Process Diagnosis in Transient Operating Regions: Fault detection and isolation in a liquid sodium cooling system

P.J. Mosterman, G. Biswas
N. Sriram
Center for Intelligent Systems
Vanderbilt University
Nashville, TN 37235
U. S. A.

T. Washio
Inst. for Scientific and
Industrial Research
Osaka University
Ibaraki, Osaka 567
JAPAN

S. Yoshikawa
Advanced Technology Division
PNC Corporation
Ibaraki-ken 311-13
JAPAN

Abstract

Diagnosis of dynamic physical systems is complex and requires close interaction of monitoring, fault generation and refinement, and prediction. We have developed a methodology for model-based diagnosis of abrupt faults in continuous systems that combines dynamic system modeling with a qualitative reasoning approach. A temporal causal system model is used to identify faults from deviant measurements and predict future behavior expressed as signatures, i.e., qualitative magnitude changes and higher order time-derivative effects. A comparison of the *transient* characteristics of the observed variables with the predicted effects helps refine initial fault hypotheses. This allows for quick fault isolation, and circumvents difficulties associated with interactions caused by feedback and dependent faults. This methodology has been successfully applied to the secondary cooling loop of fast breeder reactors. Results on a representative set of simulated faults are discussed in this paper.

1 Introduction

Automated schemes for monitoring and diagnosis for a new generation of complex physical systems are typically based on *functional redundancy* techniques that exploit static and dynamic relations between observed variables in a system for fault detection and isolation. Functional relations among system parameters can be expressed as a set of mathematical constraints, and traditional *filtering* and *observability* methods can be employed to generate *residual vectors* which can be processed by state estimation, parameter identification, and recognizing characteristic quantities [4]. Alternate schemes based on *topological methods* characterize behavior relations as directed graphs constructed from system models under *normal* operating conditions and *faulty* situations [5; 7]. Propagation of observed discrepancies in the graph implicates system components.

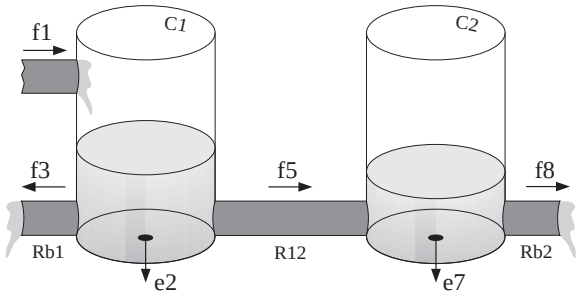

Figure 1: **The bi-tank system.**

An advantage of topological models is their compositional nature, the ability to dynamically partition the system into possible faulty subsystems given a set of observations, and focus on specific constraints that are related to hypothesized faults. A problem with topological models is that they are often incomplete, underconstrained, and *ad hoc* (not derived from physical principles). Therefore, the resultant diagnostic search space is computationally intractable and generates spurious fault candidates. It is important that topological models for diagnosis describe normal and faulty system behavior, incorporate sufficient behavioral detail to map observed deviations to system components and parameters, and generate dynamic behavior caused by faults. Faults change system parameters, therefore, the assumption of constant parameters does not hold[1] and the temporal effects of these causes have to be included.

Our models for diagnosis are derived from bond graph models [11] of the physical system which inherently enforce relevant physical constraints such as conservation of energy and continuity of power. The diagnosis

[1]For example, the behavior of a capacitor based on its constituent relation $V = \frac{q}{C}$, is $\frac{dV}{dt} = \frac{1}{C}\frac{dq}{dt} + q\frac{d}{dt}(\frac{1}{C})$. The typical relation, $\dot{q} = i = C\frac{dV}{dt}$, makes the constant parameter assumption.

