom THEN conclusion] type, where each conclusion can, in turn, serve as a symptom in the next rule until the final conclusion (the specific failure) is reached. The expert system may work on the information presented to it by the detection hardware/software or may interact with a human operator, inquiring from him or her about particular symptoms and guiding him or her through the entire logical process (Hakami and Newborn, 1983; Kumamoto et al., 1984).

Model-Based Methods

A broad class of failure detection and isolation methods makes explicit use of a mathematical model of the plant. In the following, the general structure of such methods and of the models they use will be briefly introduced, and the problems of isolability, sensitivity, and robustness will be mentioned. A more detailed treatment of these subjects will be given in the succeeding sections.

General Structure of Model-Based Methods

Most model-based failure detection and isolation methods rely on the idea of *analytical redundancy* (Chow and Willsky, 1984). In contrast to physical redundancy, when measurements from different sensors are compared, now sensory measurements are compared to analytically obtained values of the respective variable. Such computations use present and/or previous measurements of other variables and the mathematical model describing their relationship. The idea can be extended to the comparison of analytically generated quantities only, each one obtained through a different computation. In either case, the resulting differences are called *residuals*.

While residuals are zero in ideal situations, in practice, this is seldom the case. Their deviation from zero is the combined result of noise and faults. If the noise is negligible, residuals can be analyzed directly. With any significant noise present, statistical analysis (statistical testing) is necessary. In either case, a logical pattern is generated, showing which residuals can be considered normal and which ones indicate fault. Such a pattern is called the *signature* of the failure. It should be noted that most failure detection and isolation methods do not utilize the information embodied in the size of the residuals beyond their relation to test thresholds.

The final step of the procedure is the analysis of the logical patterns obtained from the residuals, with the aim of isolating the failure or failures that cause them. Such analysis may be performed by comparison to a set of patterns (signatures) known to belong to

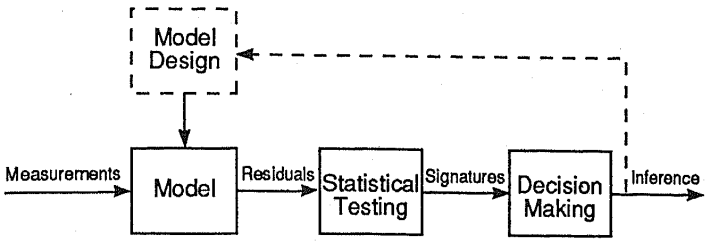

Fig. 2. Stages of model-based failure detection and isolation.

simple failures or by the use of some more complex logical procedure.

The three stages of the failure detection and isolation procedure—namely, residual generation, statistical testing, and logical analysis—are depicted in Fig. 2 (Chow and Willsky, 1984).

Note that failure detection and isolation procedures relying on system identification techniques can be considered as a special class of model-based methods. Rather than plant variables, the residuals now refer to plant parameters.

Plant and Failure Models

Most model-based failure detection and isolation methods rely on linear discrete-time models. This implies that any nonlinearity is linearized around some operating point. Also, continuous-time plants are represented by their discretized model.

Plant parameters may be varying with time. "Normal" variations are usually small and slow compared to the dynamics of the plant. Such variations will be neglected here for the sake of simplicity. Abrupt and/or significant changes, on the other hand, may and should be considered as multiplicative process faults.

The state-space model relates the state vector $\mathbf{x}(t)$ to the input vector $\mathbf{u}(t)$ and output vector $\mathbf{y}(t)$ (all functions of the discrete time t) using known system matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$. The well-known state equations are

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \quad (1)$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) \quad (2)$$

An equivalent input-output model may be presented in shift-operator form, with matrices $\mathbf{G}(z)$ and $\mathbf{H}(z)$ consisting of elements that are polynomials in the shift operator z and $\mathbf{H}$ being a diagonal matrix:

$$\mathbf{H}(z)\,\mathbf{y}(t) = \mathbf{G}(z)\,\mathbf{u}(t) \quad (3)$$

The matrices of the input-output model are related to those of the state model:

$$\mathbf{G}(z) = \mathbf{C}[\text{Adjoint } (\mathbf{I}z - \mathbf{A})]\mathbf{B}$$

$$\mathbf{H}(z) = [\text{Determinant } (\mathbf{I}z - \mathbf{A})]\mathbf{I} \quad (4)$$

Note that Eq. (3) usually can be simplified by eliminating, row by row, the common factors in the $\mathbf{H}(z)$ and $\mathbf{G}(z)$ matrices.

Introducing $\mathbf{v}(t)$ for additive process faults and $\mathbf{w}(t)$ for (additive) process noise, and $\mathbf{P}$ and $\mathbf{Q}$ for their coefficient matrices, the state equation [Eq. (1)] becomes

$$\mathbf{x}(t+1) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) + \mathbf{P}\mathbf{v}(t) + \mathbf{Q}\mathbf{w}(t) \quad (5)$$

The input-output equation [Eq. (3)] becomes

$$\mathbf{H}(z)\,\mathbf{y}(t) = \mathbf{G}(z)\,\mathbf{u}(t) + \mathbf{L}(z)\,\mathbf{v}(t) + \mathbf{M}(z)\,\mathbf{w}(t) \quad (6)$$

Here the matrices $\mathbf{L}$ and $\mathbf{M}$ are obtained in accordance with Eq. (4), with $\mathbf{B}$ replaced by $\mathbf{P}$ and $\mathbf{Q}$, respectively. Note that the presence of the new terms in Eq. (5) may influence $\mathbf{H}(z)$ and $\mathbf{G}(z)$ since $\mathbf{L}(z)$ and $\mathbf{M}(z)$ interfere with the simplification of the equations.

Introduce now $\Delta\mathbf{u}(t)$ and $\Delta\mathbf{y}(t)$ for the additive measurement fault (bias) on the input $\mathbf{u}(t)$ and output $\mathbf{y}(t)$, and $\delta\mathbf{u}(t)$ and $\delta\mathbf{y}(t)$ for the measurement noise on the same. With these, the measured input $\tilde{\mathbf{u}}(t)$ and output $\tilde{\mathbf{y}}(t)$ are

$$\tilde{\mathbf{u}}(t) = \mathbf{u}(t) + \Delta\mathbf{u}(t) + \delta\mathbf{u}(t)$$

$$\tilde{\mathbf{y}}(t) = \mathbf{y}(t) + \Delta\mathbf{y}(t) + \delta\mathbf{y}(t) \quad (7)$$

For controlled inputs, there is no sensory measurement; instead, $\mathbf{u}(t)$ is the control signal and $\mathbf{u}(t)$ is its implementation by the actuators, with $\Delta\mathbf{u}(t)$ representing any actuator malfunction and $\delta\mathbf{u}(t)$ the actuator noise.

