TRANSCEND: A System for Robust Monitoring and Diagnosis of Complex Engineering Systems

Eric Manders*, Pieter Mosterman†,
Gautam Biswas‡, Lee Barford
Integrated Solutions Laboratory
HP Laboratories Palo Alto
HPL-1999-37
February, 1999

E-mail: [manders,biswas]@vuse.vanderbilt.edu
barford@hpl.hp.com
Pieter.J.Monsterman@dlr.de

diagnosis,
signal processing,
abrupt faults,
qualitative reasoning,
change detection, fault detection,
fault identification

This paper describes TRANSCEND, our system for monitoring and diagnosis of abrupt faults in complex dynamic systems. The key to this work has been our ability to model the transient behavior in response to faults in a qualitative framework, which enables us to overcome significant complexity and convergence issues that arise in numerical processing, especially for non linear systems. In our qualitative framework, the predicted future behavior of hypothesized faults is captured in the form of signatures, and analyzed by a progressive monitoring scheme. However, generating qualitative features from real signals, the signal to symbol transformation problem, is a challenging task. This paper discusses model-driven methods for generating symbolic feature descriptions of magnitude and slope changes in noisy, continuous data. To test our integrated framework for monitoring, prediction, and diagnosis, we have developed an automobile engine testbed. Sensors installed on the engine are connected to a PC workstation through a real-time data acquisition system. Experiments have been successfully conducted on faults introduced into the cooling system of the engine.

Internal Accession Date Only
*Department of Electrical and Computer Engineering, Vanderbilt University, Nashville, Tennessee
†Institute of Robotics and Systems Dynamics, DLR Research Center Oberpfaffenhofen, Wessling, Germany
‡Department of Computer Science, Vanderbilt University, Nashville, Tennessee
© Copyright Hewlett-Packard Company 1999

1 Introduction

The increasing complexity of engineering systems, and the need to optimize their operation while keeping costs low, has led to the use of *functional redundancy* techniques for control and fault detection and isolation (FDI) applications [22]. A brute-force strategy for control and diagnosis would be to explicitly measure all system variables of interest, but a more effective strategy is to derive required variable values using functional relations from a smaller set of measurements made on the system. Observed and derived discrepancies then drive the control and diagnosis tasks. In model based approaches, a parsimonious system model provides the functional relations between measured variables and system parameters.

Typically, faults that occur in physical systems can be classified as:

- *Incipient* faults, which are slowly changing component behaviors, related to wear and tear of components and drift in control parameters.
- *Abrupt* faults, which occur suddenly and are typically linked to situations where components may break or get stuck. In our work, abrupt faults will be described by sudden changes in component parameter values.

Our work has focused on analysis of abrupt faults in continuous dynamic systems. An important characteristic of these faults is that they exhibit transient behavior that takes the system away from steady state. After a while, the system may again reach a steady state, but in many situations control actions have to be taken before that to avoid damage, and possibly to prevent catastrophic failures to the system and its environment. Besides, a number of different failures may converge to the same steady state, therefore, information for differentiating among faults may be lost if the transients are not analyzed. Moreover, an abrupt fault may not be persistent. For these *intermittent* faults, the behavior of the system may not reach the steady state corresponding to the faulty situation.

To cope with these problems, our diagnosis approach uses models that capture the system dynamics, and, therefore, can be used to predict and analyze transient behaviors that occur in response to faults. To successfully apply these models for the diagnosis task, it is necessary to compare the predicted transient behavior with the actual transient behavior and distill a small set of viable fault hypotheses based on how well the predictions and actual transients match. The accuracy of the predicted transient depends on the accuracy of the model. When precise numerical information of system parameters is available, e.g., when the system is in normal operation, precise predictions can be made. When precise information is not available, such as when faults occur, the predictions are necessarily also less precise but they must still be accurate to achieve correct diagnosis results.

We adopt a qualitative reasoning approach to diagnosis to overcome some of the difficulties associated with quantitative techniques. Faults cause unknown changes in component parameter values that make it hard to predict system behavior quantitatively after a fault occurs. Well established state and parameter estimation techniques can be applied to analyze fault situations, but are hard to apply when system models are non linear [22]. Also, the difficulties in the estimation methods are further compounded because measured signals in real systems are typically noisy, and sensor response is a function of environmental conditions and characteristics that may drift over time.

To avoid the complexity problems and other pitfalls of numerical processing of the actual measurement signals by the diagnosis algorithms, we use our system models to inform the signal processing system of *symbolic features* that will be useful for fault isolation and refinement in particular situations. In this framework, we develop signal processing and interpretation methods that can extract the relevant symbols from noisy signals. In general, these symbols can represent discretized magnitude changes (above and below normal), slopes, and other characteristics of the signal that are relevant in the context of the model. This methodology for model-driven signal to symbol transformation may contain a toolbox of signal processing and analysis techniques. An overall goal for this project is to exploit model generated predictions to choose the best signal to symbol transformation method. For example, if it is known *a priori* that a signal may contain a discontinuity, a signal processing technique can be chosen to preserve this particular characteristic.

The emphasis of our previous work has been on the development of system models and diagnosis algorithms to achieve accuracy and precision in the fault isolation task. To test and evaluate this approach we have built simulation environments and lumped parameter models of multi-tank systems and the secondary sodium cooling loop of a fast breeder reactor [13, 14, 17]. This paper describes the signal to symbol processing module of our diagnosis system in detail. We demonstrate how signal processing techniques can be applied to find the presence of specific behaviors, and discuss the primary signal analysis and interpretation issues required to make the diagnosis methodology work with actual signals and systems in realistic environments. To this end we have constructed an automobile engine testbed, and have been running a number of data collection, data analysis, and fault isolation experiments on this testbed. In our current work the methods are applied to the diagnosis of faults in the cooling system of the internal combustion engine.

The organization of the rest of this paper is as follows. Section 2 reviews the overall monitoring and diagnosis framework. Section 3 explains in detail the feature extraction methodology required to perform the fault isolation task in a qualitative framework. Section 4 describes the setup of the testbed. Section 5 discusses an experiment where we introduce a fault and use the monitoring and diagnosis system to isolate the fault. The conclusions of our work are presented in Section 6.

2 A Framework for Model Based Monitoring and Diagnosis

We review here the central ideas of TRANSCEND (TRANSient based Continuous ENgineering systems Diagnosis), our model based system for monitoring and diagnosis of complex dynamic systems. Fig. 1 presents the different components of this system. The input to the physical process under diagnostic scrutiny is the input vector u, and the observations made on the system are indicated by vector y. We track the behavior of the process using an observer system, which consists of the observer model and an observer scheme. In general, the observer model is made up of a set of ODEs that define dynamic system behavior and generate the expected behavior of the system, $\hat{y}$, given u. The observer scheme tracks the residuals, $r = y - \hat{y}$, the difference between actual observations y and predicted observations, $\hat{y}$, and uses a standard gain matrix scheme to correct small discrepancies in the system state vector x [8].

The residuals, r, are also input to the signal analysis module, which uses sophisticated signal analysis techniques that we describe in Section 3 for discrepancy detection in the measured signals. In realistic situations, the measured signals are noisy, so the discrepancy detection combines noise attenuation methods with the symbol generation task. Once discrepancies are detected, the symbol generation unit reports the values of a set of symbols to the fault generation and fault refinement units. These symbols represent (i) discontinuous changes, (ii) magnitude deviations, (iii) slope deviations, and (iv) steady state behavior.

The diagnosis model captures qualitative dependency relations between component parameters and the observed variables. The fault generation unit uses this information, m and the symbolic residuals, r_s, to generate hypothesized faults from observed deviations, f_h, and to predict their future transients and their steady state behavior, p. A progressive monitoring scheme compares actual observations against the predicted transients for each hypothesized fault, $f \in f_h$. Only those faults whose transients agree with the observations are retained in the refined set f_r. The goal is to continue monitoring till the true fault is isolated. A brief overview of the modeling methodology and the diagnosis algorithms is presented below. A more detailed description of the modeling scheme and diagnosis methodology with examples can be found in [17, 19].

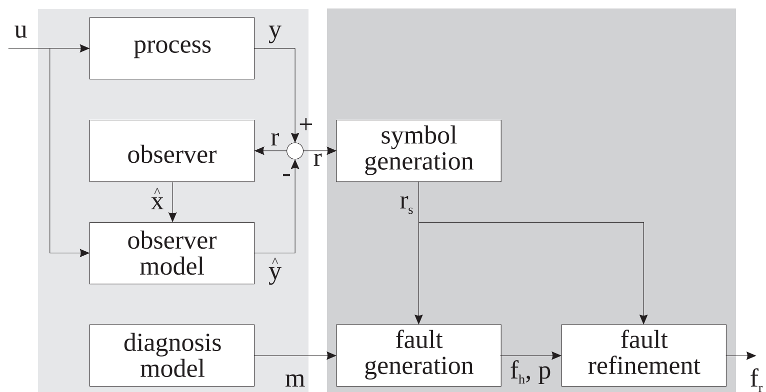

Figure 1: TRANSCEND system architecture.

2.1 Modeling Paradigm and Approach to Diagnosis

This section briefly overviews our diagnosis models and the qualitative framework we have developed for analyzing fault transients.

2.1.1 Models for Diagnosis

We have adopted the bond graph [24] methodology to system modeling. Bond graphs provide a systematic framework for building consistent and well constrained models of dynamic physical systems across multiple domains. Topological causality constraints, systematically derived from component descriptions, and structural connectivity expressed in bond graph models are exploited for effective and efficient diagnosis. The differential equation form of the observer model can be directly derived from the bond graph as a set of state equations.