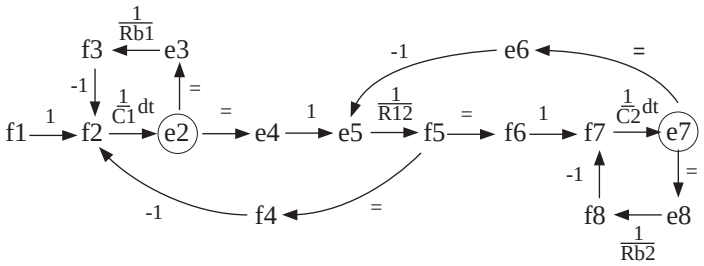

Figure 2: **Temporal causal graph.**

methodology extracts *temporal causal graphs* from bond graphs using the SCAP algorithm [11] and adds parameter and temporal information to links. The resultant graphical structure represents effort and flow variables as vertices, and relations between the variables as directed edges [7]. The bi-tank system in Fig. 1 can be modeled by the temporal causal graph in Fig. 2. Direct proportionality is indicated by a 1 sign and inverse proportionality is indicated by a -1 sign on an edge. The total inflow f_2 into tank 1 with capacity C_1 is directly proportional to inflow f_1. The outflow on the left of C_1 equals the flow through R_{b1}, f_3, which is inversely related to f_2. The outflow on the right of C_1, f_4, equals the flow f_5 through R_{12}, which is inversely related to f_2. The net flow f_2 into C_1 has an integrating relation to the amount of stored liquid in the tank, which appears as the relation between the pressure e_2 and flow f_2 as a time-derivative effect, $\frac{1}{C_1}dt$. Similarly, the flow f_5 through the connecting pipe is related to pressure e_5 by a resistive relation $\frac{1}{R_{12}}$. Resistive, proportional, and equality relations propagate effects instantaneously, whereas energy storage elements introduce integrating effects, which may result in delays in propagating changes.

2 Diagnosis of Dynamic Continuous Systems

Model based diagnosis utilizes system models to predict normal operating values, $\hat{\mathbf{y}}$, for a set of observed variables, $\mathbf{y}$, (Fig. 3). The variables monitored over time are matched against their predicted values in the desired normal mode of operation to compute *residuals* $\mathbf{r}$, as deviations from normal. An observer mechanism accounts for modeling errors, drift, and noise ($\mathbf{e}$) to ensure $\mathbf{r}$ is not spurious. When deviations occur, possible faults, $\mathbf{f}$, are generated by the diagnosis module using constraint analysis and propagation methods applied to the system model. To accurately isolate problems (identify the true fault) we predict future behaviors of the observed variables by introducing these faults into the model and continuing to monitor the observables to check for *consistency* among the predictions and observations.

Our focus in this work is on abrupt faults which are

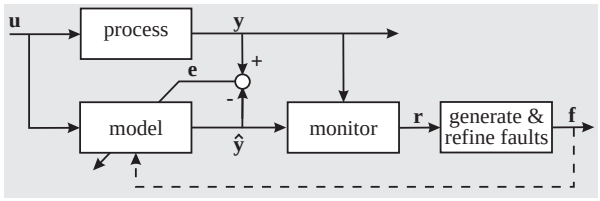

Figure 3: **Diagnosis of dynamic systems.**

dramatic and cause significant deviations from steady state operations called *transients*. For most abrupt faults, if the goal is to quickly isolate faults and return the system to normalcy, it is essential to track and analyze system behavior at frequent intervals soon after the fault occurs, so that their unique transient effects are not lost. Identifying and analyzing dynamic transients caused by faults lays great emphasis on the monitoring and prediction components of the overall diagnosis process, a feature that differentiates our efforts from a lot of the traditional work in model-based diagnosis (e.g., [2; 3; 10]).

Temporal Ordering

Transients are critical in analyzing dynamic physical systems. When faults occur, their effects introduce changes in system behavior that propagate instantaneously to some variables in the system, but have delayed effects on other parts because of the *time constants* involved in the propagation path. Relations that do not embody time constants propagate abrupt changes instantaneously. These abrupt changes occur on a time-scale that is much smaller than the time-scale of the observations, and, therefore, manifest as *discontinuities*. Because physical systems are inherently continuous, this is a sampling artifact associated with the time-scale of the observations. The temporal properties associated with changes, especially discontinuities are exploited in analyzing faults. When multiple phenomena with different time constants affect a measurement, the resultant time constant is hard to estimate without detailed computation. Moreover, faults change parameter values which affect time constants, so quantitative analysis is not al-

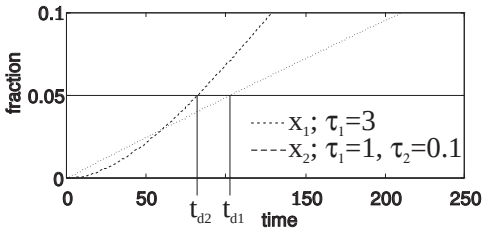

Figure 4: **Delay times for observing deviations.**

ways possible.

Consider two variables, x_1 and x_2 related to the same fault. x_1 embodies a first order effect and x_2 a second order effect. Typically a measurement is considered *normal* if it is within a certain percentage (say 2%–5%) of its nominal value. Fig. 4 shows the delay times (t_{d1} and t_{d2}) associated with the two variables in crossing the error-threshold. At times between t_{d2} and t_{d1}, x_2 is reported deviant but x_1 is reported normal. A second order effect in continuous time should respond more slowly to the failure than a first order effect, but x_2 is reported to deviate first. In a qualitative reasoning framework, the temporal ordering of first and higher order effects of deviations from normal is, in general, impossible unless the sensor system is wired and calibrated to guarantee a temporal ordering in response times. An *observation* reported to be normal at a given time may be a slowly changing signal that has not reached its threshold, therefore, they are not used to refute faults in our consistency-based approach. The only situation where a normal observation reliably refutes a fault is when it is compared against a discontinuity (abrupt change).