Finally, introduce $\Delta\mathbf{A}(t)$, $\Delta\mathbf{B}(t)$, and $\Delta\mathbf{C}(t)$ for the discrepancies between the model matrices $\hat{\mathbf{A}}$, $\hat{\mathbf{B}}$, and $\hat{\mathbf{C}}$ and the true system matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$:

$$\hat{\mathbf{A}} = \mathbf{A} + \Delta\mathbf{A}(t)$$

$$\hat{\mathbf{B}} = \mathbf{B} + \Delta\mathbf{B}(t)$$

$$\hat{\mathbf{C}} = \mathbf{C} + \Delta\mathbf{C}(t) \quad (8)$$

Similarly, let $\Delta\mathbf{G}(z, t)$ and $\Delta\mathbf{H}(z, t)$ denote the discrepancies between the input-output model $\hat{\mathbf{G}}(z)$, $\hat{\mathbf{H}}(z)$ and the true system $\mathbf{G}(z)$, $\mathbf{H}(z)$:

$$\hat{\mathbf{G}}(z) = \mathbf{G}(z) + \Delta\mathbf{G}(z, t)$$

$$\hat{\mathbf{H}}(z) = \mathbf{H}(z) + \Delta\mathbf{H}(z, t) \quad (9)$$

Such discrepancies may account for multiplicative plant faults (e.g., some parameters in $\mathbf{A}$ changing). Strictly speaking, the true system matrices vary with time if there are multiplicative faults; however, this will not be shown in the notation.

To obtain a complete description of the system with all the possible faults and noises taken into account, the true variables $\mathbf{u}(t)$ and $\mathbf{y}(t)$ expressed from Eq. (7) and the true matrices $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$ from Eq. (8) or $\mathbf{G}$, $\mathbf{H}$ from Eq. (9) are to be substituted into Eqs. (5) and (2) or into Eq. (6). We defer this step to the next section, where the generation of residuals is discussed.

Note that the failure model in Eqs. (5)–(9) depicts the failures as general functions of time. In some earlier work, explicit failure models were used, including the time of the failure's occurrence (i.e., the time of a jump) as a parameter. These techniques imply significant computational complexity—e.g., a growing bank of filters. Several approximations have also been proposed utilizing a finite bank of filters with periodic resetting. For more on these methods, which have lost much of their popularity in the past decade, refer to Willsky's 1976 survey.

Isolability, Sensitivity, Robustness

Isolability, sensitivity, and robustness are quality properties of any failure isolation procedure that strongly influence the usefulness of such procedures.

Isolability is the ability of a procedure to distinguish (isolate) certain specific faults, provided their size is large enough. The isolability concepts depend on the statistical test selected and are linked to the structure of the system matrices.

Sensitivity is a more qualitative measure, characterizing the size of faults that can be isolated under certain conditions. Sensitivity depends on the size of the respective elements in the system matrices and on the noise properties. Closely related to sensitivity is the time it takes to detect/isolate a certain failure.

Robustness is the ability of a procedure to isolate faults in the presence of modeling errors. This is a serious problem since such errors are practically inevitable and may interfere seriously with fault isolation. In particular, they appear the same as multiplicative process faults and can only be distinguished from those based on the time history of the discrepancy.

Isolability, sensitivity, and robustness can, in most cases, be influenced by the careful selection and/or transformation of the plant model. Such model design would usually happen in the off-line design phase of the

procedure. It may, however, be done on-line, based on the performance of the procedure under certain conditions. In this sense, an on-line model redesign level can be added to the structure of the failure isolation procedure (Fig. 2).

Generating the Residuals

Residuals may be generated in a number of different ways. We will discuss subsequently the straightforward use of the input-output equation and of the state equation, followed by the application of the observer, Kalman filtering, and identification concepts. We will also indicate how faults and noise appear in the residuals obtained by some of the techniques listed.

Straight Input-Output Residuals

After a slight rearrangement, Eq. (3) is suitable for generating residuals. Introduce the combined input-output vector $\mathbf{q}(t)$ and the combined system matrix $\mathbf{F}(z)$:

$$\mathbf{q}(t)^T = [\mathbf{u}(t)^T, \mathbf{y}(t)^T]$$
$$\mathbf{F}(z) = [\mathbf{G}(z), -\mathbf{H}(z)] \qquad (10)$$

Now the ideal input-output equation can be written as

$$\mathbf{F}(z)\ \mathbf{q}(t) = \mathbf{0} \qquad (11)$$

Applying this equation to the measurements $\tilde{\mathbf{q}}(t)$ with the model matrix $\hat{\mathbf{F}}(z)$ yields a set of "parity equations," where the right-hand side is not zero, in general, but $\mathbf{e}(t)$, the vector of "residuals":

$$\hat{\mathbf{F}}(z)\ \tilde{\mathbf{q}}(t) = \mathbf{e}(t) \qquad (12)$$

A simplified version of this approach, known as *static balance equations*, has been popular in the chemical engineering literature. Originally, such equations contained a static system matrix $\mathbf{F}$ and the variables in $\mathbf{q}$ were restricted to material and energy flows (Vaclavek, 1974; Almasy and Sztano, 1975; Stanley and Mah, 1977; Romagnoli and Stephanopoulos, 1981). A first-order dynamic extension is proposed by Rooney et al. (1978). The full-blown dynamic case is discussed in Gertler and Singer (1985) and Gertler et al. (1985). A detailed analysis of the properties of residuals obtained from dynamic input-output equations is given also by Chow and Willsky (1984).

The residuals $\mathbf{e}(t)$ are the result of the combined noises, additive faults, and model discrepancies (multiplicative faults/modeling errors). According to Eqs. (6), (7), and (9), this relationship can be expressed as depicted in Fig. 3 and indicated in the following equation, where the first line contains the noise terms; the second line, the additive fault terms; and, the third line, the model discrepancy term.