The bond graph representation is also directly amenable to qualitative reasoning. We exploit this to develop our models for component oriented diagnosis by establishing a correspondence between individual system components and bond graph elements [3, 12, 15]. A fault is represented by a deviation of a component parameter in the bond graph model. In our qualitative framework this means a component parameter deviating above (+) or below (−) its nominal value. An individual system component may have multiple aspects represented in the bond graph. For example, a component such as a pipe may be represented by its build up of flow momentum (I) and resistance to flow (R). Thus a modeled component may fail in as many ways as there are associated component parameters.

Dynamic characteristics of system behavior are represented as a *temporal causal graph* (TCG), which may be algorithmically derived from the bond graph model [13, 17]. The TCG is a directed graph whose edges are in one to one correspondence with component parameters and qualitative algebraic relations based on component connectivity. Component parameter values and their temporal influences on system behavior are defined as attributes of the TCG edges. System behavior variables, defined in terms of the domain independent concepts of *effort* and *flow*, correspond to TCG vertices. In addition, there may be *signal* vertices that denote modulated variables and other constraints that exist between system variables. The TCG representation provides a rich but uniform framework for representing magnitude and temporal constraints among system variables. The hypothesis generation and prediction components of the fault isolation task are then developed as standard graph traversal algorithms.

2.1.2 Diagnosis from Transients

The diagnosis of abrupt faults relies heavily on characterization of transients. *Time constants* play a key role in defining transients. For some variables, a fault will cause an instantaneous change. For other variables, energy storage elements act as buffers that introduce propagation delays and changes take longer to manifest. If observations are available from the system at rates that are faster than the smallest time constant, it becomes easier to characterize transients. A change in signal value that takes place on a time scale smaller than the rate at which observations become available is labeled as an *abrupt change*. In the model this corresponds to the modeling of a discontinuity.

A measurement is considered *normal* in a qualitative sense if its value is within a certain percentage of the nominal value for that variable. Otherwise, it is considered to be *deviating*. Normal measurements cannot necessarily be used to refute faults during transient analysis [17]. A normal observation at a given time may actually be a slowly changing value that has not crossed the threshold to be labeled as deviating. The only situation when a normal measurement can be used to refute a fault during transient analysis is if that fault is predicted to cause an *abrupt change* in the measurement. This is contingent on the fact that abrupt changes can be detected reliably. This issue is discussed in greater detail in Section 3. A temporal ordering of the manifestation of first and higher order effects in studying deviations from normal is also generally impossible. In the current approach, deviating observations are individually analyzed to generate sets of single fault hypotheses.

Individual signal features are the prime discriminating information between competing fault hypotheses. In noisy environments, the accuracy of magnitude change measurements is governed by the properties of the associated sensors and the change detection techniques employed.

Reliable techniques can also be developed for qualitative estimates of slopes or first order derivatives, i.e., increasing (+), steady (0), and decreasing (−), from measured signals. However, it is increasingly difficult to estimate higher order derivatives in the presence of noise [5]. In our system we limit the measurement of derivatives to first order. To summarize, two features are needed to describe transients. Magnitude changes (+ or −) possibly augmented by specifying a discontinuous change and the first order derivative of the signal (+ or −).

2.2 Generating Fault Hypotheses and Predicting Future Behavior

Three algorithms make up the core of the fault isolation procedure: (i) fault hypothesis generation by component parameter implication, (ii) prediction by generation of transient signatures for each hypothesized fault, and (iii) fault isolation or hypothesis refinement by monitoring. Hypothesis generation and prediction are the topic of this section. The monitoring methodology is presented in the next section.

2.2.1 Component Parameter Implication

For every recorded discrepancy between measurement and nominal value, a backward propagation algorithm is invoked on the temporal causal graph to implicate component parameters. Observed deviating values are propagated backward along the directed edges and consistent '−' and '+' deviation labels are assigned sequentially to vertices along the path based on the edge relation. A component is implicated when an edge corresponding to that system component parameter is traversed. The component parameter deviation is then labeled '−' or '+', depending on the value assigned to the last vertex and the edge relation. This procedure generates a set of hypothesized single faults that are consistent with each reported deviating observation.

2.2.2 Signature Generation

We constrain the problem space here by making the assumption that faults do not cause changes in system configuration, and the system model remains valid even after faults occur in the system. The prediction algorithm employs the system model to compute the qualitative transient behavior of the observed variables under the individual fault conditions. Transient behavior is expressed as a tuple of qualitative values for magnitude, 1^{st} order time-derivative and higher order effects. The qualitative values are similar to those of the measured values: '+', '−', '0' or '·'. The '·' implies that the value is unknown, a result of opposite qualitative influences. The tuple is called the *signature* for the variable [14, 17].

The algorithm propagates the effects of a hypothesized fault to establish a signature for all observations. Energy storage elements cause time integrating effects and introduce temporal edges in which case the cause variable affects the derivative of the effect variable. Propagation of a deviation starts with a 0^{th} order effect, i.e., a magnitude change. When an integrating edge is traversed, the magnitude change becomes a 1^{st} order change, i.e., the first derivative of the affected quantity changes. Similarly, a first order change propagating across an integrating edge produces a second order change, and so on. The highest predicted derivative order required is a design consideration. For some faults, signatures may differ only for higher order derivatives. However, higher order effects also take longer to manifest, and thus the time to isolate the fault

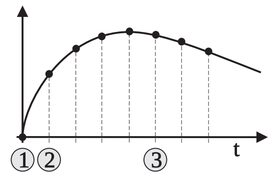

dt	0	1	2
①	0	+	-
②	+	+	-
③	+	-	-

Figure 2: Progressive monitoring

may be too long. This issue and the relation to the measurement selection problem is discussed in detail in [11].

2.3 Monitoring

The mechanism for monitoring has two aspects. First, the actual comparison of observations with the predicted behaviors to prune hypothesized faults. Second, the criteria for halting the transient monitoring based pruning of hypotheses.

2.3.1 Progressive Monitoring

Transient characteristics at the time of failure will change over time as other phenomena in the system affect the measured variables. Therefore, fault signatures evolve dynamically. For example, a fault in the system may have no effect on the initial magnitude of a variable, but it may affect the first derivative. Immediately after the fault occurs the variable value will be observed to be normal, but as time progresses the magnitude will follow the direction of the first derivative. The notion of using higher order derivatives in the analysis of measured values during monitoring is referred to as *progressive monitoring* [11, 17]. Fig. 2 illustrates the process. When an observed variable does not match a predicted *normal* value, the comparison is extended to predicted higher order derivatives in the signature. If the higher order derivatives match the observed value the hypothesized fault is still retained.

Predictions of magnitude deviations that have a high or low value before progressive monitoring is applied indicate abrupt changes, which correspond to discontinuities in the model. A discontinuity must not be confused with a magnitude deviation. A magnitude deviation emerges over time and is identified by the progressive monitoring scheme. The fault isolation process is significantly enhanced if discontinuities can be reliably detected.

2.3.2 Temporal Behavior

Two distinct phases of signals in response to fault disturbances, transient behavior and steady state behavior, carry the distinctive discriminative information for diagnosis. Transient phenomena disappear after a time interval, and so it is important to determine when the transient detection phase terminates. At this point, the system should switch to steady state detection/verification mode.

Palowitch [21] reports that signals may exhibit a *compensatory* (Fig. 3(a)) or an *inverse response* (Fig. 3(b)). A compensatory response exhibits a decreasing slope and gradually moves towards a new steady state value. For an inverse response, after an initial increase or decrease, the signal may reverse direction. An additional phenomenon resulting from abrupt faults can

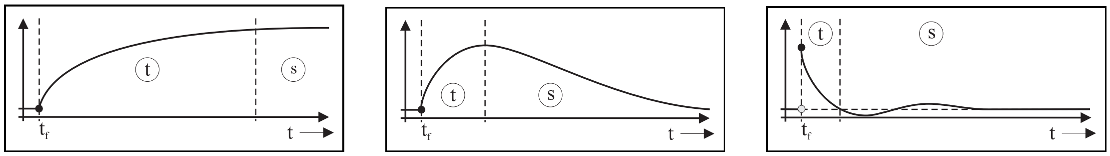

(a) compensatory (b) inverse (c) reverse

Figure 3: Elementary transient signals with different qualitative behavior. A vertical dashed line marks the point where transient monitoring is started and suspended.

be categorized as a *reverse response* (Fig. 3(c)). A reverse response occurs if a discontinuous signal exhibits overshoot, and its qualitative magnitude reverses sign. For example in Fig. 3(c) the signal goes from above normal to below normal.

In the qualitative analysis framework, the transition to steady state analysis is detected from an initial magnitude deviation by noting that:

- For a compensatory response the slope eventually becomes 0.
- For an inverse response there is no discontinuous change of magnitude associated with t_f. The switch from transient to steady state detection occurs when the magnitude and slope deviations take on opposing signs. Eventually the slope may become 0.
- For a reverse response the signal has a discontinuous initial magnitude deviation. The switch to steady state detection occurs when the magnitude changes sign.