Feature Detection

Given the difficulties associated with temporal ordering between signals, the analysis of individual signal features becomes the primary discriminative factor in transient analysis and relies on (i) magnitude deviation, (ii) slopes, and (iii) discontinuity at time of failure. Typically signal deviations are measured in terms of *magnitude changes* and estimated values of *first-order derivatives*. Studies [1] show that noisy signals and sensor errors make higher-order derivative estimation from measurements very unreliable. In this work, we make the assumption that with appropriate filtering techniques qualitative measurement deviations above/below ($\pm$) normal and slopes ($\pm$) can be estimated reliably.

Using physical principles a first observed change in a measurement is considered discontinuous if its magnitude and slope have opposite signs. The discontinuity detection algorithm has been successfully applied in hydraulic systems. Not all discontinuities take this form, though, and some may go undetected which decreases diagnosis resolution.

3 Diagnosis System Implementation

Measured deviations are processed using the temporal causal graph. Details of these algorithms for generating hypothesized faults from observed deviations, predicting future behavior, and monitoring of the predictions to refine the set of possible faults appear in [6; 8].

3.1 Initial Component Parameter Implication

When a discrepancy between measurement and nominal value is detected, a *backward propagation* algorithm operates on the temporal causal graph to implicate component parameters. Implicated component parameters are labeled $-$ (below normal) and $+$ (above normal). The algorithm propagates deviant values backward on the directed edges and consistent $\pm$ deviation labels are assigned sequentially to vertices along the path if they do not have a previously assigned value. An example is shown in Fig. 5 for a deviant right tank pressure, ${e_7}^+$ in the bi-tank system (Fig. 1). ${e_7}^+$ initiates backward propagation along $f_7 \stackrel{\frac{1}{C_2}dt}{\rightarrow} e_7^{\pm}$ and implicates C_2 below normal (C_2^-) or f_7 above normal (f_7^+). Note that this assumes that for e_7 to be high its time-derivative $\frac{de_7}{dt}$ has to have been high at some earlier time. The next step along $f_6 \stackrel{1}{\rightarrow} f_7^+$ implicates f_6^+, and $f_8 \stackrel{-1}{\rightarrow} f_7^+$ implicates f_8^-, and so on. The algorithm operates depth-first processing all instantaneous edges before propagating along first-order changes, and so on [6]. Propagation is terminated along a path when a conflicting assignment is reached. All component parameters along this path are possible faults. As discussed, observed normal measurements do not terminate the backward propagation process.

3.2 Prediction

The main task is to predict the dynamic qualitative deviations in magnitude and derivatives of the observed variables under the fault conditions. This is called a *signature*.

Dynamic Behavior

The forward propagation algorithm propagates the effect of faulty parameters along instantaneous and temporal edges in the temporal causal graph to establish a qualitative value for all measured system variables. Temporal edges imply integrating edges, and, therefore, affect the derivative of the variable on the other side of the edge. Initially, all deviation propagations are 0-order magnitude values. When an integrating edge is traversed, the magnitude change becomes a 1^{st}-order (derivative) change, shown by an $\uparrow$ ($\downarrow$) in the temporal causal graph (Fig. 6). Similarly, a first order change

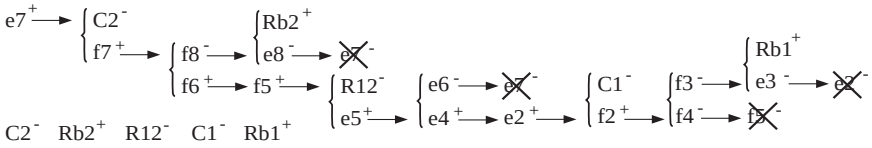

Figure 5: **Backward propagation to find faults.**

propagating across an integrating edge creates a second-order (derivative) change (↑↑ (↓↓) in Fig. 6), and so on.

Forward propagation with increasing derivatives is terminated when a signature of sufficient order is generated as determined by a measurement selection algorithm [9]. A *complete* signature contains derivatives specified to its sufficient order. When the complete signature of an observed variable has a deviant value, monitoring should report a non normal value for this variable.

When assigning values to vertices, situations may occur where the variable has an assigned deviation for the higher order derivatives but the lower order derivatives are not assigned values. Under the single-fault assumption in prediction, this implies that the lower order derivatives of the prediction for the fault under scrutiny are non-deviating.

3.3 Monitoring

This module compares reported signatures and actual observations as they change dynamically after faults have occurred. A number of heuristic mechanisms governed by the dynamic characteristics of the specific system previously discussed improve monitoring quality.