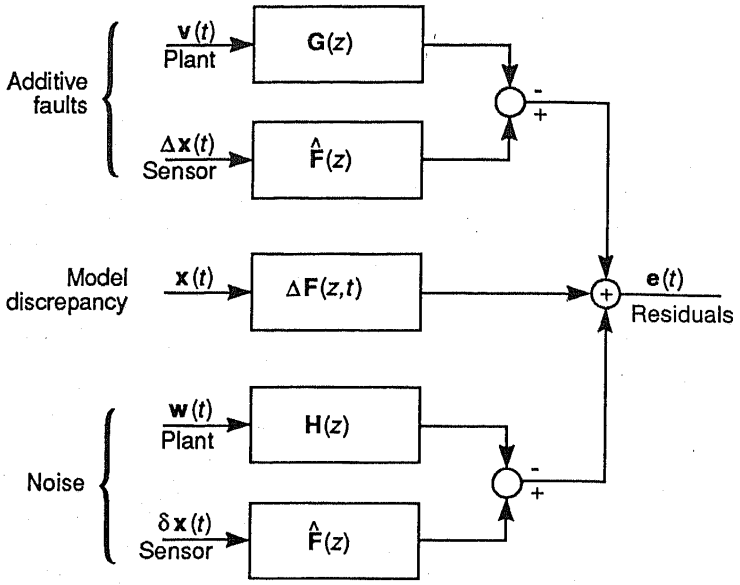

Fig. 3. Components of input-output residuals.

$$\begin{aligned}\mathbf{e}(t) &= \hat{\mathbf{F}}(z)\ \delta\mathbf{q}(t) - \mathbf{M}(z)\ \mathbf{w}(t) \\ &\quad + \hat{\mathbf{F}}(z)\ \Delta\mathbf{q}(t) - \mathbf{L}(z)\ \mathbf{v}(t) \\ &\quad + \Delta\mathbf{F}(z)\ \mathbf{q}(t) \qquad (13)\end{aligned}$$

State-Related Residuals

If the state variables are directly measurable or if they are computable from the output measurements (matrix $\mathbf{C}$ is invertible) then the residuals can be expressed directly in terms of the state variables. The nominal state can be computed by Kalman filtering or by an observer.

Matrix $\mathbf{C}$ usually is not invertible and, thus, direct comparison of the state is not possible. A workable approach in such cases is to define the residuals in terms of the output. If the residual is defined as the difference between the measured output $\tilde{\mathbf{y}}(t)$ and an estimate obtained by Kalman filtering, it is called the *innovation*:

$$\mathbf{e}(t) = \tilde{\mathbf{y}}(t) - \mathbf{C}\hat{\mathbf{x}}(t|t-1) \qquad (14)$$

Here $\hat{\mathbf{x}}(t|t-1)$ is the Kalman estimate of the state. If the model is perfectly accurate and the noise is white with zero mean, then the innovation sequence of a fault-free system is also white with zero mean. This property can be utilized to construct a number of statistical tests, as described in the next section, and it is mostly responsible for the early popularity of the method (Willsky, 1976).

Another popular approach has been the use of parallel observers. If the state vector (or part of it) is estimated by two observers based on different (although possibly overlapping) sets of outputs, a residual can be defined as shown, where $\hat{\mathbf{x}}_1$ and $\hat{\mathbf{x}}_2$ are the two estimates:

$$\mathbf{e}(t) = \hat{\mathbf{x}}_1(t) - \hat{\mathbf{x}}_2(t) \qquad (15)$$

Since the observer errors are $\mathbf{e}_i(t) = \hat{\mathbf{x}}_i(t) - \mathbf{x}(t)$, $i = 1, 2$, the residuals are, in fact,

$$\mathbf{e}(t) = \mathbf{e}_1(t) - \mathbf{e}_2(t) \qquad (16)$$

The observers can be constructed so that their errors $\mathbf{e}_1$ and $\mathbf{e}_2$ depend on some noise and modeling errors in identical or at least similar ways, thus leading to a robust fault isolation scheme (Patton and Willcox, 1985; Wunnenberg et al., 1985). Similar schemes can be constructed using parallel Kalman filters (Tylee, 1983). The application of parallel observers to physically redundant sensors has been reported by Stuckenberg (1985). The combination of observers with identification techniques was proposed by Watanabe and Himmelblau (1983).

In distributed systems, such as flexible structures, the plant variables depend not only on time but also on the spatial position. To handle this situation, the variables may be decomposed using time- and space-dependent factors. With r denoting the spatial position, $\Psi_i(r)$ the space-dependent so-called eigenfunctions, and $v_i(t)$ the time-dependent modal coordinates, the outputs may be expressed as

$$y(r, t) = \sum_{i=1}^{\infty} \Psi_i(r)\ v_i(t) \qquad (17)$$

The modal coordinates may be estimated from a set of observations $y(r_j, t)$ belonging to discrete spatial points (modal filtering). Using the estimated modal coordinates and Eq. (17), estimates of the observations $y(r_j, t)$ are obtained that, compared with the actual

observations, produce the output residuals (Baruh and Choe, 1987).

Identification-Based Methods

In identification-based methods, a residual-like quantity is defined in relation to the plant parameters. The plant is identified in a fault-free reference situation, then repeatedly on-line. The results of the latter are compared to the reference values and a parameter error (residual) is formed.

In the works of Isermann and coworkers (Isermann, 1984 and 1985; Geiger, 1984), the underlying physical parameters of the plant are computed from the identification results, and the comparison is performed on these. Rault et al. (1984) compute recursively, the covariance matrix of the identified parameters and compare it to reference covariances. Kumamaru et al. (1985) work with parameter estimates obtained from repeated identification runs and compare the consecutive results to each other, without resorting to a reference value.

Identification appears as an auxiliary method in a number of approaches. In Watanabe and Himmelblau (1983), a least-squares identifier is used to estimate the multiplicative process faults, based on variable estimates obtained by an observer. In Halme and Selkainaho (1984), plant parameters are estimated for an adaptive Kalman filter that, in turn, generates innovation-type residuals.

Statistical Testing

Whichever of the approaches discussed in the previous section is used, the resulting residual vector is the combined effect of faults and noise (and of modeling errors, but these will be neglected at this point). If there is no noise, any nonzero residual is an indication of a fault condition so that the logical analysis of the fault situation can be performed directly on the residuals. However, the assumption of zero noise is not a realistic one in most cases; even the round-off errors in the residual computation may cause nonzero residuals. Therefore, in general, the effect of faults on the residuals has to be separated from that of noise. This is done by statistical testing, making use of the fact (or assumption) that the noise is random with zero mean while failures are deterministic or semideterministic (as explained in the Introduction).

Direct Parallel Testing

The most straightforward approach is the direct parallel testing of the residuals. Following each computation of the residuals, a separate test is applied to each element of the residual vector. Based on the outcomes of the individual tests, a Boolean signature vector $\epsilon(t)$ is formed so that $\epsilon_i(t) = 1$ if $e_i(t)$ fired the test, and $\epsilon_i(t) = 0$ if it did not.

Such a test is especially easy to apply if the residuals are generated from an input-output model. This is because, in this case, the residuals of the fault-free system are outputs of a discrete moving average (nonrecursive) system driven by the noise. This implies that the fault-free distribution of the residuals can be derived from the noise distribution. In particular, if the noise is normally distributed, it is enough to obtain the noise-free variances of the residuals.

