A Comprehensive Framework for Model Based Diagnosis

Pieter J. Mosterman *
Institute of Robotics and
System Dynamics
DLR Oberpfaffenhofen
P.O. Box 1116
D-82230 Wessling
Pieter.J.Mosterman@dlr.de

Gautam Biswas †
Department of Computer Science
Vanderbilt University
Nashville, TN 37235
biswas@vuse.vanderbilt.edu

Eric J. Manders ‡
Department of Electrical and
Computer Engineering
Vanderbilt University
Nashville, TN 37235
manders@vuse.vanderbilt.edu

Abstract

This paper reviews our work on monitoring, prediction, and fault isolation methods for complex dynamic systems affected by abrupt faults. The key to this work has been our ability to model the transient behavior in response to these faults in a qualitative framework, where the predicted transient effects of hypothesized faults are captured in the form of *signatures* that specify future behavior for the fault with higher order time-derivatives. The dynamic effects of faults are analyzed by a *progressive monitoring* scheme that avoids direct measurement of second and higher order derivatives from real signal values, which are unreliable when signals are noisy. However, generating qualitative characteristics from real time varying signals is still a challenging problem. More recently, we have been investigating statistical techniques for reliable change detection and labeling in noisy signals. We discuss some of these techniques, and study the tradeoff between speed of detection and sensitivity of analysis. The integrated framework for monitoring, prediction, and fault isolation is being tested on real data generated from an automobile engine test bed.

Introduction

Diagnosis of faults in engineering systems is the process of detecting anomalous system behavior and then isolating the cause for the deviant behavior. Typically the cause is a faulty control setting or a faulty component in the system. In this work, we limit ourselves to finding faulty components in a system. The fault identification and isolation tasks require a model of normal operation of the system, and a number of observable variables. To detect faults, the observed variable values and the system model can be employed to estimate parameters associated with system components. This is referred to as *parameter estimation* (Isermann 1989; Mosterman & Biswas 1998). When parameters deviate from their normal or expected values, the component associated with the parameter is considered to be faulty.

Alternately, the system model can be used in combination with the observed variables to estimate the internal state of the system. This is called *state estimation* (Frank 1990). If the state values deviate from their normal or expected operating values, a fault has occurred in the system. For the fault isolation task, the system model can be extended to include a number of failure types in the components, actuators, and sensors. Normal operation is linked to component parameter values being in the normal range. When a fault occurs, deviations in component parameter values result in deviations in the internal state vectors and nonzero residuals for the states related to the fault modes. Patterns of deviations can be linked back to particular component failures.

Dynamic system behaviors can be derived from first principles based on laws of physics and the system configuration. This approach first compiles the equations that capture physical behavior of each element and then derives a reduced system of first-order differential equations of the form

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx, \end{aligned} \tag{1}$$

where $x(t)$ is the state vector, $u(t)$ is the input vector, and $y(t)$ the output vector. A, B, and C are system parameter matrices. Straightforward substitutions can be applied to the state-space model to generate a system's input-output model:

$$A_n y^n + A_{n-1} y^{n-1} + \cdots + A_0 = B_n u^n + B_{n-1} u^{n-1} + \cdots + B_0. \tag{2}$$

The input-output model for dynamic behavior includes time-derivative terms. By algebraic manipulations, the input-output equations can be rewritten so a measured variable is a expressed as a function of input variables and systems parameters only, but this form will include higher order derivatives for complete system specification.

The time-varying characteristics of dynamic systems are a function of their first and higher order time derivatives. We discuss reliable methods for extracting

*Pieter J. Mosterman is supported by a grant from the DFG Schwerpunktprogramm KONDISK.

† Gautam Biswas is supported by grants from PNC, Japan and Hewlett-Packard Co.

‡Eric J. Manders is supported by a grant from Hewlett-Packard Co.

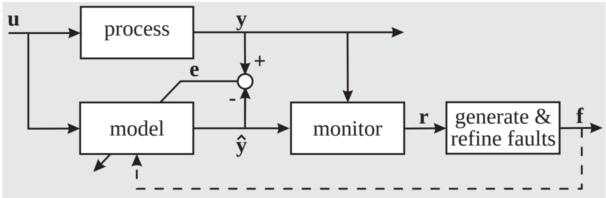

Figure 1: **Diagnosis of dynamic systems.**

this information in later sections of this paper. Traditional systems engineering approaches to fault identification and isolation have been based on quantitative analysis techniques. This requires a fairly accurate model of the system under normal operation and noise reduction techniques to identify signal features. We present a qualitative reasoning methodology that overcomes some of the problems of quantitative techniques. This methodology in conjunction with robust statistical techniques for change detection and slope estimation provides a powerful framework for fault isolation in dynamic systems.