When any of these situations are detected, transient verification for that particular signal is suspended (stage t in Fig. 3), and steady state detection is initiated (stage s in Fig. 3). Steady state is detected when a first order derivative becomes 0 for a sufficient period of time. The sufficient period of time is usually based on design information. Techniques specific to steady state detection are not discussed in this paper

3 Signal to Symbol Transformation for Qualitative Diagnosis

In the previous section we identified the following features that must be abstracted from the measured signals for the fault isolation task:

- Magnitude changes in terms of
  - qualitative deviations from nominal (+ and −), and
  - whether a discontinuous change occurred.
- The first order derivative of the signal (+ or −).

This section discusses methods that we have employed for extracting the symbolic features from real signals. We refer to this as the *signal to symbol transformation*.

3.1 Assumptions

Before discussing how to extract specific features from a signal, we point out some general concerns in signal analysis. In general, two key issues govern the signal to symbol transformation task on real data:

1. the discrete time representation of signals, and
2. the presence of noise that contaminates the signals.

Continuous time signals must be discretized in time for data capture and analysis on digital systems. The sampling rate required to correctly sample a bandlimited continuous time signal is given by the Nyquist theorem. However, this is a theoretical lower bound that applies to reconstruction of the signal from its samples by a method that cannot actually be implemented. When we wish to do signal *analysis*, and in particular on-line feature extraction, we are interested in local signal behavior, i.e., signal behavior around a point in time. For this problem, the sampling theorem can at best be a starting point to select a sampling rate. Typically we must use oversampling, that is, sampling at a rate that exceeds the Nyquist rate. This results in a signal representation that is more robust for local signal analysis techniques. This problem has been identified in a variety of signal analysis application (e.g., [27]). Empirical studies are required to determine an appropriate oversampling factor. In the case of FDI even establishing the bandwidth of the system is non trivial. The highest frequency component derived from the system model is an approximation that depends on the level of detail included in the model. Besides, in faulty situations, when system component values may change abruptly, the time constants associated with the changes, are, in general, unknown.

Noise in data leads to a compromise between speed of feature extraction and confidence in the derived result. More signal samples are required to obtain more robust feature values. A crucial point here is that we are interested in feature extraction, meaning signal *analysis*, not signal *restoration*. Noise models that are used in signal processing are often chosen to be Gaussian. A noise model that is not completely accurate does not necessarily invalidate a feature extraction method. However, measurement data often contains *outliers*, sparse samples with values very different from the assumed distribution. Removing outliers is generally beneficial and can be achieved through a median filter of short length.

3.2 Magnitude Changes and Discontinuities

Discrepancy detection initiates the fault generation and refinement process and requires a compromise between false alarms and missed alarms. To reduce the incidence of false alarms, we take a conservative approach, and employ a threshold ($2\% - 5\%$ change from the expected) to make the decision whether a change has occurred or not. The isolation task requires information about the nature of the change, i.e., whether a magnitude change is an abrupt change (or not), and the direction of the change: above ($+$) or below ($-$) the expected value. To derive this information from a signal several approaches are studied:

- Reconstructing the signal from noisy measurements using a transform technique to establish a signal representation in which the transient information is expressed by a sparse number of coefficients. Time-frequency techniques fall into this category.

- Using a signal model to build a change detector by computing a residual of the time series from the model.
- Using statistical hypothesis testing methods, where a change in a parameter value, such as the mean value of a set of samples, above a specified threshold with sufficient statistical confidence implies an abrupt change.

This section discusses some of the method selection criteria for detecting a step change in a constant signal contaminated by Gaussian noise.

3.2.1 Signal Reconstruction

Successive Over-Relaxation (SOR) [4] is a method based on the reconstruction of a signal using splines where discontinuous changes in the spline function are allowed but they incur a penalty. Suppose a signal, $y(n)$, consists of a step edge, $s(n)$, and Gaussian white noise. It is desired to approximate the underlying signal so that noise is reduced, but the discontinuity is not smoothed. The conflicting requirements can be resolved by minimizing a cost function where discontinuous changes are penalized:

$$\mathcal{E} = \sum_i (u_i - y_i)^2 + \lambda^2 \sum_i (u_i - u_{i-1})^2 (1 - l_i) + \alpha \sum_i l_i,$$

where $u(n)$ represents the approximation to $s(n)$ and $l_i \in \{0,1\}$, ($l_i = 1$ implies a discontinuity between adjacent samples). $\mathcal{E}$ must be minimized with respect to the two unknowns u_i and l_i. The first term captures the distance of the data to the estimated signal. The other two terms must be considered simultaneously, since l_i acts as a switch which controls the smoothing. If $l_i = 0$, the two adjacent approximating points are to be joined smoothly and a cost of $\lambda^2(u_i - u_{i-1})^2$ is accrued. If $l_i = 1$, a discontinuity between approximating points exists and a constant penalty of α is imposed. The discrete first derivative, $(u_i - u_{i-1})$, is a measure of smoothness, and λ^2 controls the degree to which the smoothness is enforced. The constant α is necessary for controlling the susceptibility to discontinuities; increasing α, implies a higher resistance to spurious edges.

Minimizing the cost function is non-trivial since it is non-convex. Blake and Zisserman [4] suggest a graduated non-convexity (GNC) algorithm to minimize an approximation of $\mathcal{E}$. We can obtain an expression for the smallest step that can be detected as a function of α and λ, and an expression for a lower bound on α so that no spurious steps are detected. A gradient in the signal that exceeds $\sqrt{\alpha/2\lambda^3}$ is susceptible to multiple discontinuities (i.e., ramp functions are replaced by a series of steps). Fig. 4(a) shows the results of the algorithm on the step change signal. The l_i parameter is not shown since the location of the discontinuity can be clearly seen in the reconstructed signal.

3.2.2 Wavelet Transform

An abrupt change in a signal can be interpreted as a local high frequency component of the signal. The Fourier transform is the standard method to analyze the frequency behavior of a signal. However, with the Fourier transform it is not possible to assess the frequency behavior localized in time. Several modern transform techniques have been developed through which

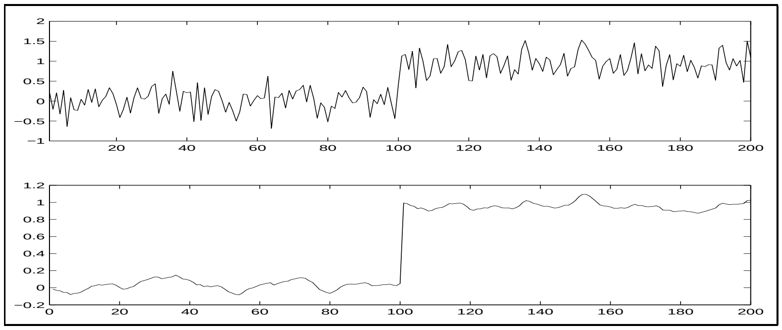

(a) Successive Over-Relaxation (SOR). Penalty for a discontinuity $\alpha = 0.5$, smoothness constraint $\lambda = 4$.

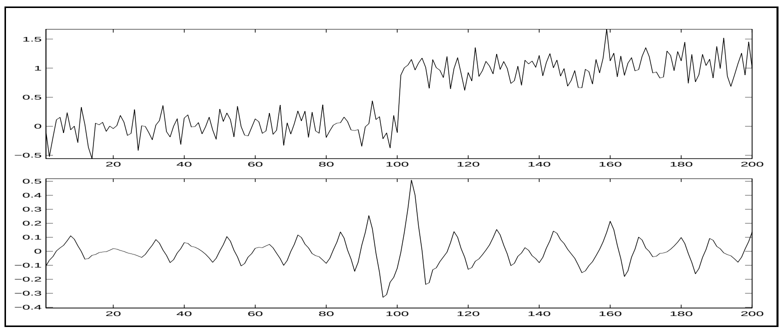

(b) Discrete Wavelet Transform (DWT) with the Daubechies-3 wavelet, with 5 levels of decomposition. Shown is level d3. The change is detected by a threshold at this level.

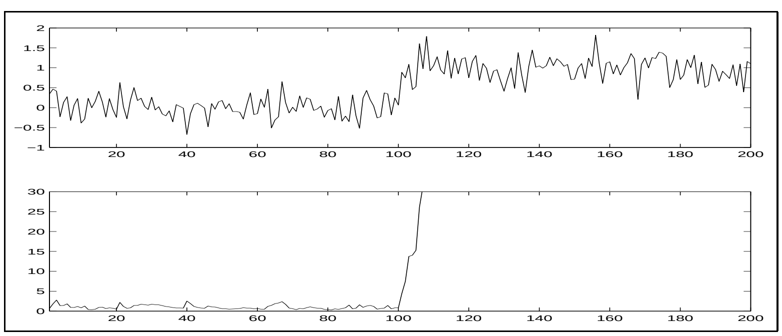

(c) Generalized Likelihood Ratio (GLR) innovation function. The change is detected by a threshold on the innovation function.

Figure 4: Abrupt change detection in a unit step function with Gaussian noise ($\sigma = 0.3$) using three different methodologies. The step occurs at x = 100 in all cases.

we can describe the local frequency behavior of a signal. This class of methods is known as *time-frequency* or *time-scale* techniques. The *wavelet transform* is one of those methods. In the wavelet transform, the signal is decomposed into basis functions that are localized in time and frequency [9]. The wavelet transform in the discrete time domain can be computed very fast with the Discrete Wavelet Transform (DWT). We have built our abrupt change detector with the DWT algorithm.