Sensitivity

A heuristic parameter is the sensitivity or the time step employed in the monitoring process, which is a function of the different rates of response the system exhibits. Too large a time step may not distinguish between discontinuities and smooth changes. Too small a time step may unnecessarily burden the real-time diagnosis processor. To ensure reliability, we heuristically estimate the time step as a significant fraction of the smallest time constant in the system. The upper bound on this fraction is chosen to be $\frac{1}{3}$, based on its convergence characteristic in forward Euler numerical simulation. No lower bound is required, since the diagnosis algorithms require a measured deviation of magnitude and slope before their measured value is used in refutation mechanisms. This makes them insensitive to small time steps. For example, discontinuity detection relies on the signal slope. Until a deviation of this slope from normal is detected, no assessment of discontinuity is made.

Progressive Monitoring

Transients generated by failures are dynamic, therefore, the signatures of the observed variables change over time.

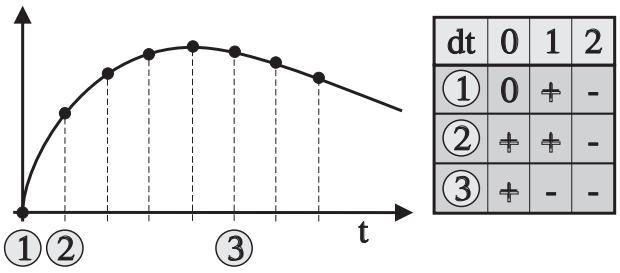

dt	0	1	2
①	0	+	-
②	+	+	-
③	+	-	-

Figure 7: **Dynamic behavior of higher order derivatives.**

For example, a variable may have a magnitude reported normal and a 1^{st} derivative which is above normal. Over time the variable value will go above normal. Incorporating effects of higher order derivatives in the comparison process is referred to as *progressive monitoring*. It replaces derivatives that do not match the observed value with the value of derivatives of one order higher in the signature. Fig. 7 shows time stamps marked 1, 2, and 3, where a lower order effect is replaced by manifested higher order effects. If the predicted deviation of higher order derivatives do not match the observed value, the fault is rejected.

Progressive monitoring is activated when there is a discrepancy between a predicted value and a monitored value that deviates (this applies to 0^{th} and higher order derivatives). At every time point, it is checked whether the next higher derivative could make the prediction consistent with the observation. If this next higher derivative value is normal the next higher derivative value is considered, until there is either a conflict in prediction and observation, a confirmation, or an unknown value is found.

Temporal Behavior

Fault detection from transients becomes invalid after feedback effects have affected behavior significantly. The monitoring process needs to decide when after time of failure, t_f, to suspend transient verification. Signals may exhibit *compensatory* or *inverse responses* (Fig. 8) [10]. For compensatory responses the `slope becomes` 0. For inverse responses, the `magnitude and slope deviations have opposing sign` assuming there was `no discontinuous`

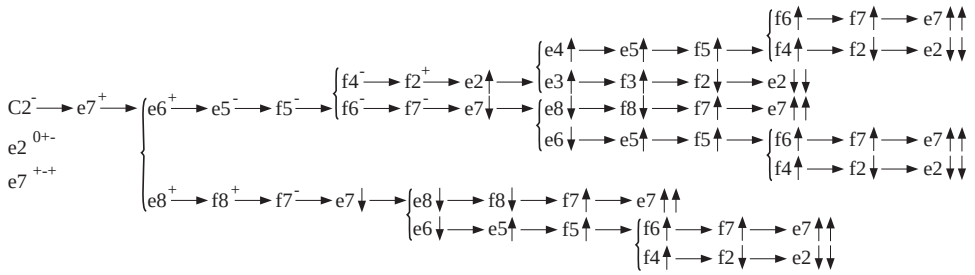

Figure 6: **Forward propagation to establish a signature.**

`change at` t_f. If a discontinuous magnitude change were present, the transient at t_f could manifest as a decrease of this magnitude resulting in a slope with opposite sign. However, this is not an inverse response since the transient effects are exactly those as exhibited at t_f. A *reverse response* constitutes a third phenomenon that occurs for `discontinuous changes at` t_f, and signal overshoot causes `magnitude deviation to reverse sign`. Qualitative observations of magnitude and slope detect these behaviors from an initial magnitude deviation. When these situations are detected, transient verification (stage **t**, Fig. 8) for that particular signal is suspended and steady state detection is activated (stage **s**, Fig. 8).

3.4 Monitoring and Diagnosis Example

We demonstrate the use of the monitoring scheme on a particular fault, a sudden increase in outflow resistance R_{b2}. Fig. 9 shows the results of progressive monitoring, where at times the signatures of the observed variables are modified because of higher-order effects. For example, the signature of R_{b1}^{+} for e_7 changes from $0, 0, 1$ at the beginning of step 9 to $1, 1, 1$. The 2^{nd} order derivative, which is positive, is assumed to have affected the magnitude to make the candidate consistent with the observation $1, 1, .$ in step 9. Discontinuity detection was not employed. When discontinuity detection was used, the same result was obtained in three steps [7]. The diagnosis engine as described correctly detected and isolated all single-fault parameter deviations if pressure in one tank and outflow of the other were measured. Similar results were obtained on a three-tank system [7].