Diagnosis of Dynamic Systems

Fig. 1 illustrates a generic model based approach to fault detection and isolation (Frank 1990; Isermann 1989). A set of variables, called *measurements*, are monitored at frequent intervals during normal operation. Dynamic behaviors from system models are utilized to predict operating values for a chosen set of system variables in a given mode of operation. The diagnosis system maps these measurements, $\mathbf{y}$, that deviate from predicted normal behavior, $\hat{\mathbf{y}}$, onto a system model (Fig. 1). Analysis of discrepancies or residuals, $\mathbf{r}$, in the context of the model helps to *generate* one or more hypothesized *root-causes*, $\mathbf{f}$, that explain the measured deviations. Hypothesized faults suggest modifications to the system models which are then employed to predict future system behavior. Continued monitoring and comparison with these predictions helps *refine* the initial fault set, $\mathbf{f}$. Faults whose predictions remain consistent with the observations determine the root-causes for the observed problems. The goal is to continue the monitoring, comparison, and refinement process till the exact set of faults occurring in the system are isolated. The overall process of monitoring, generating hypothetical faults, prediction, and fault isolation with explicit system models as the core of the analysis scheme is referred to as *model based diagnosis*.

Two basic approaches to diagnosis are traditional quantitative methods developed in systems theory, and qualitative approaches that we have developed based on energy models and ideas from systems theory.

Quantitative Approaches

There are two basic numeric approaches to diagnosis: (i) state estimation (Frank 1990) and (ii) parameter estimation (Isermann 1989).

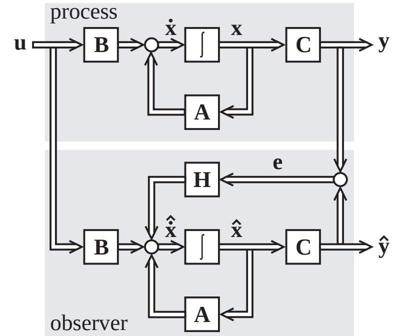

Figure 2: **A general observer scheme.**

State Estimation State estimation relies on accurate numerical models of the system under investigation to reconstruct the internal state from known input and output signals. The state vector captures complete system behavior, so this methodology should capture all possible faults hypothesized by the model. Fig. 2 illustrates an observer method (Luenberger 1979) applied to estimate the state. The observer runs in parallel with the system. The actual parameters of the process defined by the A, B, and C matrices (see Eq. (1)) are estimated in the observer, and a feedback mechanism (the matrix H, which determines the rate of convergence) ensures the estimated state $\hat{x}$ converges to the true value based on computed discrepancies. The idea is to get the estimated output $\hat{y}$ to equal the actual output y, but due to parameter uncertainties and unmodeled real process behaviors, this idealized situation is never reached, and there is a continuous adaptation of the estimated state.

For the diagnosis task, the system equations are extended by a number of additional terms that represent specific types of failures. If the corresponding state variables exceed threshold values, the related fault can be hypothesized to have occurred. A unifying representation, in discrete form, is given by (Frank 1990)

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{E}\mathbf{d}_k + \mathbf{K}\mathbf{f}_k \\ \mathbf{y}_k &= \mathbf{C}\mathbf{x}_k + \mathbf{F}\mathbf{d}_k + \mathbf{G}\mathbf{f}_k \end{aligned} \quad (3)$$

In the model, $\mathbf{d}_k$ represents a disturbance term due to noise, and $\mathbf{f}_k$ represents the effects introduced by the fault term. The $\mathbf{K}$ matrix models the actuators and component faults, and the $\mathbf{G}$ matrix models sensor faults.

Parameter Estimation Parameter estimation relies on an input-output model of the system (Eq. (2)). Measurement snapshots from the system determine the input u^i and output y^j values during system operation. Depending on the number of A_i and B_i parameters, a number of snapshots are required to solve for all the required parameters in Eq. (2). To reduce noise effects, more snapshots are often taken than there are parameters, and a least mean square (LMS) approach provides the best fit (Isermann 1989).

Parameter estimation methods, based on determining parameter values from input-output relations work well for incipient faults where processes undergo very slow changes and the system tends to remain in steady state so the influence of system state is little (Frank 1990). If more dynamic behavior is present, the system state cannot be ignored. This is addressed by using an observer as in Fig. 2 to estimate the system state or by taking measurements so that the system state is part of the output variable set, y_i.