Further simplification results if the noises are white and mutually independent. From Eq. (13), the fault-free residuals are

$$\mathbf{e}(t) = \hat{\mathbf{F}}(z)\ \delta\mathbf{q}(t) - \mathbf{M}(z)\ \mathbf{w}(t) \qquad (18)$$

If $\mathbf{S}_q$ and $\mathbf{S}_w$ are the covariances of noises $\delta\mathbf{q}$ and $\mathbf{w}$, respectively, and n is the order of the system, then the covariance matrix of the residuals, $\mathbf{S}_e$, is obtained as

$$\mathbf{S}_e = \sum_{i=0}^{n} [\hat{\mathbf{F}}_i \mathbf{S}_q \hat{\mathbf{F}}_i^T + \mathbf{M}_i \mathbf{S}_w \mathbf{M}_i^T] \qquad (19)$$

The diagonal elements of $\mathbf{S}_e$ are the individual residual variances. Based on these and on a selected level of confidence, the thresholds for each residual can be obtained using normal distribution. To avoid frequent false alarms, usually a high level of confidence (99 percent or higher) is selected.

If the noise covariances are not available, the residual variances may be estimated from observations of fault-free periods. In this case, t-distribution is to be used (Vaclavek, 1974), although, with a large sample size, this approaches normal distribution.

The parallel test is relatively easy to administer and yields a distinctive binary signature that may serve as a basis for failure isolation. A disadvantage of this approach is that it does not utilize the additional information represented by the off-diagonal elements of the covariance matrix $\mathbf{S}_e$.

Multivariate Testing

Obviously, the elements of the fault-free residual vector are not independent of each other, even if the noise is white and uncorrelated. To utilize this additional information, Newman and Perkins (1979) suggested a multivariate statistical test. In the ρ-dimensional space of the residuals $\mathbf{e}(t) = [e_1(t), \ldots, e_\rho(t)]^T$, any constant probability density $p(e_1, \ldots, e_\rho) = \text{const}$ "isotherm" is described by a closed hypersurface. Selecting a level of confidence implies choosing one such surface. If the point defined by $\mathbf{e}(t)$ is outside the limit surface, the system is declared faulty. Beyond the obvious difficulty of administering this test to higher-dimensional residual vectors, the main disadvantage of this approach is that it provides only a faulty/not faulty decision that does not facilitate the isolation of the failure.

Compound Scalar Testing

Another approach to testing the residuals is to introduce a single scalar statistic, e.g., as

$$\lambda(t) = \mathbf{e}^T(t)\ \mathbf{S}_e^{-1}\mathbf{e}(t) \qquad (20)$$

(Romagnoli and Stephanopoulos, 1981). If the noises are normally distributed and there are no faults, then the statistic λ follows the chi-square distribution with ρ degrees of freedom (ρ is the number of residuals).

This test is easy to administer, however, similar to the multidimensional test, it provides a single binary inference. Romagnoli and Stephanopoulos proposed a sequential procedure, whereas equations are systematically eliminated until the test indicates a fault-free system. An important component of their method is an algorithm to obtain the new λ following each elimination step recursively, without complete recomputation.

Sequential Likelihood Ratio Test

This test, which is rather popular in aerospace applications, compares the hypothesis H_1 of nonzero residual mean to the null hypothesis H_0 of zero mean. The decision is based on the likelihood ratio

$$\lambda_i(t) = \ln \frac{p\{e_i(0), \ldots, e_i(t)|H_1\}}{p\{e_i(0), \ldots, e_i(t)|H_0\}} \qquad (21)$$

Here e_i is the ith element of the residual vector. The numerator and the denominator, respectively, are the probabilities of the observed time series under the two hypotheses. If the residuals are independent and normally distributed, the logarithmic likelihood ratio is very easy to compute: it becomes the difference of two sums.

The sequential likelihood ratio test has been applied to Kalman filter innovations by Chien and Adams (1976), Willsky and Jones (1976), Bonivento and Tonielli (1984), and Uosaki (1985), and to residuals from input-output parity equations by Deckert et al. (1977).

In most methods, the likelihood ratio is compared to two limits (normal/faulty), with a gray range between them. When it lies in the gray range, there is no decision and further innovation values are taken. In this sense, the test is sequential.

One difficulty in applying this test is the need for a numerical value of the innovation

mean in H_1. This is either a preset nominal value (Chien and Adams, 1976) or is obtained by maximum likelihood estimation from the innovations (Willsky and Jones, 1976).

Another difficulty is the accumulation of normal-state information in the likelihood ratio before the actual situation changes from H_0 to H_1, resulting in a delayed detection of the change. Chien and Adams propose a scheme in which such accumulations are reset automatically to zero. Uosaki's solution is a backward evaluation of the likelihood ratio, starting with $e_i(t)$ and continuing backward until a decision is reached. Willsky and Jones include the time of the change as an explicit but unknown variable, making the problem formulation more precise but the algorithm much more complex.

The test limits are derived either from specified probabilities for false/missed alarms (Uosaki) or from the expected time between two false alarms and the expected time of detecting a failure.

Failure isolation may be enhanced by running a set of parallel tests, with different alternative hypotheses H_1, H_2, etc., based on the same innovations sequence (Bonivento and Tonielli).

The likelihood ratio test is performed in compound scalar form by Basseville et al. (1987) and Rougee et al. (1987) to detect parametric changes in the transfer function denominator, without knowing the numerators, in multivariable systems. A vector "instrumental statistic" is generated that characterizes the goodness-of-model fit. The test is applied to determine if this vector has zero mean; the mean is estimated for the alternative (nonzero) hypothesis. From the instrumental statistic vector, a single scalar statistic is formed and tested against the appropriate chi-square distribution. This test is statistically efficient but provides a single binary decision; therefore, it does not facilitate failure isolation. For this latter purpose, parallel models are used, each one derived under a specific failure hypothesis, and separate tests are applied to each of the models.

Bayesian Approach

Some failure detection and isolation schemes utilize an a priori probability distribution of the occurrence of a set of failures. Such a priori distribution may be obtained from the observation of an extended history of the plant or may be assumed as design parameters.

In the nonlinear filtering scheme developed by Loparo and coworkers (Loparo et al., 1986; Eckert et al., 1986), failure events are modeled as sudden shifts in parameters governed by a random jump process with an a priori known distribution. Complete system models are determined for each of the possible failure modes. A nonlinear detection filter calculates the conditional probability of the system operating in each of such failure modes. While no finite-dimensional exact solution to the problem is known, reasonable approximations are possible using a finite set of Kalman filters and the solution of a stochastic differential equation. The algorithm is rather complex computationally, however, in simulation studies, it exhibits impressive robustness properties.

Isolability Conditions

One of the main questions to be asked in connection with any failure detection and isolation method is whether it is capable of isolating specific failure types from each other. The answer lies primarily in the structure of the model used in residual generation and, to a lesser degree, in the statistical test applied. In the following section, the concepts of deterministic (zero-threshold) and statistical (high-threshold) isolability will be introduced, together with a transformation technique to bring the model to the desired structure. To simplify the discussion, only bias faults in an input-output model will be considered. The concepts can be readily extended to process faults in the same model structure. They apply to the state-variable formulations as well, but the details depend on the method of residual generation. It will be assumed that no modeling errors are present.