4 The Liquid Sodium Cooling System

Our transient-based diagnosis scheme has been successfully applied to a number of different hydraulic systems (see *http://129.59.101.251:8080/* for a working version of the system) [7]. As an example, we illustrate its application to the secondary cooling loop of a fast breeder reactor. The need for a qualitative approach in this system is motivated by its high-order (six), nonlinearity, and the non availability of precise and real-time numerical simulation models. The precision of flow sensors is limited because of signal noise, and achieving hardware redundancy by installing flow sensors is expensive.

Heat from the reactor core is transported to the turbine by a primary and secondary cooling system. Liquid sodium is pumped through an intermediate heat exchanger to transport heat from the primary cooling loop to the feed water loop by means of a superheater and evaporator vessel (Fig. 10). Pump losses are modeled by R_1. The coil in the intermediate heat exchanger that accounts for flow momentum build-up is represented by a fluid inertia, I_{IHX}. The two sodium vessels are capacitances, C_{EV} and C_{SH}. An overflow column, C_{OFC}, maintains a desired sodium level in the main motor. All connecting pipes are modeled as resistances.

The derivation of the causal relations of the sodium pump (Fig. 11) are based on a modulation factor g between input angular velocity, f_8, and output flow rate, f_9, $g = af_8 - bf_9$. Details are presented in [6]. This factor is directly proportional to f_8 and inversely proportional to f_9, $g(f_8, -f_9)$. The dependency of g on f_8 and $-f_9$ can be explicitly modeled by edges between these variables and the affected variables. In case of the dynamic behavior, the affected variables are input torque, e_8, and output pressure, e_9, and the corresponding edges are added to the causal graph (Fig. 12).

The dependency on system variables of the modulation factor results in nonlinear, quadratic, behavior $e_8 = af_8f_9 - bf_9^2$, and, therefore, the relation on the edge between f_9 and e_8 is unknown. A sensitivity analysis of this relation is shown in Fig. 13 and reveals that depending on the values of f_8 and f_9, the sensitivity of e_8 to f_9 is positive or negative. Given the nominal values of the steady state operation of the system, which is parameter dependent, the weight of $f_9 \rightarrow e_8$ can be determined as a direct (1) or inverse (-1) influence. However, once a deviation occurs, f_8 and f_9 may differ from their nominal values and a different operating point may be reached. Since these new values are caused by failure, and, therefore, unknown, the influence may reverse and is unknown as well. Because this can only occur if e_8 is