Another issue related to measurement selection stems from the quantitative nature of parameter estimation, which relies on numerical values of magnitude and higher order derivatives of output variables. Often these are hard to obtain in a noisy environment, especially second and higher order derivatives (Chantler *et al.* 1996). Typically, this problem is circumvented by partitioning the model into areas so that only first order derivatives are required to estimate their internal parameters. Isermann states *"the measured variables should be selected such that the process is divided in first order elements, or, in other words, all state variables should be measurable"* (Isermann 1989). In general, a dynamic system has a number of state vectors. In other work, we have shown that the choice of state variables influences the diagnosability of a system (Mosterman & Biswas 1998). Parameter estimation is computationally expensive. It also suffers from convergence problems that are typical of numerical approaches.

A Qualitative Approach to Fault Isolation

Our model-based diagnosis group at Vanderbilt University[1] has been working on a set of real-world diagnosis problems for a number of years (Biswas, Kapadia, & Yu 1997; Mosterman & Biswas 1998; Mosterman *et al.* 1997). An important realization during this period of work has been the shortcomings of a quantitative numeric schemes for diagnostic analysis:

- Imprecision of models – most real world models of complex systems are nonlinear and building precise numeric models of these systems that are accurate over a range of behaviors is a very difficult task (e.g., our experience with models of the secondary sodium cooling loop of fast breeder reactors (Mosterman *et al.* 1997)),
- Computational burden – generating system behaviors from numeric models is computationally expensive, especially since complex dynamic models often have convergence problems due to numeric imprecision (e.g., our experiences with the Space Station thermal bus models (Biswas, Kapadia, & Yu 1997)).

These and other issues have also been discussed by others (e.g., (Frank 1990; Isermann 1989)).

To overcome these problems, we have developed an integrated methodology that combines monitoring, fault detection, fault isolation, prediction, and hypothesis refinement in a qualitative parameter estimation framework. Details of our methodology have been presented in (Mosterman & Biswas 1997; 1998). A difficult problem that arises in our framework is the conversion of real, noisy numeric signals into a qualitative $+, 0, -$ (above normal, normal, and below normal) form. We address this problem by investigating statistical techniques for reliable change and slope detection and labeling of measured signal values.

Measurement Selection Systems theory notions form the basis for developing a systematic qualitative parameter estimation diagnosis methodology that integrates the entire model based diagnosis process of measurement selection, monitoring, fault detection, fault hypothesis generation, and hypothesis refinement by progressive monitoring (Mosterman 1997; Mosterman & Biswas 1996; 1997; 1998). Like quantitative parameter estimation, the qualitative approach assumes system variables and their 1^{st} order derivatives can be derived from the noisy observations. However, instead of precise numerical values, only qualitative $-, 0, +$ values are used for the analysis. In order to avoid problems of obtaining higher order derivatives, a *measurement selection algorithm* (Mosterman, Biswas, & Sriram 1997) based on the model structure first determines sets of measurements that can uniquely identify all possible faults given only magnitude and slope deviations of the possible measurements. The measurement selection process achieves *complete diagnosability* (Mosterman & Biswas 1998).

Dynamic Behavior Given qualitative measurement information (i.e., magnitude and slope deviations of individual variables) parameter deviations are estimated (Mosterman & Biswas 1997; 1998). Theoretically, if one assumes that the measurement set satisfies complete diagnosability, two measurement snapshots (one to determine magnitude deviations and the second to infer slopes) should be sufficient for fault isolation. However, noise in the data requires the imposition of a band around nominal values. Only when a signal value goes outside this band, the signal can be reported to be above or below normal. Fig. 3 shows two variables, a first order effect, x_1, and a second order effect, x_2, and their delay times, t_{d1} and t_{d2}, respectively in crossing the error-threshold. At times between t_{d2} and t_{d1}, x_2 is reported to deviate but x_1 is reported normal. Although x_2 embodies a second order effect with a 0 value 1^{st} order derivative at the point of failure, it crosses the error threshold before a first order effect. This is contrary to expectations where a first order effect is expected to dominate (i.e., be much faster than) a second order effect.