As the wavelet basis function we selected the Daubechies-3 [7] wavelet, and for our test signal made a 5 level decomposition. An abrupt change is detected by applying a threshold on the wavelet coefficients at a specific level. Both the level and the threshold are computed from the noise in the signal and the depth of the decomposition [26]. Fig. 4(b) shows the results of this technique. With increased noise levels in the signal the decomposition level in which to detect the transient will represent a larger time scale and the ability to localize the change point is diminished. Recent work improve on this by using statistical models in the wavelet-domain [6].

The wavelet transform based detector does not assume an underlying signal model. This results in a detector that is also suitable for changes that are not discontinuous changes in the signal but also other localized high frequency components such as sharp peaks (cusps).

3.2.3 Statistical Signal Processing

The signal is assumed to be an independent random variable sequence y_k with probability density functions $p_{\theta_0}(y)$ and $p_{\theta_1}(y)$ before and after the change, respectively. The parameter θ is a signal model parameter. The change detection algorithm is based on statistical hypothesis testing of a change in a parameter of the distribution. There has been extensive development of a systematic framework using this method [2].

The central quantity is the log-likelihood ratio, defined by:

$$S_j^k = \sum_{i=j}^{k} \ln \frac{p_{\theta_1}(y_i)}{p_{\theta_0}(y_i)}$$

When θ_1 is unknown, the decision function is the Generalized Likelihood Ratio (GLR), determined by:

$$g_k = \max_{1 \leq j \leq k} \sup_{\theta_1} S_j^k$$

For the detection of an abrupt change of unknown magnitude in the mean μ_0 of a Gaussian process we can construct this function as [2]:

$$g_k = \frac{1}{2\sigma^2} \max_{1 \leq j \leq k} \frac{1}{k-j+1} \left[\sum_{i=j}^{k} (y_i - \mu_0) \right]^2 ,$$

where σ is the standard deviation and μ_0 the mean of the signal before the change. Fig. 4(c) shows the output of this decision function. The stopping rule is given by:

$$t_a = \min\{k : g_k \geq h\}$$

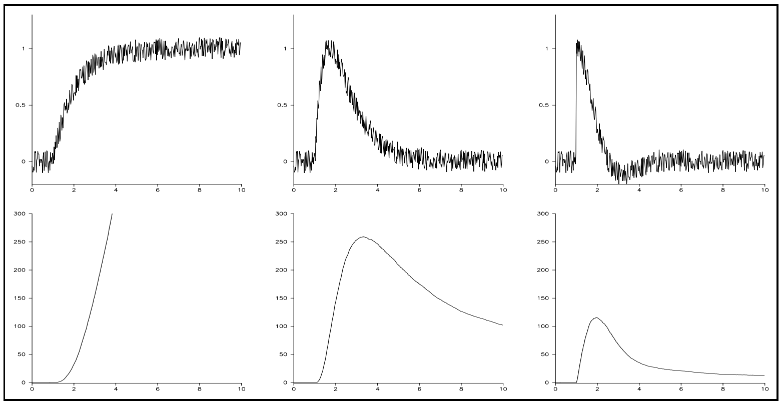

Figure 5: Generalized Likelihood Ratio (GLR) applied to compensatory, inverse, and reverse behaviors.

where h is a predefined threshold and t_a is the detection time.

The GLR is computationally intensive, but a more computationally attractive approximation can be used. Since the GLR algorithm directly uses the statistical signal model it is difficult to interpret results when the method is applied to abrupt changes in signals that do not match the model. Fig. 5 shows the result of applying the GLR detector on the elementary transients discussed earlier. For the compensatory response the GLR is still meaningful in detecting the new steady state (the new hypothesis for the distribution), but this response does not exhibit a discontinuity. For reverse and inverse responses the signal model is clearly not correct and the interpretation of the decision function is unclear.

3.2.4 Evaluation

In our experimental study in Section 5, we have chosen the statistical signal processing method because its model corresponds to the type of signals generated by the fault that we analyze in our experiments. As discussed earlier, we realize that the best discontinuity detection method is clearly dependent on the signal behavior *following* the discontinuity. If we can derive the possible behaviors for a measurement by analyzing the system model, a more reliable and accurate detector can be employed. This is the motivation for developing a tool box with multiple schemes. All the three methods discussed above work reasonably well, but the spline-based SOR method is computationally the most expensive way to analyze the signal. This becomes an important factor in real-time analysis.

3.3 Slope Estimation

The qualitative estimate of the slope, i.e., increasing (+), decreasing (−), and steady (0), is the second feature we extract from the real signal. The slope information is used only *after* the initial discrepancy in magnitude of the signal is detected and the computation of the slope can be deferred until that time. This implies that the slope estimate only applies to the continuous segment of the signal, since any discontinuous behavior coincides with the initial discrepancy.

3.3.1 First Order Difference

To establish a suitable estimator we start with the ideal discrete-time differentiator that has a frequency response:

$$H(e^{j\omega}) = j\omega, \quad |\omega| \leq \pi$$

The filter that exhibits this response is noncausal and of infinite length. Consequently, an actual differentiator can only approximate the ideal behavior. The simplest approximation is the first order difference operator $y(n) = x(n) - x(n-1)$, for which the frequency response is:

$$H(e^{j\omega}) = 1 - cos\,\omega + j\;sin\,\omega$$

This approximates the ideal response for low frequencies, but when ω approaches π, the response deviates significantly from the ideal filter (e.g., [5]). We can improve the performance of the first order difference operator through oversampling of the signal. In effect, the discrete time signal becomes a better instantaneous approximation of the continuous time signal, and, therefore, the performance of the operator improves as well. However, when the signal is noisy, the performance is generally unacceptable.

3.3.2 Linear Approximation

An intuitive approach that may alleviate the noise problem is to compute a piecewise linear approximation of the signal data. A least squares regression on each segment of data gives an approximation of the slope for that segment. This method is straightforward to implement, and it was used in some of our experiments described in Section 5.

3.3.3 Finite Impulse Response

A better approach that we have used recently, derives a Finite Impulse Response (FIR) filter that provides an unbiased minimum variance derivative estimate. This method assumes that the signal is a piecewise polynomial with the same degree as the order of the desired derivative [25]. Figure 6 shows a comparison of the three methods. It is clear that the difference operator is not usable for noisy signals. The performance of the FIR filter and the piecewise linear approximation are comparable for this case. However, the latter method implicitly decimates the derivative signal by a factor equal to the segment length.

4 The Automotive Engine Cooling System

Our work on model-based diagnosis of abrupt faults in continuous dynamic systems encompasses the following characteristics:

- Process under test: The process or device under test is a dynamic physical system.
- Process models: The dynamic behavior of the process can be modeled by a lumped parameter model that captures both normal and faulty system behavior. System behavior is governed by energy exchange between the components of the system. System component failures of interest can be mapped to the parameters of the model.

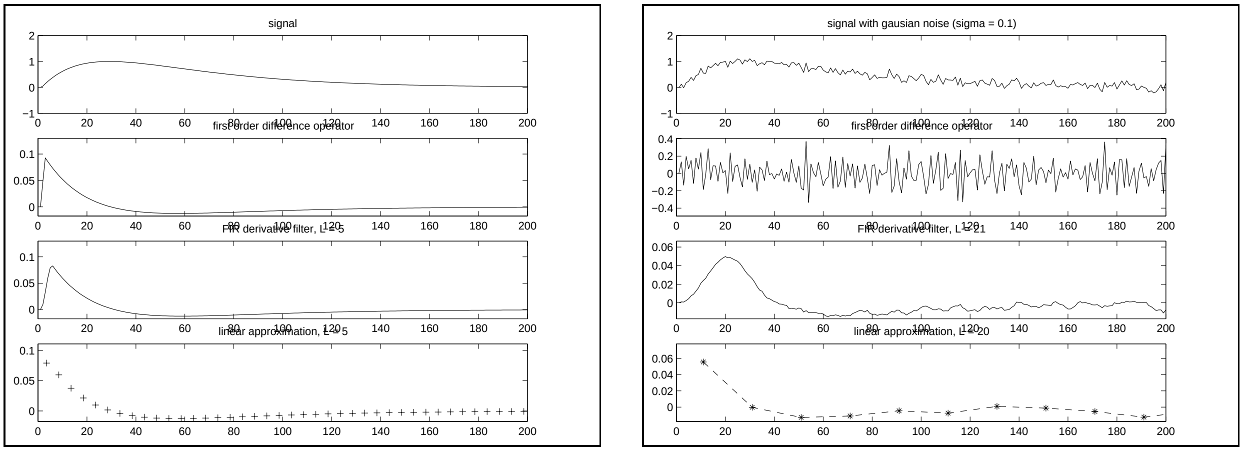

(a) Noise free case.

(b) With additive Gaussian noise ($\sigma = 0.1$).

Figure 6: Comparison of methods for derivative computation for the reverse response in the noise free case (a) and with noise (b).

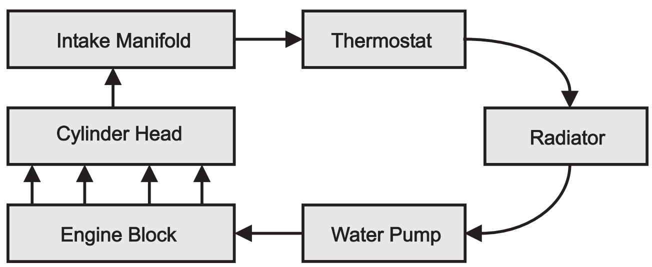

Figure 7: The engine cooling system block diagram.