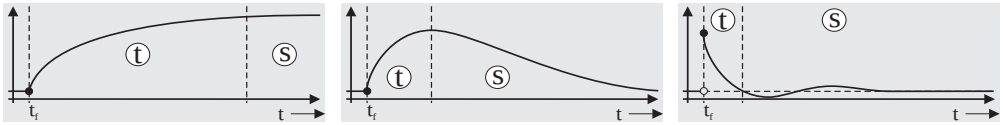

Figure 8: **Qualitative signal transients.**

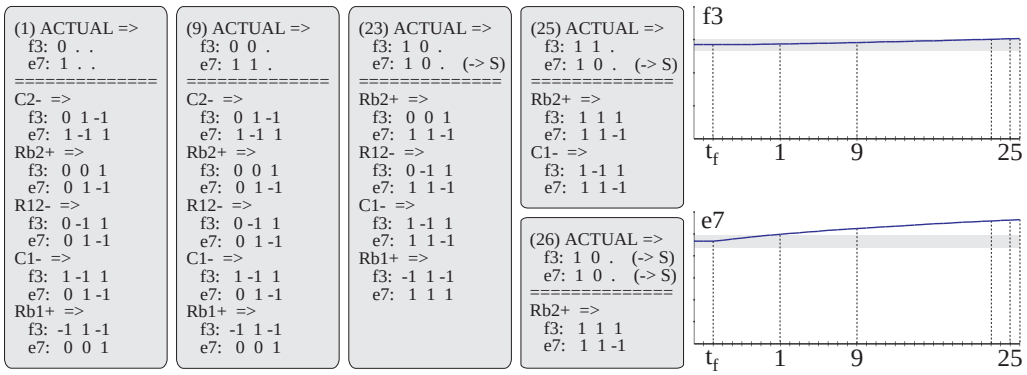

Figure 9: **Progressive monitoring for fault R^+_{b2}.**

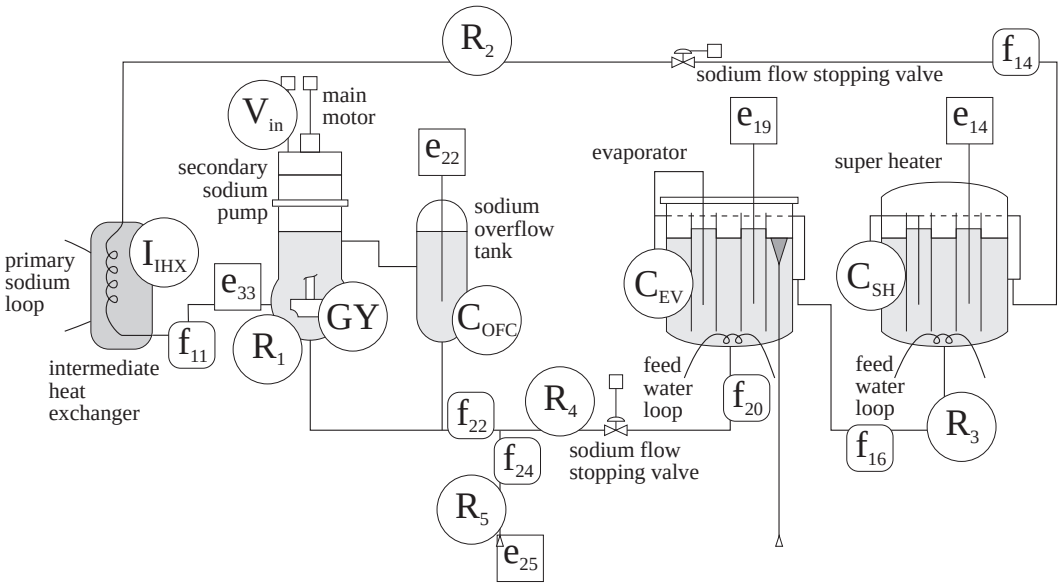

Figure 10: **Secondary sodium cooling loop.**

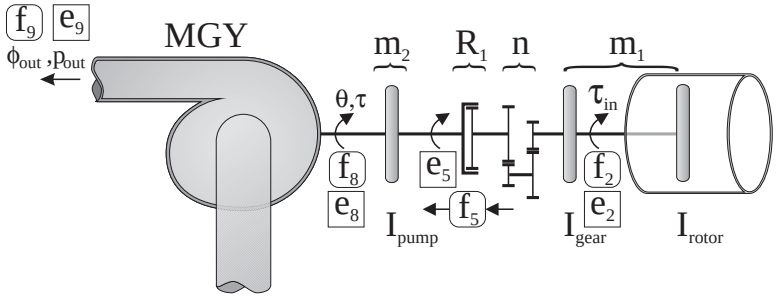

Figure 11: **Synchronous ac motor that drives a pump.**

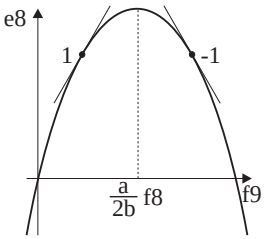

Figure 13: **Detailed sensitivity analysis of $\frac{\partial e_8}{\partial f_9}$.**

predicted to be high based on the proportional influence (-1 or 1), only a predicted decrease in e_8 is unambiguous, and, therefore, propagated. A predicted increase in e_8 is propagated as *unknown*.

4.1 Simulation Results

Design documents were used to choose relative parameter values so the behaviors generated would be qualitatively accurate. The monitoring sample rate corresponded to 20 seconds. Component failures were modeled by changing model parameters by a factor of 5. Capacitive and inductive failures produced abrupt change of pressure and flow, respectively. The margin of error was set at 2% for practical reasons, and signatures were generated up to 3^{rd} order derivatives. A straightforward implementation of the diagnosis algorithms in Visual Basic Pro 3.0 performed fault hypothesis generation and prediction in 9 seconds on a 200 MHz Pentium[2] Pro machine with 32 MB RAM. Steady state was difficult to detect and not used.

The variables $\{f_2, f_8, f_{11}, e_{14}, e_{19}, e_{22}, e_{33}\}$ in Fig. 10 are typical measurements. Simulation results (Table 1) showed that most faults could be accurately diagnosed in a reasonable number of time steps, k. R_3^-, R_4^+, and C_{EV}^- were the exceptions. To detect C_{EV}^-, flow of sodium through the overflow mechanism would have to be monitored. This is a configuration change that we will deal with in future work. R_3^- and R_4^+ were detected but not isolated because the overflow mechanism was not modeled in the temporal causal graph. If this overflow mechanism is modeled, R_3^- and R_4^+ are accurately diagnosed as shown by the entries in parentheses (see also [9]). Precision in diagnosis may improve by considering predicted effects of order higher than 3, but as noted before, they take longer to manifest which may then cause cascading faults to appear. In real situations, cascading multiple faults are more likely than independent multiple faults. Cascading faults are best handled by quick analysis of transients to establish root-causes and then suspend diagnosis when other faults begin to influence the measured transients. In spite of the loss of precision, the results are more practical from a computational viewpoint.