This implies that because of different time constants in the system, deviation information becomes available at different times for different measurements, it is unrealistic to assume that system diagnosis can be completed in two snapshots. After measurement deviations

[1] *http://www.vuse.vanderbilt.edu/~biswas/Research/macs*

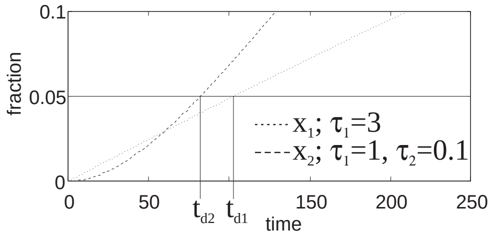

Figure 3: **Delay times for observing deviations.**

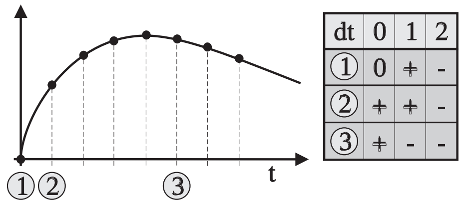

Figure 4: **Progressive monitoring.**

are reported, the system needs to be tracked for a period of time till sufficient information from all other measurements becomes explicit. At times, measurement value changes may remain within the specified margin of error for a period of time (Mosterman 1997), and, therefore, a deviating qualitative value is not observed for a number of steps. To avoid problems, we do not use normal measurements for refuting faults during fault isolation or fault refinement.

Loading effects in the system may cause the dynamic behavior of a fault to change with time. We have developed a *progressive monitoring* scheme (Mosterman & Biswas 1997) to deal with these situations. Consider the signal behavior in Fig. 4. Qualitative signal behavior at the point of failure shows a normal magnitude and positive slope. After a period of time, the positive slope causes the magnitude of the signal to cross the specified threshold, and the measurement is reported to be above normal. This is reconciled by moving the first derivative value into the magnitude slot (step 2) and the observation and estimated behavior are still considered to be consistent. Some time later (Step 3) the second derivative effect causes the slope of the signal to become negative. Like before, this is reconciled by moving the second derivative value in the predicted behavior into the first derivative slot. Therefore, progressive monitoring allows the tracking of transient behaviors due to faults using only magnitude and slope changes.

We have also demonstrated in previous work (Mosterman & Biswas 1998) the importance of other features, such as discontinuity detection in identifying faults. For example, capacitive failures in fluid systems, cause abrupt changes in pressure values, but resistive failures do not. If abrupt changes can be reliably detected, the fault isolation task becomes much more efficient.

Qualitative Landmark Behaviors It is important to realize that transients caused by faults can be tracked

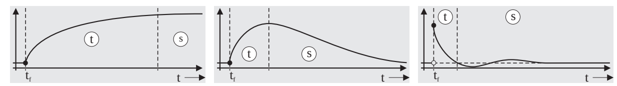

Figure 5: **Typical signal transients in physical systems that exhibit different qualitative behavior over time.**

for short periods of time before they are masked. This is mainly because of loading and compensatory effects which result in significant differences in the qualitative behavior patterns. The implication is that diagnosis should be suspended after a period of time. This is done by identifying three landmark behaviors illustrated in Fig. 5 (Mosterman & Biswas 1997):

- *compensatory response*, after a period of dynamic behavior, the signal settles for a steady state,
- *inverse response*, after an initial transient, signal behavior reverses before settling for steady state, and
- *reverse response*, due to an abrupt change at the time of failure the signal overshoots or undershoots the initial, normal, value.

When any of these situations are detected, transient analysis is suspended for that signal only (stage t in Fig. 5), and steady state detection is activated (stage s in Fig. 5). Therefore, after a period of time, some signals may be processed in the transient mode, whereas others are processed in the steady state mode. Steady state is often hard to detect, and, therefore, activated only as a user-selected option in our monitoring, prediction, and diagnosis framework.

Monitoring

Our qualitative fault detection and isolation methodology requires higher order derivatives to characterize transients in response to abrupt component failures in the system. The progressive monitoring scheme discussed in the previous section provides a framework by which higher order derivative effects are linked to time delays in the manifestation of the behavior in the observed signal, therefore, predicted higher order effects can be analyzed by looking at the sequence of magnitude and first order changes in the observed signal. The major task of the monitoring component then becomes the extraction of qualitative magnitude and slope values and to detect abrupt changes.

Determining Signatures from Real Data

There are two significant issues in dealing with real data in monitoring and fault isolation tasks. The first is the presence of noise in signals, and the second is the discrete time representation of signals for digital signal processing.

Noise in data leads to a fundamental tradeoff between speed and accuracy in diagnostic analysis. On the one hand, one would like to use additional measurement samples so that the effects of noise can be filtered or smoothed out. On the other hand, additional measurement samples imply longer time periods before feature available for processing by the fault analysis routines. In some cases, this may result in the characteristic changes caused by the fault being lost because of interactions and compensatory effects. In other situations, the inability to catch the fault early may result in catastrophic failures in the system.