Incidence Matrices

The structure of a model matrix $\mathbf{F}$ is characterized by its incidence matrix

$$\Phi = \text{Inc}\ (\mathbf{F}) \tag{22}$$

The incidence matrix is a Boolean matrix with its elements defined as

$$\begin{aligned} \phi_{ij} &= 1, \quad \text{if} \quad f_{ij} \neq 0 \\ &= 0, \quad \text{if} \quad f_{ij} = 0 \end{aligned} \tag{23}$$

If $\mathbf{F}$ is the system matrix of an input-output model as defined in Eqs. (3) and (10), then $\phi_{ij} = 1$ means that a bias on the jth variable influences the ith residual and $\phi_{ij} = 0$ means that it does not.

Deterministic (Zero-Threshold) Isolability

Consider Eq. (13), where the residuals are described in an input-output setting. Assume that the measurement biases $\Delta\mathbf{q}(t)$ are the only possible failures and there is no noise. Thus,

$$\mathbf{e}(t) = \hat{\mathbf{F}}(z)\ \Delta\mathbf{q}(t) \tag{24}$$

In such a deterministic (noise-free) situation, the Boolean signature vector $\epsilon(t)$ is obtained simply as

$$\begin{aligned} \epsilon_i(t) &= 1, \quad \text{if} \quad e_i(t) \neq 0 \\ &= 0, \quad \text{if} \quad e_i(t) = 0 \end{aligned} \tag{25}$$

Obviously, the signature vector in response to a single fault $\Delta q_j \neq 0$ is

$$\epsilon(t|\Delta q_j) = \varphi_{\cdot j} \tag{26}$$

where $\varphi_{\cdot j}$ is the jth column of the incidence matrix φ. This implies that, for the detection of any single bias fault Δq_j, its column $\varphi_{\cdot j}$ must be nonzero; for the isolation of the same, $\varphi_{\cdot j}$ must be different from all other columns of Φ.

The preceding condition, which we call *deterministic* or *zero-threshold isolability*, is rather obvious and has been recognized by several authors (e.g., Ben-Haim, 1980; Gertler and Singer, 1985).

Statistical (High-Threshold) Isolability

Consider now Eq. (24), with measurement noise $\delta\mathbf{q}(t)$:

$$\mathbf{e}(t) = \hat{\mathbf{F}}(z)\ \Delta\mathbf{q}(t) + \hat{\mathbf{F}}(z)\ \delta\mathbf{q}(t) \tag{27}$$

Assume that a direct parallel test is applied to the residuals. The signature vector $\epsilon(t)$ is now obtained as

$$\begin{aligned} \epsilon_i(t) &= 1, \quad \text{if} \quad |e_i(t)| \geq \eta_i \\ &= 0, \quad \text{if} \quad |e_i(t)| < \eta_i \end{aligned} \tag{28}$$

Here η_i is the threshold for the ith residual. If a bias Δq_j is large enough, it will trigger the test on all the residuals it influences, i.e., Eq. (26) holds. This situation will be referred to as *full firing*. In general, there is more penalty on a false alarm than on missing a small fault. Therefore, the thresholds will be set relatively high. [In Satin and Gates (1978), for example, the thresholds are set so that missed detections are 20 times as likely as false alarms.] On the other hand, the thresholds are determined by the fault-free distribution of the residuals; therefore, the limit value of a fault Δq_j that triggers the test will, in general, be different for each residual. Thus, a fault of intermediate size may not trigger the test on all the residuals it influences. This phenomenon, called *partial firing*, results in a degraded signature $\epsilon(t|\Delta q_j)$, in that some of the $1 - s$ in $\varphi_{\cdot j}$ will be replaced by $0 - s$. Now, if such a degraded signature is identical with another column $\varphi_{\cdot l}$, then the fault Δq_j is misisolated as Δq_l.

A model structure that guarantees that partial firing does not lead to misisolation is called a *statistically* or *high-threshold isolable* structure (Gertler and Singer, 1985; Gertler et al., 1985). A sufficient condition for statistical isolability of single bias failures is that Φ has a *column-canonical* structure (i.e., each column has the same number of $0 - s$, each in a different configuration).

Barring other disturbing effects (such as modeling errors or extreme noise values), a fault in a statistically isolable structure

- Is either detected and properly isolated (if large enough); or
- It leads to a signature that does not belong to any known fault (if intermediate size); or
- It is completely missed (if too small).

Model Transformation

System models in their original (primary) form seldom satisfy the conditions of deterministic isolability, let alone statistical isolability. To attain the desired structure, a model transformation may be performed. A secondary model thus obtained may have the same number of equations (residuals) as the primary model, or fewer or more. In the latter case, some of the equations are linearly dependent on the others but their incidence structure is different.

Model transformation may be looked on as a reshuffling of zeros in the parity equations. The effect of some variables (faults) is eliminated from certain equations at the expense of introducing other variables. In the case of Eq. (3), this implies eliminating input variables and introducing more than one output per equation. The number of zeros per equation cannot be increased by this transformation.

Technically, the transformation requires the solution, for the transforming matrix $\mathbf{T}(z)$, of the matrix equation

$$\mathbf{F}^x(z) = \mathbf{T}(z)\,\mathbf{F}(z) \qquad (29)$$

The desired model $\mathbf{F}^x(z)$ is specified in terms of its incidence matrix Φ^x. Each zero in Φ^x determines a linear algebraic equation; the solution of these yields the elements of $\mathbf{T}(z)$ [the solution is separable according to the rows of $\mathbf{T}(z)$]. The number of elements in $\mathbf{T}(z)$ must be at least as much as the number of zeros in Φ^x.

As is easy to see, this model transformation is governed by combinatorial rules. For anything but the most trivial systems, an extremely large number (hundreds or thousands) of secondary models may be derived, each satisfying the isolability conditions.

Zeros and/or linear dependence (even if it concerns parts of rows/columns) in the primary model matrix reduce the number of feasible secondary models. In particular, full-column dependence indicates that isolability cannot be attained by analytical means; additional sensors are necessary.

The preceding transformation technique was introduced in Gertler and Singer (1985) and Gertler et al., (1985). Some aspects were first developed in Shutty, 1984.

A somewhat different approach to model generation/transformation was reported by Chow and Willsky (1984). In their paper, linearly independent sets of input-output parity equations are derived, based on the underlying state-space model and observability considerations. Further parity equations are then generated as linear combinations of elements from a basic set. A special case is discussed by Massoumnia and Vander Velde (1988), where each parity equation contains only one output and several inputs, or one input and several outputs. Ben-Haim (1983) applied a different generalization (transformation) to the state vector to tailor the signature structure of statevariable-based residuals to certain types of failure.