[2]Pentium is a trademark of Intel Corporation.

4.2 Output Additive Noise

To investigate the effects of noise on the measured signals, 2% uniformly distributed noise was added to the output values of the model, for example to model discretization noise. The maximal deviation as a result of this is on the derivative which is the difference between two measurements. This maximum difference is $2 * 2\%$, and, consequently, the margin of *normal* on the qualitative values was chosen 5%. Therefore, the noise effects are well within the qualitative margin of error, and similar results as in the ideal case were expected and found. Table 2 shows a case of typical results, which compared to Table 1, indicates that the diagnosed faults are less precise but still accurate. In some simulation runs, the addition of noise resulted in a more precise diagnosis since previously unnoticed deviations within the original 2% qualitative margin, exceed the qualitative threshold due to additional noise. The only dissonant fault is C_{OFC}^+. For this fault, the modeled faulty behavior of e_{33} is within the noise level, and, therefore, is not detected (see Fig. 14). An increased fault factor results in correct diagnosis when the simulated faulty behavior exceeds the noise level. This exemplifies the trivial situation that in order to diagnose a fault, based on simple thresholding, its effect has to exceed the noise level. More sophisticated methods may be used to distill these abrupt changes from a noisy signal.

5 Discussion and Conclusions

In summary, our results indicate that the qualitative topology-based diagnosis scheme that integrates monitoring, prediction, and fault isolation works well for complex dynamic systems. Successful diagnosis was achieved by setting monitoring parameters in keeping with the dynamic characteristics of the system, and by tracking transients effectively soon after failures occurred. Future work will require the development of more sophisticated feature detection schemes, and a stronger focus on measurement selection to establish *distinguishability* among possible faults. We are presently working on a measurement selection algorithm based on the detectable features that relies on minimal graph coverage techniques.

Acknowledgements

We thank Phillipe Dague for his insightful comments on this work. We also thank PNC Corp., Japan, and Hewlett Packard Research Labs for partially funding this research.

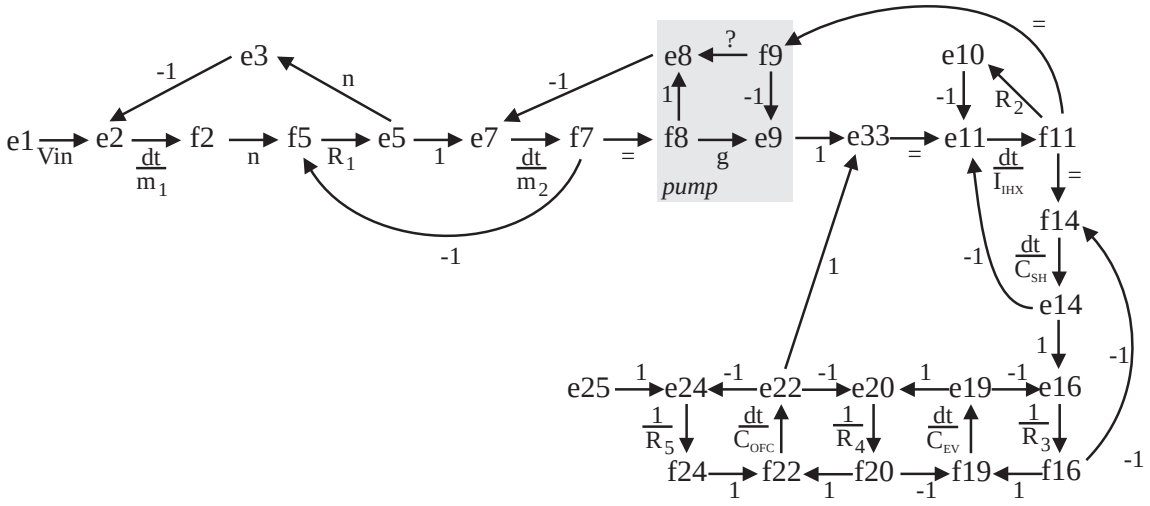

Figure 12: **Temporal causal graph of dynamic behavior.**

Fault	Diagnosis	k	Fault	Diagnosis	k
R_1^+	R_1^+	58	R_1^-	R_1^-	43
R_2^+	R_2^+	27	R_2^-	R_2^-	46
R_3^+	R_3^+	1255	R_3^-	$\emptyset$ (R_3^-)	699 (699)
R_4^+	R_5^+ (R_4^+, R_5^+)	3429 (3429)	R_4^-	R_4^-	43
R_5^+	R_1^+, R_2^-, R_3^+ R_4^+, R_5^+	2	R_5^-	R_3^-, R_4^-, R_5^-	687
C_{SH}^+	C_{SH}^+	73	C_{SH}^-	C_{SH}^-	16
C_{EV}^+	C_{EV}^+	45	C_{EV}^-	-	-
C_{OFC}^+	C_{OFC}^+	9	C_{OFC}^-	C_{OFC}^-	3
I_{IHX}^+	I_{IHX}^+	16	I_{IHX}^-	I_{IHX}^-	2