Fault isolation from transients is based on qualitative signatures. Therefore, we can view the mapping of the signal into a signature as signal interpretation problem (signal to symbol transformation).

Noise reduction Real signals always contain measurement noise. A linear FIR filter that implements a passband will attenuate features with frequency behavior outside the passband as well as the noise. In particular, any abrupt change in the signal, will typically be smoothed out, and the locality of the change point is lost.

Nonlinear filtering may preserve high bandwidth features while at the same time attenuating noise. Statistical order filters (rank order filters) and morphological filters are very well known examples. A common realization of a statistical order filter is the *median filter*. It has the property of preserving abrupt changes in the signal while removing noise components smaller than half the length of the filter kernel. The median filter is particularly well suited for removing outliers which do not fit within commonly used noise models. Removal of outliers requires only a short filter and the shape of the signal is relatively unchanged.

Morphological filters have a filter kernel based on signal shape and not signal statistics. This makes it possible to design a filter which will preserve features with specific shapes (Giardina & Dougherty 1988). These filters may introduce artifacts in the data as a result of the nonlinear operations. The design consideration here is whether or not this will influence the signal to symbol transformation that follows the noise reduction step. If the features of importance to the signature derivation are preserved, distortion may be acceptable.

Selection of the sampling rate The objective of monitoring dynamic behavior is to track the behavior with minimum delay. Again, there is a tradeoff between the detection delay and reliability.

The sampling rate influences the reliability of extracted features. The Nyquist theorem defines the sampling rate required to perfectly reconstruct a signal from its samples. The theoretical result is based on acausal methods, which cannot be used in an on-line realization. Typically, *oversampling*, i.e., use of a sampling rate that exceeds the Nyquist rate, is used in digital signal processing to develop discrete algorithms for sampled continuous systems. It results in a more robust description of the signal in causal systems. This can be very practical in real-time feature detection applications.

Discrepancy Detection

Discrepancy detection is a crucial component of the monitoring system. We trade sensitivity to changes in the signal for robustness to reach a compromise between false alarms and missed alarms. In addition, we need to label the nature of the change, whether it is a first-order change, or an abrupt change. The detection process implies the use of a threshold to make the decision whether a change has occurred or not. A naive approach is to compare the measured signal value to the nominal value directly. This, however, would give a poor performance in the presence of noise, and requires further analysis to label the nature of the change as well.

In the context of this paper we will look at three sophisticated methods for the detection of abrupt changes in more detail. For this purpose the change detection problem is abstracted to the detection of a unit step change in the presence of normally distributed noise. The signal is 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 three methods are based on splines, statistical signal processing and the wavelet transform respectively. Fig. 6 illustrates each method.

Successive Over-Relaxation Successive Over-Relaxation (SOR) (Blake & Zisserman 1987) is a method based on the concept of splines. An energy function is constructed that is minimized to fit a polynomial to the data (Fig. 6(a)). The SOR algorithm (weak-string in the one dimensional case) has a penalty associated with a break in the function. So this algorithm actually will find a continuous time representation for the signal data. The value of the penalty can be thresholded to obtain the desired sensitivity under noise. A higher order version of the algorithm can detect discontinuities in the first derivative, i.e, changes in slope (weak-rod in the one dimensional case).

Statistical signal processing This approach is based on statistical hypothesis testing. Much work has recently been done in developing a systematic framework using this method (Basseville & Nikiforov 1993). The central quantity in statistical hypothesis testing is the logarithm of the likelihood ratio, defined by:

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

The generalized likelihood ratio (GLR) will detect an abrupt change of unknown magnitude in the mean μ_0 of a Gaussian process σ. The decision function is given by (Basseville & Nikiforov 1993):

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

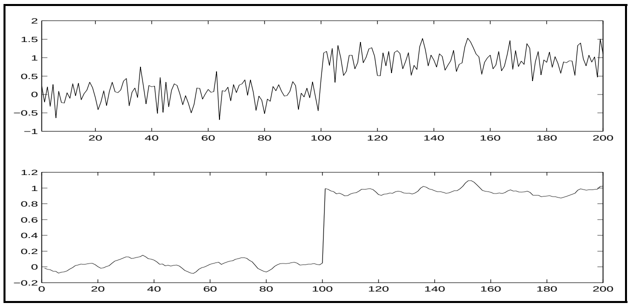

(a) **Abrupt change detection using Successive Over Relaxation (SOR) (weak string algorithm). The method reconstructs the signal, with the breakpoint indicated, and an associated penalty for the break**