- Testbed environment: The physical structure of the process is such that sensors can be introduced to measure system variables without significantly altering the normal behavior of the system. In addition, it is important that abrupt faults can be introduced into the system in a controlled manner and they should not cause irreversible or catastrophic faults in a short period of time.

Based on these requirements, we have constructed a testbed at Vanderbilt University around the V-8 internal combustion engine of a Chevrolet automobile. The system is engineered to include a set of pressure, temperature, and flow sensors in the cooling system. Data collection and analysis is performed on an Intel Pentium processor workstation. This section discusses the diagnosis model for the cooling system, and the details of our experimental setup.

4.1 Model of the Cooling System

An automotive cooling system uses a liquid coolant pumped through passageways in the engine block by a centrifugal pump that is driven by the crank axis. The hot coolant flows through the radiator (heat exchanger) where it is cooled by air blown through the radiator by a fan. The coolant is then pumped back into the lower part of the engine block where it moves towards the back. The coolant then flows through the passageways in the cylinder heads up into the intake manifold and towards the front of the engine. Fig. 7 shows a block diagram of the fluid path. If the engine is operating at its desired temperature, the thermostat is open and the coolant circulates back to the radiator. If the engine has not reached its desired temperature, the thermostat is closed and the warm coolant is pumped back into the engine block. When the thermostat is closed it facilitates quick warm up of the engine. When it is open it works to maintain the desired temperature.

The bond graph model [24] of the cooling system, presented in Fig. 8, represents a lumped parameter model that encompasses the mechanical, hydraulic, and thermal subsystems. The hydraulic part, by forcing coolant through the engine blocks, conveys heat away by convection. The heat transfer rate is, therefore, a function of the fluid flow rate. The influence from the thermal subsystem to the hydraulic subsystem, such as fluid expansion with increasing temperature, are not modeled. This implies that discrepancies detected on measurements in the thermal part cannot implicate parameters in the hydraulic part. The primary components of the system are shown as shaded boxes in Fig. 8.

The crank axis is represented by an ideal source of torque, S_e. This source drives the centrifugal pump, represented by the modulated gyrator, MGY, that models the conversion of mechanical energy to fluid energy [11]. The inertia of the rotor in the pump is represented by mass m_1. It is assumed that all hoses and similar fluid paths provide a linear resistance to flow, shown as fluid resistances in the figure. Fluid storage is represented by capacitances. Resistance R_{l-hose} represents the fluid energy losses that occur during coolant flow through the lower hose. The small outlet of the radiator introduces a fluid inertia, $I_{rad-out}$. A T-split coupling in the lower hose used to introduce a leak. Even when it is closed, it allows a small flow of coolant, modeled by resistance R_{leak}. On entering the engine block, the coolant is channeled into passageways, whose fluid resistance is represented by parameter R_{hy-blk}. The fluid capacitance is represented as a capacitor, C_{hy-blk}. The thermostat valve is modeled as a combination of a fluid inertia, I_{tstat}, and fluid resistance, R_{tstat}. When the thermostat valve is open, the coolant flows through the upper hose, represented by fluid resistance, R_{u-hose}, to the radiator modeled as capacitance, C_{hy-rad}.

In our lumped parameter model, conductive heat transfer occurs at two primary locations. First, heat is transferred from the combustion chamber of the engine to the liquid coolant. This is modeled as a constant flow of entropy, modeled by the source, S_f. The heat capacitance of the coolant in the cylinder head and block, modeled by C_{th-blk}, is a function of the specific heat of the coolant, the thermal capacitance per unit mass. The overall capacity to absorb heat is a function of the specific heat and the mass flow rate of the coolant. The mass flow rate is accounted for by a modulated transformer, MTF_2, whose modulation factor is a function of the coolant flow rate determined in the fluid subsystem. The heat capacitance of the coolant in the radiator, modeled by C_{th-rad} and MTF_1, in a similar fashion is a function of the specific heat of the coolant and its mass flow rate, determined by the fluid flow rate in the radiator. Heat transfer from the engine block to the radiator is attributed to the mass transfer of coolant from the engine block to the radiator, and is modeled by a modulated flow source, MSf_1, which is a function of the fluid flow rate in the upper hose. Analogously, convective heat transfer from the radiator to the engine block through the lower hose is modeled by MSf_2. The second location where conductive heat transfer takes place is in the radiator, where the mass of coolant liquid, is cooled by the air blown through the radiator by the fan. The heat transfer occurs through the

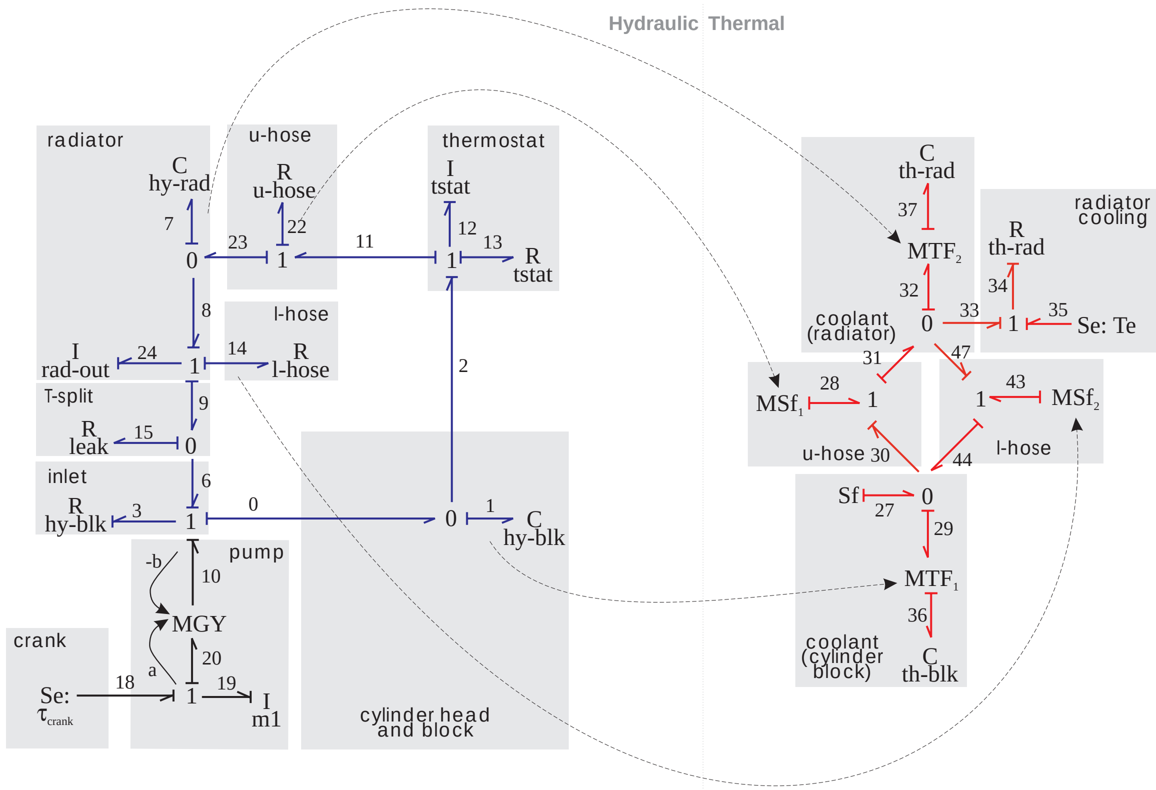

Figure 8: Bond graph model for the engine cooling system.

resistive junction R_{th-rad}, which is a function of thermal conductivity and radiator geometry. The outside air is represented as an infinite sink at constant temperature, T_a, and is modeled by an effort source. In the model, there is an algebraic loop between resistances R_{hy-blk} and R_{leak}. Our diagnosis algorithms can handle algebraic loops but the loop nonetheless results in many deviations that are unknown in a qualitative sense. The non-linearity that results in unknown edge relations in the MGY model of the pump ([17]) are resolved by assuming straight veins for the water pump.

4.2 Experimental Setup

This section describes the cooling system instrumentation, along with sensor location, possible faults that may occur, and the general experimental procedure.

4.2.1 System Under Test and Instrumentation System

Fig. 11(a) shows the complete testbed. The instrumentation system consists of an Intel Pentium microprocessor based computer running the Microsoft Windows NT operating system. This machine is equipped with a data acquisition board with 8 differential inputs, and a maximum acquisition rate of 333kS/s (Data Translation, DT3001-PGL). An external enclosure holds a screw terminal interface to the data acquisition board and is fitted with connectors for the sensors. The enclosure can also house additional signal conditioning electronics, although none is in place at this time. The screw terminal itself provides cold junction compensation (CJC) for thermocouples. To eliminate interference from the ignition electrical system, all wiring from the sensors to the enclosure is shielded.

4.2.2 Diagnosable Faults

Several faults can be introduced into the cooling system without damaging the engine, provided the temperature of the engine block does not exceed certain limits. Possible faults are:

- The thermostat may fail. This may cause a mode switch to happen. Failure may occur either in the open or closed position.
- The belt that drives the water pump may fail. This results in the fan and pump no longer being driven in which case the coolant becomes too hot.
- A hose may get punctured, causing coolant to leak quickly.
- The radiator may start leaking. This is typically a slow leak.
- Metal deposits in the coolant may clog the radiator outlet.
- The water pump may fail, either catastrophically or gradually through wear.