Table 1: **Fault detection and identification, k is the number of time steps.**

Fault	Diagnosis	Samples	Fault	Diagnosis	Samples
R_1^+	R_1^+	84	R_1^-	R_1^-	38
R_2^+	R_2^+	38	R_2^-	R_2^-	68
R_3^+	R_3^+	2580	R_3^-	$\emptyset$ (R_3^-, R_5^+)	547 (543)
R_4^+	R_5^+ (R_4^+, R_5^+)	2514 (709)	R_4^-	R_4^-	86
R_5^+	R_1^+, R_2^-, R_3^+ R_4^+, R_5^+	2	R_5^-	R_1^-, R_2^+, R_3^- R_4^-, R_5^-	2
C_{SH}^+	C_{SH}^+	128	C_{SH}^-	C_{SH}^-	21
C_{EV}^+	C_{EV}^+	80	C_{EV}^-	-	-
C_{OFC}^+	$\emptyset$	18	C_{OFC}^-	C_{OFC}^-	6
I_{IHX}^+	I_{IHX}^+	20	I_{IHX}^-	I_{IHX}^-	2

Table 2: **Fault detection and isolation under measurement noise.**

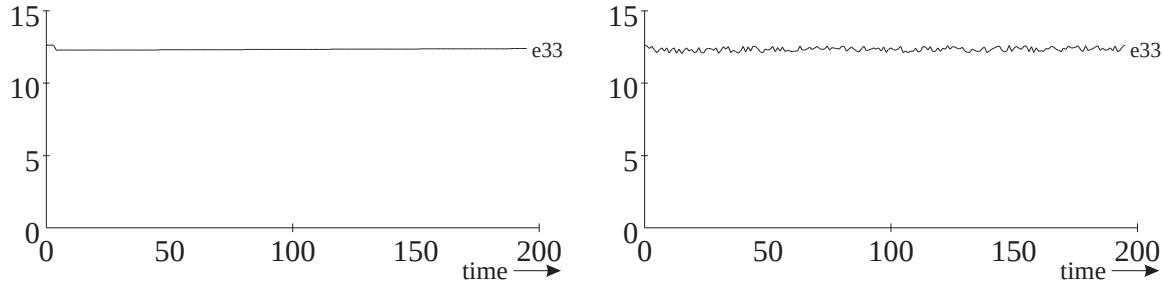

Figure 14: **Response of e_{33} to C^{+}_{OFC} with and without noise.**

References

[1] M.J. Chantler et al. The Use of Quantitative Dynamic Models and Dependency Recording for Diagnosis. *7th Intl. Principles of Diagnosis Workshop*, pp. 59-68, Val Morin, Canada, Oct. 1996.

[2] J. deKleer and B.C. Williams. Diagnosing multiple faults. *Artificial Intelligence*, 32:97–130, 1987.

[3] D. Dvorak and B. Kuipers. Model-based Monitoring of Dynamic Systems. *Proc. 11th IJCAI*, pp. 1238-1243, Detroit, MI, 1989.

[4] P. Frank. Fault diagnosis: A survey and some new results. *Automatica*, 26:459–474, 1990.

[5] A. Misra, J. Sztipanovits, and R. Carnes. Robust diagnostic system: Structual redundancy approach. In *Proceedings of the SPIE's International Symposium on Knowledge-Based Artificial Intelligence Systems in Aerospace and Industry*, pages 249–260, Orlando, FL, April 1994.

[6] P.J. Mosterman. *Hybrid Dynamic Systems: A hybrid bond graph modeling paradigm and its application in diagnosis.* PhD dissertation, Vanderbilt University, 1997. *http://www.vuse.vanderbilt.edu/~pjm/thesis.html.*

[7] P.J. Mosterman and G. Biswas. An integrated architecture for model-based diagnosis. *7th Intl. Principles of Diagnosis Workshop*, pp. 167-174, Val Morin, Canada, Oct. 1996.

[8] P.J. Mosterman and G. Biswas. Diagnosis of Continuous Valued Systems in Transient Operating Regions. *IEEE Trans. on Systems, Man, and Cybernetics,* 1997. in review.

[9] P.J. Mosterman, G. Biswas, and N. Sriram. Measurement selection and diagnosability of complex dynamic systems. In *Eighth International Conference on Principles of Diagnosis*, October 1997. in review.

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

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