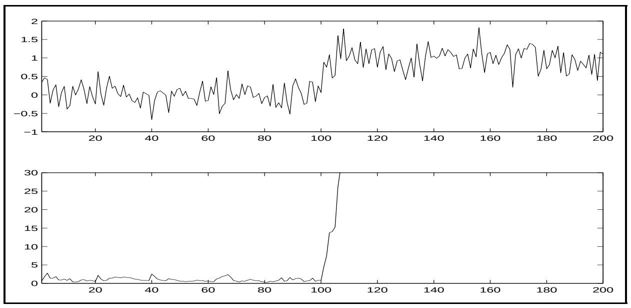

(b) **Abrupt change detection with a Generalized Likelihood Ratio (GLR) innovation function. The change is detected by a threshold on the innovation function**

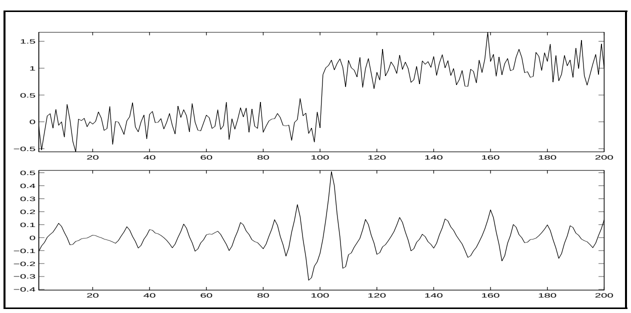

(c) **Abrupt change detection with a Discrete Wavelet Transform (DWT) (Daubechies-3, with 5 levels of decomposition), level d3 shown. The change is detected by a threshold on the wavelet decomposition at this level**

Figure 6: **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**

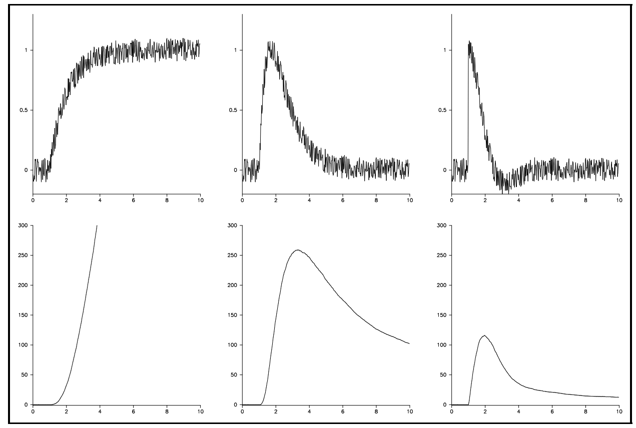

Figure 7: **Generalized Likelihood Ratio (GLR) change detector applied to transients for compensatory, inverse, and reverse behaviors.**

where σ is the standard deviation and μ_0 the mean of the signal before the change. The stopping rule is given by $t_a = \min\{k : g_k \geq h\}$, where h is the predefined threshold and t_a is the detection time. Fig. 6(b) shows the result.

Since the GLR algorithm directly uses the statistical hypothesis it is difficult to interpret results when the method is applied to abrupt changes that do not match this model. Fig. 7 shows the result of applying the GLR detector on the transients discussed in the previous section. The GLR detector still works quite well as a change detector in this example because the sampling rate is chosen to be sufficiently high. Labeling the change based on the value of the innovation function alone would not be reliable.

Wavelet Transform The wavelet transform is the most generic of the methods discussed. It does not make any assumptions about properties of the signal. The ability of the wavelet transform to describe the local behavior of a signal makes it a candidate for a good change detection algorithm (Wang 1995). Change is detected by applying a threshold on a specific component in the wavelet decomposition of the signal as can be seen in Fig. 6(c). This component will have a maximum around the location of the change. With increased noise the component to use will represent a larger time scale and the ability to localize the change point is diminished. Recent work tries to improve on this by using wavelet-domain statistical models (Crouse, Nowak, & Baraniuk 1998).