Sensitivity and Robustness

Because an isolable model structure guarantees isolation if a failure is large enough, it is also of interest what failure size is really necessary to trigger the tests. This question may be answered by sensitivity analysis. Sensitivity aspects also may play a role in the selection of the desired model; filtering may be applied to improve sensitivity within a certain model framework.

Closely related to sensitivity is the aspect of robustness of the failure detection and isolation algorithm in the face of modeling errors. Such errors are present almost inevitably and may interfere seriously with the isolation of failures.

Modeling errors appear as discrepancies between the true parameters of the plant and the parameters of the model (they may reflect real parameter discrepancies or an improper model structure). Such modeling errors are indistinguishable from certain multiplicative process faults, especially on the basis of momentary signatures. The temporal behavior of the signatures may provide some clue since changes may most likely be caused by failures, although a change of the operating point also may result in a modeling error in the linear approximation of a nonlinear system.

Modeling errors also may interfere with the isolation of additive failures since they contribute to the residuals and may falsify the failure signatures. With the extension of the isolability concept to certain modeling errors, the failure detection and isolation algorithm can be desensitized with respect to those errors. Robustness considerations also may be taken into account in the model selection/transformation procedure.

Measure of Sensitivity

In a statistical testing framework, a possible measure of sensitivity is the marginal value of a fault that triggers the test, if the momentary value of the noise is zero. (Strictly speaking, this is an inverse sensitivity measure.) If the residuals are tested in parallel, this marginal triggering value of a fault is different, in general, for each residual:

$$\tau_{ij} = \eta_i / f_{ij} \qquad (30)$$

Here τ_{ij} is the marginal value of the jth fault that triggers the test on the ith residual and η_i is the threshold for the ith residual. The threshold is related to the fault-free variance of the residual that depends on the noise and on the system matrices; see, e.g., Eq. (19). The parameter f_{ij} characterizes the influence of the jth fault on the ith residual; for an additive fault, it is constant. In a static system, f_{ij} is the respective element of the system matrix $\mathbf{F}$. In a dynamic input-output model [Eq. (12) or (13)], this may be the steady-state gain of the respective element of the matrix polynomial $\mathbf{F}(z)$:

$$f_{ij} = [f_{ij}(z)]_{z=1} \qquad (31)$$

It is not certain, however, that the maximum fault effect on the residual will occur in steady state: for example, f_{ij} may be of a differentiating nature (Sundar, 1985).

Even if we disregard the preceding uncertainties in the definition, the triggering values are difficult to work with since they constitute a matrix characteristic of failure sensitivity. On the other hand, the selection of any scalar measure of sensitivity implies (arbitrary) preferences. Thus, such sensitivity considerations do not lead to a straightforward framework for failure detection and isolation algorithm design.

Improved Sensitivity via Filtering

Failure sensitivity can be improved by filtering the residuals (Evans and Wilcox, 1970; Gertler and Singer, 1985). The idea is to reduce the effect of (higher-frequency) noise while maintaining the effect of (lower-frequency) faults. A simple first-order filter is

$$e_{Fi}(t) = a e_{Fi}(t-1) + (1-a)\,e_i(t) \qquad (32)$$

This can be designed so that the fault-free variance of $e_{Fi}(t)$ is much smaller than that of $e_i(t)$, although their mean is the same. In fact, the variance can be reduced to any desired value, at the expense of a delay in the detection/isolation.

This type of filter works well if the residual equations are static or, in the case of dynamic equations, if the residuals depend on the faults in a proportional fashion. If the relationship is of a differentiating nature, filtering may reduce failure sensitivity. Therefore, there may be a need for a number of parallel filters with different coefficients (Sundar, 1985). On-line adaptation of the filter, based on the momentary signature pattern, may be an even more promising approach. This may be achieved in a rule- and knowledge-based framework.

Enhanced Robustness via Algebraic Cancellation

Robustness of the failure detection and isolation algorithm can be improved by desensitizing residuals with respect to certain modeling errors. This may be achieved by explicit algebraic cancellation of some terms in the residual equations. This is basically the extension to modeling errors of the concept of isolability and model transformation.

Residuals obtained from input-output relationships cannot be desensitized, by algebraic cancellation, to uncertainties (errors) in the $\mathbf{F}(z)$ matrix, only to those of some underlying parameters (Gertler et al., 1985). Such underlying parameters may be those of a state-space model or of a physical system model. Since the "coefficients" of the uncertainties to be eliminated are process variables (measurements), the cancellation is dependent on the operating point, i.e., part of the model transformation may have to be performed on-line.

The cancellation concept has been applied to observer-based residuals as well. If the residuals are obtained as the difference between two observer outputs, the observers can be designed so that some uncertainties of the state model matrices are ineffective (Patton and Willcox, 1985). This method, too, is subject to limitations depending on the frequency (speed of change) of plant variables. Observers can be designed to exhibit exact zero sensitivity with respect to modeling errors in the state equation, provided that all plant inputs and some outputs are guaranteed to be fault-free (Watanabe and Himmelblau, 1982).

Reachable Measurement Intervals

The problem of modeling errors has been approached in a straightforward manner by Horak and coworkers (Horak and Goblirsch, 1986; Horak and Allison, 1987). Given the input variables and the uncertainty range of the parameters of the state-space model, the reachable intervals of the output variables are explicitly computed. The technique is based on Pontryagin's optimum principle. To ease the computational complexity of the algorithm, several approximations are employed. The concept of reachable intervals is easily extended to include plant and measurement noise. Once the reachable intervals have been obtained, they serve as dynamic thresholds for the momentary measurements. The method is attractive for failure detection in the presence of modeling errors but needs further development before it can effectively support failure isolation.

Design for Robustness

As pointed out earlier, usually a large number of different models can be generated by transformation, all of which describe the same physical system. These models have different structure and parameters and, of course, different sensitivity and robustness properties. This great multitude of models makes it possible, at least in principle, to select the one that is the best from the point of view of sensitivity and/or robustness.

There are two major reasons why this selection is rather difficult to realize. First, it is generally impossible to formulate a single measure describing the required sensitivity and robustness properties, especially since these requirements usually vary with time (with the failure/noise/modeling error situation). Second, even if the requirements are clearly formulated, the large number of possible models rules out an analytical solution.

Chow and Willsky (1984) gave a minimax formulation for the robust model design problem, seeking the best parity equations as the ones resulting in the smallest expected residual under the worst modeling error for a fault-free system. This formulation is conceptually clear but too complex for practical design. Approximate solutions in simple cases are, however, possible.