Since we are dealing with a combined thermal and fluid flow problem, it is important to collect temperature and pressure values at various points in the cooling system loop. We have installed sensors at expert selected measurement points. The selection was made based on

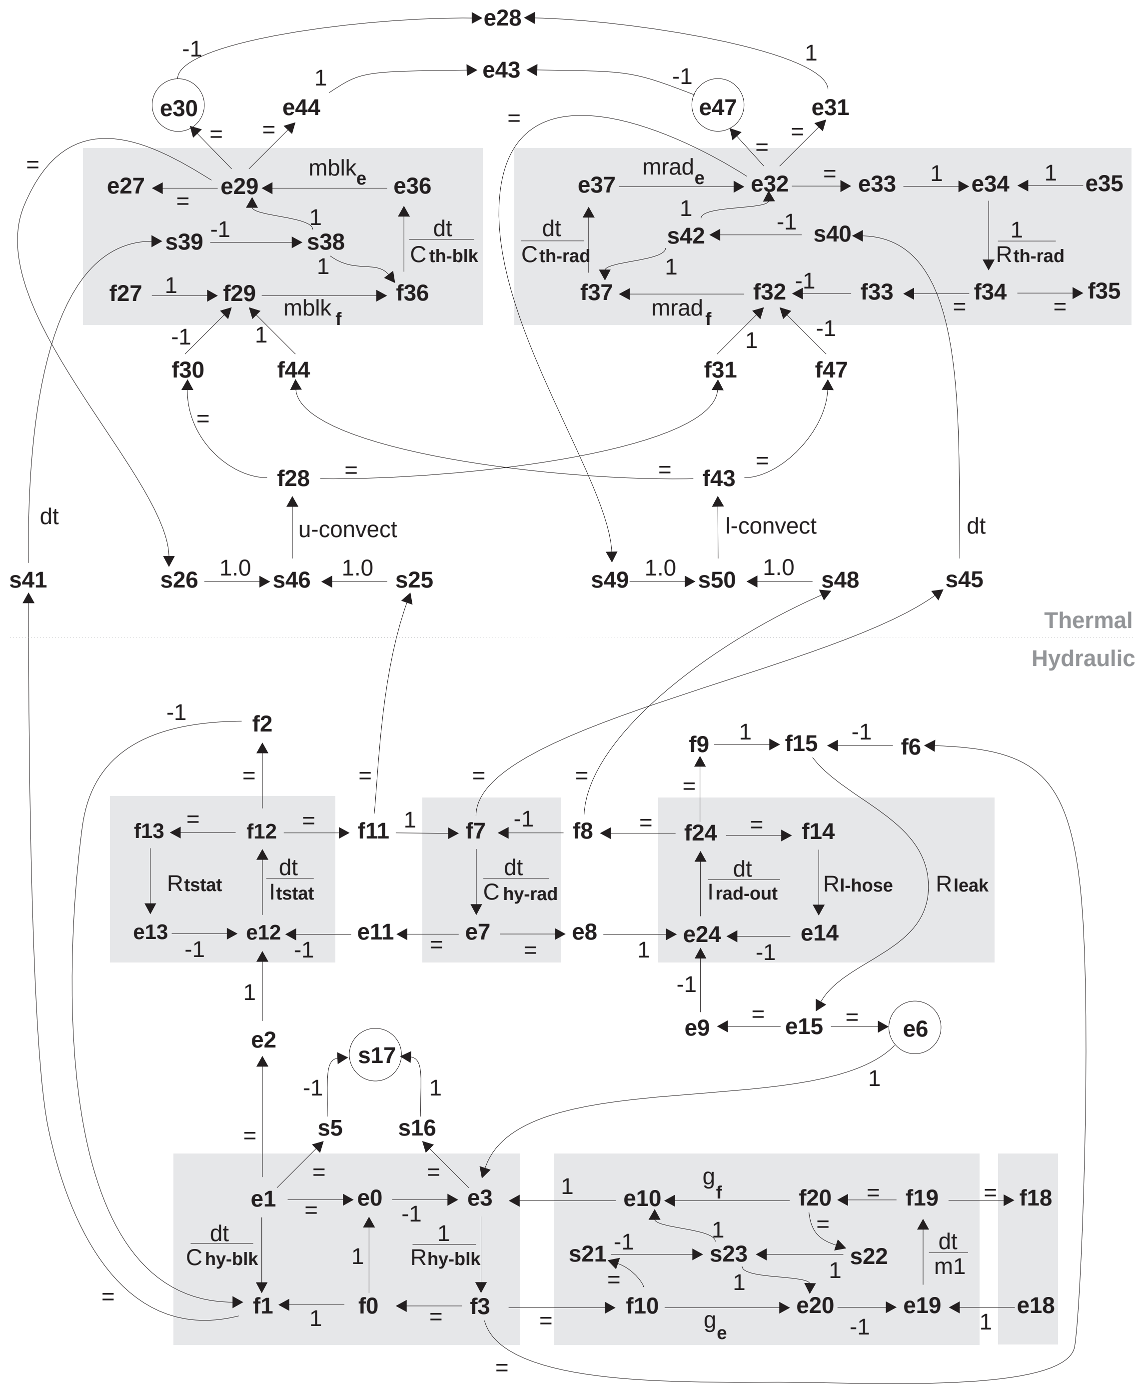

Figure 9: Temporal causal graph for the engine cooling system. Vertices that correspond to measured variables are circled.

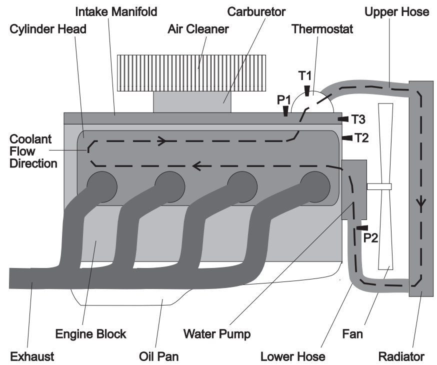

Figure 10: Engine schematic with suggested sensor placement.

the discriminating ability and the possibilities for ease of installation on the engine. In future work, we will apply formal measurement selection methods. Fig. 10 shows the location of these measurement points on the engine block. Table 1 relates the list of faults to the measurements that are affected by a fault occurrence.

Fig. 11(b) shows a detail of the engine with the installed sensors. Thermocouples are used for temperature measurements. A probe style thermocouple (T1) has been installed in the thermostat housing, immediately downstream from the thermostat. A second probe style thermocouple (T3) is installed just upstream from the thermostat in the intake manifold. This location is very close to the cylinder heads, where the coolant reaches its maximum temperature. The sensor is installed in an existing opening in the intake manifold, normally used for the coolant loop that is connected to the automobile heat exchanger. Both probe style thermocouples are immersed in engine coolant. A 'bolt-on' type thermocouple (T2) is fixed to the cast iron engine block but the measurement is currently not a part of the model, and, therefore, not used for diagnosis. An amplified voltage output pressure transducer (P1) is installed in the intake manifold, in an existing opening next to the thermostat housing. This puts the pressure measurement immediately downstream of the thermostat. A second amplified voltage pressure transducer (P2) is installed in the lower radiator hose, to measure the pressure close to the radiator outlet.

The measured variables {T1,T2,P1,P2} correspond to vertices {e47,e30,s17,e6}, respectively in the TCG of the model(Fig. 9). Associating actual physical measurements with variables in the model is not always straightforward. Pressure variables correspond to effort vertices in the

Thermostat failure: open	T1
Thermostat failure: closed	T3, P1
Belt failure	T1, T3
Punctured hose (fast leak)	P1, P2
Radiator leak (slow leak)	T1, T3, P1
Radiator obstruction	P1, P2
Water pump failure: gradual	T1, T3, P1
Water pump failure: abrupt	T1, T3

Table 1: Faults in the cooling system and implicated transducers.

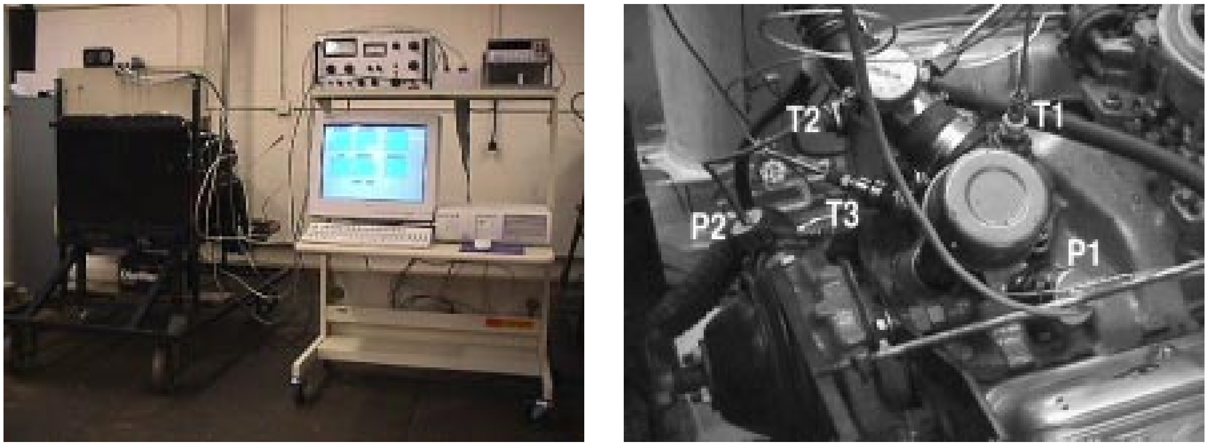

(a) Chevrolet V-8 engine mounted on a rolling frame and personal computer based instrumentation system.