Slope estimation

The second component of signature derivation is estimation of the slope of the signal after the initial change. The simplest way to do this is a discrete approximation using a difference operator. This approach is extremely sensitive to signal noise because the difference operator acts as a high pass filter (e.g., see (Chantler *et al.*

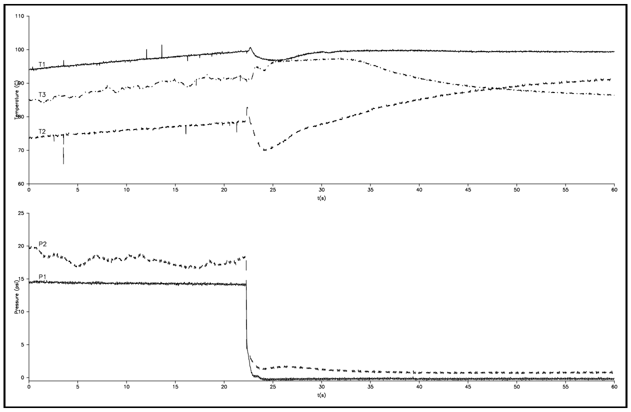

Figure 9: **Measurements from the engine cooling system. A failure in the cooling system (ruptured lower hose) occurs at approximately t = 22(s)**

1996)). In the presence of noisy signals, it is unrealistic to assume that successful diagnosis using first order derivatives can be based on two samples in time. Moreover, the use of higher order derivatives is almost always impossible (unless dedicated transducers such as accelerometers are available). A more reliable method for estimating derivatives is to use statistical model fitting methods. There are number of methods for modeling transients by exponential functions (Landaw & DiStefano III 1984). Regression methods are computationally expensive. However, less complex methods may be used for generating qualitative values.

Experimental Setup

At Vanderbilt University, we have assembled an engine testbed that consists of a Chevrolet V-8 engine, a Hewlett Packard personal computer with an analog to digital converter board running the Microsoft Windows NT operating system, and a set of sensors. Our current focus is on real-time monitoring, fault detection and isolation of faults in the engine cooling system. Presently, we have installed two pressure sensors, and three thermocouples on the engine and in the cooling system loop for this purpose. In the near future we plan to install a flow sensor (hot wire anemometry) just downstream of the thermostat housing in the upper hose. Fig. 8 shows the testbed system.

Fig. 9 shows an example of one minute of acquired data. The sampling time in this experiment is 0.02 (s). At approximately t = 22 (s) a failure is introduced; a ruptured lower hose results in a large abrupt loss of liquid coolant in a matter of a few seconds. The engine was shut off soon after this event to prevent engine damage. The very fast transients in the pressure signals and slower transients in the thermocouple signals are clear. From the data we can also see that the noise contamination of the thermocouple signals disappears after the engine is shut off and that there is both shot noise (sparse signal peaks) and more normally variable noise present.

We are in the process of gathering more experimental datasets like this and applying the methods outlined in the previous section to these types of signals. The nature of the system is such that transients in the pressure signals will always be faster than the transients in the temperature signal, and it may be possible to exploit this fact in the monitoring strategy. In fact, we do not expect to be looking for discontinuities in the temperature data. Also, in large faults like this the transient is clearly identified visually. It will be in the detection of smaller transients, with a poorer signal to noise ratio that will really test our methodology.

Conclusions

This paper reviews our previous work on monitoring and fault isolation in dynamic physical systems (Mosterman & Biswas 1997; 1998). The link between our approaches and quantitative parameter estimation techniques in systems theory is established. The difference is that our fault isolation and refinement techniques are developed in a qualitative modeling and analysis framework. The advantages of working in a qualitative reasoning framework are clearly outlined. Signatures based on magnitude and derivative effects and progressive monitoring schemes are employed to track dynamic behavior responses to abrupt changes. To a large extent this mitigates some of the problems of dynamic tracking of system behaviors based on qualitative simulation schemes (e.g., (Dvorak & Kuipers 1989)).

However, generating qualitative characteristics from real time varying noisy signals is a challenging problem. In this paper, we outline the key characteristics of a monitoring methodology to facilitate the capture of changes in the magnitude and slope of noisy signals. Preliminary work in the use of statistical filters has been presented. Our goal is to develop the engine test bed to the level that we can collect fault data for real-time diagnosis with our integrated monitoring and fault isolation systems.

References

Basseville, M., and Nikiforov, I. 1993. *Detection of Abrupt Changes: Theory and Applications.* Englewood Cliffs, NJ: Prentice Hall.

Biswas, G.; Kapadia, R.; and Yu, X. 1997. Combined qualitative-quantitative steady state diagnosis of continuous-valued systems. *IEEE-SMC* 27A(2):167–185.

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

Chantler, M. J.; Daus, S.; Vikatos, T.; and Coghill, G. M. 1996. The use of quantitative dynamic models and dependency recording for diagnosis. In *Seventh International Conference on Principles of Diagnosis*, 59–68.