Lou et al. (1986) report an elegant solution for the case when model uncertainty is represented by a finite set of slightly different possible model variants. In the framework of a geometric interpretation of the problem, those parity equations are sought that are maximally "orthogonal" to some normalized average of the model variants. The solution implies the singular-value decomposition of a composite model matrix and provides the "most robust" parity equations, together with a scalar measure of their global robustness.

If the design is performed with isolability in mind, the structural constraints somewhat limit the model choice. Still, there is a combinatorial multitude of structurally suitable models, leaving ample room for robustness considerations. In this framework, specific robustness measures, defined as inverse sensitivity relative to modeling errors in certain underlying parameters, seem more appropriate than global measures (Gertler, 1986b). A set of models, each one best for one of these measures, can be found in a discrete search procedure. Such search algorithms are sufficiently effective for off-line model design. The resulting models are then used on-line, in parallel. While most failures may, under normal circumstances, be isolated based on a single model, the use of multiple models will significantly improve the overall robustness of the isolation procedure (Gertler and Singer, 1989).

Conclusion

The main features of model-based failure detection and isolation methods have been surveyed in this paper. Several techniques to generate residuals from plant measurements and to obtain failure signatures via the statistical testing of the residuals have been discussed. It has been pointed out that the major quality issues of failure detection and isolation algorithms are isolability, sensitivity, and robustness. Isolability is related primarily to the structure of the residual-generating model and can be achieved by appropriate model transformation. Sensitivity and robustness requirements may vary with the failure/noise/modeling error situation and are best taken into consideration in a design framework that makes use of the model redundancy afforded by transformation.

References

Almasy, G. A. and T. Sztano (1975). "Checking and Correction of Measurements on the Basis of Linear System Model," *Problems of Control and Information Theory*, Vol. 4, pp. 57-69.

Baruh, H. and K. Choe (1987). "Sensor-Failure Detection Method for Flexible Structures," *Journal of Guidance, Control, and Dynamics*, Vol. 10, pp. 474-482.

Basseville, M. and A. Benveniste (Ed.) (1986). *Detection of Abrupt Changes in Signals and Dynamic Systems*, LNCIS No. 77, Springer, Berlin.

Basseville, M., A. Benveniste, G. Moustakides, and A. Rougee (1987). "Detection and Diagnosis of Changes in the Eigenstructure of Nonstationary Multivariable Systems," *Automatica*, Vol. 23, pp. 479-489.

Ben-Haim, Y. (1980). "An Algorithm for Failure Location in a Complex Network," *Nuclear Science and Engineering*, Vol. 75, pp. 191-199.

Ben-Haim, Y. (1983). "Malfunction Isolation in Linear Stochastic Systems: Application to Nuclear Power Plants," *Nuclear Science and Engineering*, Vol. 85, pp. 156-166.

Bonivento, C. and A. Tonielli (1984). "A Detection-Estimation Multifilter Approach with Nuclear Application," *Proceedings of the 9th World Congress of IFAC*, Budapest, Hungary, pp. 1771–1776.

Chien, T. T. and M. B. Adams (1976). "A Sequential Failure Detection Technique and Its Application," *IEEE Transactions on Automatic Control*, Vol. AC-21, pp. 750–757.

Chow, E. Y. and A. S. Willsky (1984). "Analytical Redundancy and the Design of Robust Failure Detection Systems," *IEEE Transactions on Automatic Control*, Vol. AC-29, pp. 603–614.

Cikanek, H. A., III (1986). "Space Shuttle Main Engine Failure Detection," *IEEE Control Systems Magazine*, Vol. 6, pp. 13–18.

Deckert, J. C., M. N. Desai, J. J. Deyst, and A. S. Willsky (1977). "F-8 DFBW Sensor Failure Identification Using Analytic Redundancy," *IEEE Transactions on Automatic Control*, Vol. AC-22, pp. 795–803.

Eckert, S., K. A. Loparo, and F. Roth (1986). "An Application of Non-Linear Filtering to Instrument Failure Detection in a Pressurized Reactor," *Nuclear Technology*, to appear.

Evans, F. A. and J. C. Wilcox (1970). "Experimental Strapdown Redundant Sensor Inertial Navigation System," *Journal of Spacecraft and Rockets*, Vol. 7, pp. 1070–1074.

Geiger, G. (1984). "Fault Identification of a Motor-Pump System Using Parameter Estimation and Pattern Classification," *Proceedings of the 9th World Congress of IFAC*, Budapest, Hungary, pp. 1783–1788.

Gertler, J. and G. Almasy (1971 and 1973). "Balance Calculations through Dynamic System Modelling," *IFAC Symposium on Digital Simulation of Continuous Processes*, Gyor, Hungary; also *Automatica*, Vol. 9, pp. 79–85.

Gertler, J. (1979). "A Constrained Minimum Variance Input-Output Estimator for Linear Dynamic Systems," *Automatica*, Vol. 15, pp. 353–358.

Gertler, J. and D. Singer (1985). "Augmented Models for Statistical Fault Isolation in Complex Dynamic Systems," *Proceedings of the 1985 American Control Conference*, Boston, MA, pp. 317–322.

Gertler, J., D. Singer, and A. Sundar (1985). "A Robustified Linear Fault Isolation Technique for Complex Dynamic Systems," *Preprints of the 7th IFAC Symposium on Digital Computer Applications to Process Control*, Vienna, Austria, pp. 493–498.

Gertler, J. (1986a). "Failure Detection and Isolation in Complex Process Plants. A Survey," *Proceedings of the IFAC Symposium on Microcomputer Applications to Process Control*, Istanbul, Turkey, pp. 13–25.

Gertler, J. (1986b). "A Numerical-Structural Approach to Failure Detection and Isolation in Complex Plants," *Proceedings of the IEEE Conference on Decision and Control*, Athens, Greece, pp. 1576–1580.

Gertler, J. and D. Singer (1989). "A New Structural Framework for Parity Equation-Based Failure Detection and Isolation," *Automatica*, to appear.

Hakami, B. and J. Newborn (1983). "Expert Systems in Heavy Industry: An Application of ICLX in a British Steel Corporation Works," *ICL Technical Journal*, Nov. 1983, pp. 347–359.

Halme, A. and J. Selkainaho (1984). "Instrument Fault Detection Using an Adaptive Filtering Method," *Proceedings of the 9th World Congress of IFAC*, Budapest, Hungary, pp. 1765–1770.

Himmelblau, D. M. (1978). "Fault Detection and Diagnosis in Chemical and Petrochemical Processes," *Chemical Engineering Monograph 8*, Elsevier.

Horak, D. T. and D. M. Goblirsch (1986). "Outputs in Systems with Bounded Parameter Uncertainties: Application to Failure Detection," *Proceedings of the 1986 American Control Conference*, Seattle, WA, pp. 301–308.