(b) Engine detail with sensors. The fan is on the bottom left, the carburetor on the top right (the air cleaner has been removed).

Figure 11: Chevrolet V-8 engine experimental setup.

hydraulic part of the TCG. However, in our case P1 does not correspond with an actual vertex in the casual graph. Because P1 does not measure the hydraulic pressure at the bottom of a fluid column, but rather the force exerted at the top, it corresponds to the pressure difference between the pressure variables e3 and e1, indicated by the signal vertex s17.

4.2.3 Running Diagnosis Experiments

Our test bed is set up to run the engine system to a steady state operation (or close to that), then to introduce a desired fault, and to collect fault data for diagnostic analysis. The observer model we have implemented only requires nominal values during normal operation from the design specification. This is because the system is stationary during normal operation. Note that this means that accurate estimation of actual system parameters is not required. For more complex systems that exhibit dynamic behavior during normal operation an observer model, which relies on accurate estimates is required to track behavior in real time.

However, parameter estimation for complex systems is a difficult and time consuming process.

Since the nominal steady state values for the measured signals in our test bed are well known, we use these as the reference values for diagnostic analysis. Fault data captured by our data acquisition system is processed by a median filter to remove outliers, and the statistical discontinuity detection and slope estimation techniques are applied to process the faulty data, and convert them into a sequence of symbol values. The processed signals are then analyzed by our diagnosis algorithms. Progressive monitoring is applied to narrow the possible fault hypotheses till the transient data is exhausted. In our experiments, we do not allow the system to run to a new steady state for fear of damage to the engine. The next section describes in detail the results of a diagnosis experiment we have conducted by introducing a punctured hose fault on the experimental test bed.

5 Experimental Results with a Punctured Hose

To simulate a punctured hose, we have inserted a T-split coupling in the lower radiator hose that allows us to drain coolant from the system by attaching a valve to the open end of the coupling. The coupling has a large inner diameter to enable a large outflow. A large hose puncture is simulated by a lever operated gate valve attached to the coupling. This valve can be switched from closed to open very quickly to make the fault look abrupt. We do not drain all the coolant from the engine during the experiment to prevent damage. The valve is closed after a large quantity of coolant is drained.

Fig. 12 shows the results of the experiment in which we introduced a punctured lower hose fault. The stationary running engine was very close to its steady state operation (a small temperature gradient can still be seen in the temperature data). It turns out that the nominal steady state pressure values are lower than expected. This is the result of a small leakage in the closed gate valve. The valve is opened at t = 5 sec. and remains open for several seconds. In this time period, a large volume of coolant is drained from the system. Closing the valve causes transients in both pressure and temperature data that will be ignored for this study and automatically cause transient validation to be suspended based on the three characteristic behaviors in Fig 3.

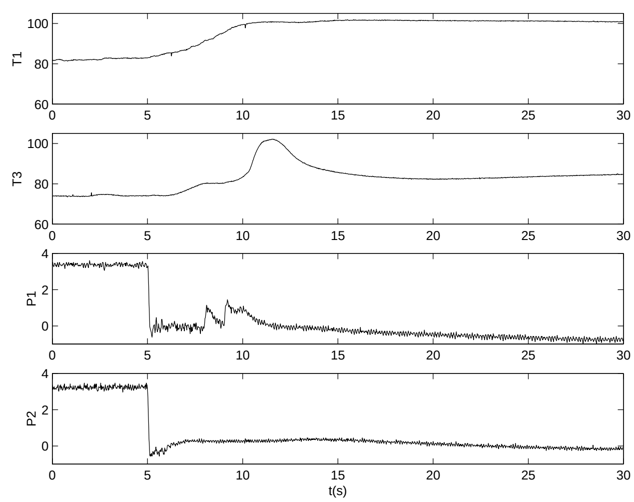

Figure 12: Abrupt loss of coolant through punctured lower radiator hose at t = 5(s).

The sampling time chosen for the experiment was 0.02 sec. The signals were first passed through a median filter of length 5 to remove the few outliers in the temperature data. In this experiment, symbol values for the measurement are computed every second, using 50 samples for the least squares linear approximation of the signal to determine the slope. A step in the monitoring/diagnosis algorithm is executed for every new set of symbolic measurements that are input by the symbol generation unit of the monitoring scheme. Using knowledge from the system model that abrupt changes cannot occur in the temperature signals, abrupt change detection is applied to the pressure signals only. This is an example of domain knowledge in the context of a set of possible faults being systematically incorporated into the TRANSCEND framework to facilitate the signal analysis task.

time	T1	T3	P1	P2
4	(0,.)	(0,.)	(0,.)	(0,.)
5	(0,.)	(0,.)	(0,.)	(0,.)
6	(0,.)	(0,.)	(-,.,*)	(-,.,*)
7	(+,+)	(+,+)	(-,-)	(-,+)
8	(+,+)	(+,+)	(-,-)	(-,0)

Table 2: Results of the signal to symbol transformation on the data around the occurrence of the fault.

Table 2 shows the results of the signal to symbol transformation from step 4 to step 8 (the fault occurs at step 5). The slope values are marked as *unknown* (indicated by a '.') until an initial deviation has occurred. The tuple is extended with a '$\star$' to indicate that an abrupt change was detected during a step. Fig. 13 shows the results generated by the hypothesis generation and refinement procedures. The non linearities and the algebraic loop in the engine model resulted in unknown signatures for second and higher order derivatives for all of the predicted faults. Therefore, only magnitude and slope values were predicted for this model. The steps are indexed starting at 0, which is when the initial discrepancy is detected (in both P1 and P2). This is actually time step 6 in Table 2. The resulting candidates generated by the hypothesis generation algorithm and their corresponding fault signatures generated by the prediction algorithm are listed in the first column of Fig. 13. For some components, both an increase and decrease in fault magnitudes are hypothesized. This contradiction in a qualitative sense is caused by the non linearity in the MGY element. The non linearity causes an unknown value for a vertex during back propagation for candidate hypothesis generation, and this results in branching behavior during candidate generation, both $+$ and $-$ deviations are hypothesized.

Detection of abrupt changes results in the elimination of a large number of hypothesized faults at step 1. At step 2 the $R_{hy-blk}-$ fault is eliminated by applying the progressive monitoring algorithm. The predicted *slope* is negative for measurement e30, which does not match the observed magnitude deviation. Without abrupt change detection, $R_{l-hose}+$ and $C_{hy-rad}+$ cannot be eliminated at step 1 nor at any later time step and will stay in the set of hypothesized faults.

The final result at step 2 in Fig. 13 shows that the diagnosis is accurate because it includes the actual fault $R_{leak}-$. Three spurious candidates are generated. Deviations because of a change in R_{tstat} are unknown for all observations. This is because of the non linearities in the model. In order to be able to discriminate R_{tstat} from other faults, further observations are required. Similarly, $I_{rad-out}+$ can not be distinguished from $R_{leak}-$ given this set of observations. In other work [18, 20], we have shown how a systematic measurement selection algorithm can be employed to establish diagnosability of the system, i.e., which fault sets can be discriminated. Furthermore, this measurement selection algorithm can be used to generate the minimal set of required measurements for complete diagnosability, i.e., all faults can be uniquely discriminated. This analysis has not yet been performed on this system.

step 0

	actual		
	s17:	−	0
	e6:	−	0
	e30:	0	0
	e47:	0	0
$C_{hy-blk}-$	s17:	−	·
	e6:	+	·
	e30:	0	+
	e47:	·	·
g_f-	s17:	−	·
	e6:	+	·
	e30:	0	+
	e47:	·	·
m_1+	s17:	−	·
	e6:	+	·
	e30:	0	+
	e47:	·	·
$R_{leak}-$	s17:	−	·
	e6:	−	·
	e30:	0	+
	e47:	·	·
$I_{rad-out}+$	s17:	−	·
	e6:	−	·
	e30:	0	·
	e47:	0	·
$I_{tstat}+$	s17:	0	−
	e6:	0	+
	e30:	0	·
	e47:	0	·
$R_{l-hose}+$	s17:	0	−
	e6:	0	−
	e30:	·	·
	e47:	·	·
$C_{hy-rad}+$	s17:	0	−
	e6:	0	−
	e30:	·	·
	e47:	·	·
$R_{tstat}+$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·
$R_{hy-blk}-$	s17:	−	·
	e6:	−	·
	e30:	0	−
	e47:	·	·
g_f+	s17:	+	·
	e6:	−	·
	e30:	0	−
	e47:	·	·
m_1-	s17:	+	·
	e6:	−	·
	e30:	0	−
	e47:	·	·
$C_{hy-blk}+$	s17:	+	·
	e6:	−	·
	e30:	0	−
	e47:	·	·
$I_{tstat}-$	s17:	0	+
	e6:	0	−
	e30:	0	·
	e47:	0	·
$R_{tstat}-$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·

step 1

	actual		
	s17:	−	−
	e6:	−	+
	e30:	+	+
	e47:	+	+
$R_{leak}-$	s17:	−	·
	e6:	−	·
	e30:	0	+
	e47:	·	·
$I_{rad-out}+$	s17:	−	·
	e6:	−	·
	e30:	0	·
	e47:	0	·
$R_{tstat}+$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·
$R_{hy-blk}-$	s17:	−	·
	e6:	−	·
	e30:	0	−
	e47:	·	·
$R_{tstat}-$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·

step 2

	actual		
	s17:	−	−
	e6:	−	0
	e30:	+	+
	e47:	+	+
$R_{leak}-$	s17:	−	·
	e6:	−	·
	e30:	0	+
	e47:	·	·
$I_{rad-out}+$	s17:	−	·
	e6:	−	·
	e30:	0	·
	e47:	0	·
$R_{tstat}+$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·
$R_{tstat}-$	s17:	·	·
	e6:	·	·
	e30:	·	·
	e47:	·	·

Figure 13: Fault Detection and Isolation for a ruptured lower hose fault ($R_{leak}-$).