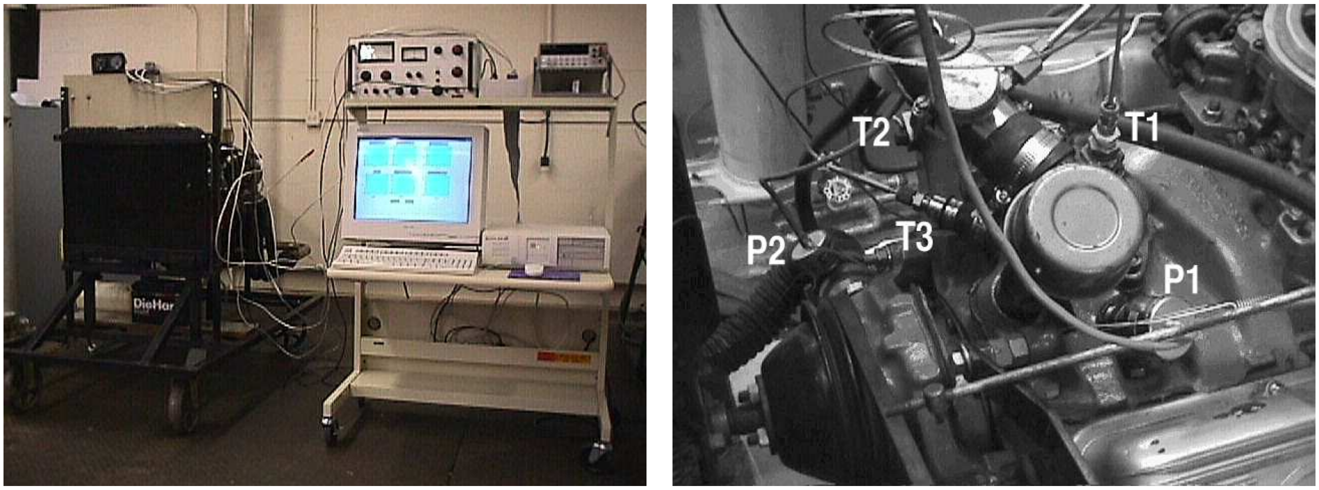

(a) **Chevrolet V-8 engine mounted on a rolling frame and personal computer based instrumentation system.**

(b) **Engine detail with currently installed sensors: pressure sensors are installed in the intake manifold (P1) and the lower hose (P2) and thermocouples are installed in the thermostat housing (T1), on the cylinder head surface (T2) and just upstream of the thermostat (T3).**

Figure 8: **Chevrolet V-8 engine experimental setup.**

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

Dvorak, D., and Kuipers, B. 1989. Model-based monitoring of dynamic systems. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 1238–1243.

Frank, P. 1990. Fault diagnosis: A survey and some new results. *Automatica: IFAC Journal* 26(3):459–474.

Giardina, C., and Dougherty, E. 1988. *Morphological Methods in Image and Signal Processing.* Prentice-Hall.

Isermann, R. 1989. A review on detection and diagnosis illustrate that process faults can be detected when based on the estimation of unmeasurable process parameters and state variables. *Automatica: IFAC Journal* 20(4):387–404.

Landaw, E., and DiStefano III, J. 1984. Multiexponential multicompartimental and noncompartimental modeling ii. data analysis and statistical considerations. *Am. J. Physiol.* 665–677.

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

Mosterman, P. J., and Biswas, G. 1996. An Integrated Architecture for Model-Based Diagnosis. In *Seventh International Conference on Principles of Diagnosis*, 167–174.

Mosterman, P. J., and Biswas, G. 1997. Monitoring, Prediction, and Fault Isolation in Dynamic Physical Systems. In *AAAI-97*, 100–105. Rhode Island: AAAI Press, 445 Burgess Drive, Menlo Park, CA, 94025.

Mosterman, P. J., and Biswas, G. 1998. Diagnosis of continuous valued systems in transient operating regions. *IEEE Transactions on Systems, Man, and Cybernetics.* in review.

Mosterman, P. J.; Sriram, N.; Biswas, G.; Washio, T.; and Yoshikawa, S. 1997. Process Diagnosis in Transient Operating Regions: Fault detection and isolation in a liquid sodium cooling system. In *IJCAI-97 QR Applications in Engineering.*

Mosterman, P. J.; Biswas, G.; and Sriram, N. 1997. Measurement Selection and Diagnosability of Complex Dynamic Systems. In *Eighth International Conference on Principles of Diagnosis*, 79–86.

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

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