Horak, D. T. and B. H. Allison (1987). "Experimental Implementation and Evaluation of the RMI Failure Detection Algorithm," *Proceedings of the 1987 American Control Conference*, Minneapolis, MN, pp. 1803–1810.

Isermann, R. (1984). "Process Fault Detection Based on Modelling and Estimation Methods," *Automatica*, Vol. 20, pp. 387–404.

Isermann, R. (1985). "Process Failure Diagnosis Based on Modelling and Identification Methods," *IFAC-IFIP-IMACS 7th Conference on Digital Computer Applications to Process Control*, Preprints, Vienna, Austria, pp. 49–58.

Kumamaru, H. et al. (1985). "Fault Detection of Dynamical Systems Based on a Recognition Approach to Model Discrimination," *Proceedings of the 7th IFAC/IFORS Symposium on Identification and Process Parameter Estimation*, York, UK, pp. 1625–1630.

Kumamoto, H. et al. (1984). "Application of Expert System Techniques to Fault Diagnosis," *The Chemical Engineering Journal*, Vol. 29, pp. 1–9.

Loparo, K. A., F. Roth, and S. Eckert (1986). "Nonlinear Filtering for Systems with Random Structure," *IEEE Transactions on Automatic Control*, Vol. AC-31, pp. 1064–1068.

Lou, X. C., A. S. Willsky, and G. C. Verghese (1986). "Optimally Robust Redundancy Relations for Failure Detection in Uncertain Systems," *Automatica*, Vol. 22, pp. 333–344.

Mah, R. S. et al. (1976). "Reconciliation and Rectification of Process Flow and Inventory Data," *Ind. Eng. Chem., Process Des.*, Vol. 15, pp. 175–183.

Massoumnia, M. A. and W. E. Vander Velde (1988). "Generating Parity Relations for Detecting and Identifying Control System Component Failures," *Journal of Guidance, Control, and Dynamics*, Vol. 11, pp. 60–65.

Newman, R. S. and J. D. Perkins (1979). "Some Tests of a Method for Fault Detection on Chemical Plants," Conference on the Profitable Use of Statistics in Chemical Engineering, Manchester, UK.

Patton, R. J. and S. W. Willcox (1985). "Comparison of Two Techniques of I.F.D. Based on a Non-Linear Stochastic Model of an Aircraft," *Proceedings of IFAC Symposium on Identification and System Parameter Estimation*, York, UK, pp. 711–717.

Pau, L. (1981). *Failure Diagnosis and Performance Monitoring*, Marcel Dekker.

Rault, A. et al. (1984). "Industrial Processes Fault Detection and Localization," *Proceedings of the 9th World Congress of IFAC*, Budapest, Hungary, pp. 1789–1794.

Romagnoli, J. A. and G. Stephanopoulos (1981). "Rectification of Process Measurement Data in the Presence of Gross Errors," *Chemical Engineering Science*, Vol. 36, pp. 1849–1863.

Rooney, T. B. et al. (1978). "Estimating Plant Flows and Inventories Using Material Balance Relations," IFAC World Congress, Helsinki, Finland.

Rougee, A., M. Basseville, A. Benveniste, and G. Moustakides (1987). "Optimum Robust Detection of Changes in the AR Part of a Multivariable ARMA Process," *IEEE Transactions on Automatic Control*, Vol. AC-32, pp. 1116–1120.

Satin, A. L. and R. L. Gates (1978). "Evaluation of Parity Equations for Gyro Failure Detection and Isolation," *Journal of Guidance and Control*, Vol. 1, pp. 14–20.

Shutty, J. (1984). Senior Project. Case Western Reserve University, Cleveland, OH.

Stanley, G. M. and R. S. Mah (1977). "Estimation of Flows and Temperatures in Process Networks," *AIChE Journal*, Vol. 23, pp. 642–650.

Stanley, G. M. and R. S. Mah (1981). "Observability and Redundancy in Process Data Estimation," *Chemical Engineering Science*, Vol. 36, pp. 259–272.

Stuckenberg, N. (1985). "Sensor Failure Detection in Flight Control Systems Using Deterministic Observers," *Proceedings of IFAC Symposium on Identification and System Parameter Estimation*, York, UK, pp. 705–710.

Sundar, A. (1985). "Process Fault Detection Using the Augmented System Model Approach," Master's thesis, Case Western Reserve University, Cleveland, OH.

Tylee, J. L. (1983). "On-line Failure Detection in Nuclear Power Plant Instrumentation," *IEEE Transactions on Automatic Control*, Vol. AC-28, pp. 406–415.

Uosaki, K. (1985). "Failure Detection Using Backward SPRT," *Proceedings of IFAC Symposium on Identification and System Parameter Estimation*, York, UK, pp. 1619–1624.

Vaclavek, V. (1974). "Gross Systematic Errors or Biases in the Balance Calculations," Papers of the Prague Institute of Technology, Prague, Czechoslovakia.

Watanabe, U. and D. M. Himmelblau (1982). "Instrument Fault Detection in Systems with Uncertainties," *International Journal of System Science*, Vol. 13, pp. 137–158.

Watanabe, U. and D. M. Himmelblau (1983). "Fault Diagnosis in Non-Linear Chemical Processes," *AIChE Journal*, Vol. 29, pp. 243–249.

Willsky, A. S. (1976). "A Survey of Design Methods for Failure Detection in Dynamic Systems," *Automatica*, Vol. 12, pp. 601–611.

Willsky, A. S. and H. L. Jones (1976). "A Generalized Likelihood Ratio Approach to the Detection and Estimation of Jumps in Linear Systems," *IEEE Transactions on Automatic Control*, Vol. AC-21, pp. 108–112.

Wunnenberg, J., R. N. Clark, and P. M. Frank (1985). "An Application of Instrument Fault Detection," *Proceedings of IFAC Symposium on Identification and System Parameter Estimation*, York, UK, pp. 699–704.

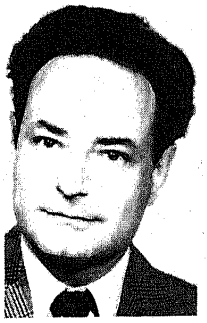

Janos J. Gertler was educated in Hungary. He received the Candidate of Science (Ph.D.) and Doctor of Science degrees, both in control engineering, from the Hungarian Academy of Sciences. For 10 years, he served as Deputy Director of the Institute for Computing and Automation of the Hungarian Academy of Sciences. He has held academic positions at Case Western Reserve University and at the Polytechnic University of New York. Presently, he is Professor of Electrical and Computer Engineering at George Mason University in Fairfax, Virginia. His research interests concern failure detection, intelligent control, and real-time control systems. He is Editor-in-Chief of the *IFAC Symposium Proceedings* and Associate Editor of the *IEEE Control Systems Magazine*.