6 Conclusions

This paper describes TRANSCEND, our system for robust monitoring and diagnosis of complex dynamic systems. The key to successful detection and isolation of abrupt faults is the analysis of transients soon after the fault occurs. Quantitative parameter estimation techniques are hard to implement because (i) building accurate numeric models of complex non linear systems over a range of behaviors is a very difficult task, and (ii) the computational burden in solving for parameter values numerically caused by imprecision and convergence problems may become prohibitive. To overcome this problem we developed a qualitative framework for hypothesizing component parameter faults, and then comparing the predicted future behaviors of the faults expressed as signatures with observed behaviors using a progressive monitoring scheme. However, generating qualitative features from real time-varying signals, the signal to symbol transformation problem is a challenging task. This paper discusses a set of model-driven techniques that we have employed for extracting qualitative magnitude changes and slope values from real, noisy signals. Experimental studies have been conducted to study the effectiveness of the different methods, and the monitoring and diagnosis framework has been applied in an automobile engine testbed environment. Results on fault detection and isolation experiments on the cooling system of the engine are reported.

Our experimental studies have brought out a number of important issues that we need to cover in greater detail. For one, building accurate and precise models of complex, non linear systems is a difficult task. We are currently conducting a number of experiments and parameter estimation studies to accurately determine parameter values of components in the system model. In future work, we would like to extend our modeling paradigm to include hybrid models (continuous models with discrete transitions) to simplify the modeling process when system behavior occurs at multiple time scales [10]. The use of hybrid models will necessitate the development of more sophisticated observer schemes [16] and also influence the signal interpretation task.

Currently, our analysis switches to a completely qualitative framework once a discrepancy has been detected in the observations. In future, with more sophisticated signal analysis and interpretation techniques, we could incorporate more numerical information in situations, for example, where we estimate the time constant of a transient. This information can be valuable in terms of further discriminating among possible fault hypotheses. A number of methods have been proposed for identifying exponential functions [23]. Also, as discussed earlier, the signal analysis techniques employed imply a trade-off between timeliness and sensitivity of detection and estimation. Since we use statistical techniques for analysis, it is possible to also extract probabilistic information that indicates the confidence is a certain feature value extracted from the signal. This information incorporated in the diagnosis scheme may help make the fault refinement process more reliable and efficient.

As we look at more complex behavior phenomena, it may be useful to estimate higher order derivatives. Whether higher order derivatives can be reliably computed depends on the signal and techniques may have to be developed to make decisions in real time on the order of derivatives to extract from a signal based on signal characteristics. Also, it may become necessary to process modal characteristics of the observed signals for diagnostic analysis [1]. Higher order modal characteristics will necessitate the definition of more sophisticated signature models (as opposed to looking at magnitude changes and higher order changes as individual features as is done now), and correspondingly more sophisticated signal processing techniques. More systematic studies will have to be conducted on the different signal processing techniques plus noise models that may be active in various situations. Our engine testbed and the computational architecture we have created provides us with a framework for extending our research in a multitude of directions as described above in the future.

7 Acknowledgements

Eric J. Manders and Gautam Biswas have been parially supported be a grant from Hewlett Packard Laboratories, USA. Gautam Biswas is currently on leave at the Knowledge Systems Laboratory, Stanford University, Palo Alto, CA 94305. Pieter J. Mosterman is supported by a grant from DFG Schwerpunktprogram KONDISK. Joel Barnett was instrumental in the construction of the testbed at Vanderbilt University.

References

[1] M. Basseville. In situ damage monitoring in vibration mechanics: diagnostics and predictive maintenance. *Mechanical systems and signal processing*, 7(5):401–423, 1993.

[2] M. Basseville and I. Nikiforov. *Detection of abrupt changes: theory and applications.* Prentice-Hall, 1993.

[3] G. Biswas and X. Yu. A formal modeling scheme for continuous systems: Focus on diagnosis. In *Proceedings IJCAI-93*, pages 1474–1479, Chambery, France, Aug. 1993.

[4] A. Blake and A. Zisserman. *Visual Reconstruction.* MIT Press, 1987.

[5] M. J. Chantler, S. Daus, T. Vikatos, and G. M. Coghill. The use of quantitative dynamic models and dependency recording for diagnosis. In *Seventh International Conference on Principles of Diagnosis*, pages 59–68, Val Morin, Canada, Oct. 1996.

[6] M. Crouse, R. Nowak, and R. Baraniuk. Wavelet-based statistical signal processing using hidden markov models. *IEEE-SIGP*, 46(4):886–902, Apr. 1998.

[7] I. Daubechies. *Ten Lectures on Wavelets.* Society for Industrial and Applied Mathematics Press, vol. 61 of CBMS-NSF Regional Conference Series in Applied Mathematics, Philadelphia, 1992.

[8] D. G. Luenberger. *Introduction to Dynamic Systems: Theory, Models, & Applications.* John Wiley & Sons, New York, 1979. ISBN 0-471-02594-1.

[9] S. Mallat and W. L. Hwang. Singularity detection and processing with wavelets. *IEEE Transactions on Information Theory*, 38:617–643, Mar. 1992.

[10] P. Mosterman and G. Biswas. A theory of discontinuities in physical system models. *Journal of the Franklin Institute: Engineering and Applied Mathematics*, 335B(3):401–439, Jan. 1998.

[11] P. J. Mosterman. *Hybrid Dynamic Systems: A hybrid bond graph modeling paradigm and its application in diagnosis.* PhD dissertation, Vanderbilt University, 1997.

[12] P. J. Mosterman and G. Biswas. A formal hybrid modeling scheme for handling discontinuities in physical system models. In *AAAI-96*, pages 985–990, Portland, Oregon, Aug. 1996. AAAI Press, 445 Burgess Drive, Menlo Park, CA, 94025.

[13] P. J. Mosterman and G. Biswas. An integrated architecture for model-based diagnosis. In *Seventh International Conference on Principles of Diagnosis*, pages 167–174, Val Morin, Canada, Oct. 1996.

[14] P. J. Mosterman and G. Biswas. Monitoring, prediction, and fault isolation in dynamic physical systems. In *AAAI-97*, pages 100–105, Rhode Island, Aug. 1997. AAAI Press, 445 Burgess Drive, Menlo Park, CA, 94025.

[15] P. J. Mosterman and G. Biswas. A theory of discontinuities in dynamic physical systems. *Journal of the Franklin Institute*, 335B(3):401–439, Jan. 1998.

[16] P. J. Mosterman and G. Biswas. Building hybrid observers for complex dynamic systems using model abstractions. In *Hybrid Systems: Computation and Control*, Nijmegen, Netherlands, Mar. 1999.

[17] P. J. Mosterman and G. Biswas. Diagnosis of continuous valued systems in transient operating regions. *IEEE Transactions on Systems, Man and Cybernetics*, 1999. to appear.

[18] P. J. Mosterman, G. Biswas, and N. Sriram. Measurement selection and diagnosability of complex dynamic systems. In *Eighth International Conference on Principles of Diagnosis*, pages 79–86, Mont. St. Michel, France, Oct. 1997.

[19] P. J. Mosterman, G. Biswas, and J. Sztipanovits. Hybrid modeling and verification of embedded control systems. In *Proceedings of the* 7^{th} *IFAC CACSD '97 Symposium*, pages 21–26, Gent, Belgium, Apr. 1997.

[20] S. Narasimhan, P. J. Mosterman, and G. Biswas. A systematic analysis of measurement selection algorithms for fault isolation in dynamic systems. In *Ninth International Conference on Principles of Diagnosis*, pages 94–101, Cape Cod, Massachussetts, USA, May 1998.

[21] B. Palowitch. *Fault Diagnosis of Process Plants using Causal Models*. PhD dissertation, Massachusetts Institute of Technology, Aug. 1987.

[22] R. Patton, P. Frank, and R. Clark, editors. *Fault Diagnosis in Dynamic Systems: Theory and Applications*. Prentice-Hall, UK, 1989.

[23] J. Petersson and K. Holmstrom. Methods for parameter estimation in exponential sums. Technical Report IMa-TOM-1997-8, Malardalen University, 1998.

[24] R. C. Rosenberg and D. Karnopp. *Introduction to Physical System Dynamics*. McGraw-Hill Publishing Company, New York, New York, 1983.

[25] O. Vainio, M. Renfors, and T. Saramaki. Recursie implementation of fir differentiators with optimum noise attenuation. *IEEE Transactions on Instrumentation and Measurement*, 46(5):1202–1207, Oct. 1997.

[26] Y. Wang. Jump and sharp cusp detection via wavelets. *Biometrika*, 82(2):385–397, 1995.

[27] I. T. Youn. Sampling density and quantitative microscopy. *Analytical and Quantitative Cytology and Histology*, 10(4):269–275, Aug. 1988.