A FOCUSED, CONTEXT-SENSITIVE APPROACH TO MONITORING

Richard J. Doyle
Suzanne M. Sellers
David J. Atkinson
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

Abstract

We address two issues which arise in the task of detecting anomalous behavior in complex systems with numerous sensor channels: how to adjust alarm thresholds dynamically, within the changing operating context of the system, and how to utilize sensors selectively, so that nominal operation can be verified reliably without processing a prohibitive amount of sensor data. Our approach involves simulation of a causal model of the system, which provides information on expected sensor values, and on dependencies between predicted events, useful in assessing the relative importance of events so that sensor resources can be allocated effectively.

1. The Monitoring Problem

Timely detection of anomalous behavior is essential for the continuous safe operation and longevity of aerospace systems. The pilot of a jet aircraft must be aware of any conditions which may affect thrust during the critical moments of takeoff. The thermal environment onboard Space Station Freedom must be carefully controlled to provide uninterrupted life support for the crew. The Mars Rover must react quickly to an unpredictable environment or the mission may come to an abrupt conclusion.

Monitoring a physical system involves a number of problem-solving tasks. Dvorak, in his survey of work on expert systems for monitoring and control [Dvorak 87], lists among these tasks recognizing abnormal conditions, combining sensory information into a picture of the global state of a system, isolating faults, predicting both normal and faulted behavior, and maintaining safe operation in the presence of faults. In addition, decisions must often be made in limited time, and with partial information.

The monitoring problem becomes more difficult when the behavior of a physical system involves interactions among components or interaction with an environment. Under these conditions, correct operation becomes context-dependent; it is not possible to determine *a priori* a set of sensor values which always imply nominal operation. Moreover, when the number of sensors in a physical system becomes very large, the ability to combine sensor data into a picture of the global state of a system becomes compromised. Studies of plant catastrophes have revealed that information which might have been useful in preventing disaster was typically available but was not prominent enough within the overwhelming morass of data presented to operators.

In this paper, we concentrate on the initial step in the monitoring process---detecting anomalous behavior quickly and reliably. We do not address here the equally important steps of tracking faulted behavior and determining control actions to continue operation in the presence of faults. Within this focus, we address two important issues: (1) how to adjust nominal sensor value expectations dynamically, taking into account the changing operating context of the system, and (2) how to utilize sensors selectively, determining which subset of the available sensors to use at any given time to verify nominal operation efficiently, without processing a prohibitive amount of data.

2. Two Issues

The traditional approach to verifying the correct operation of a system being monitored involves associating alarm thresholds with sensors. Fixed threshold values

for each sensor are determined ahead of time by analyzing the designed nominal behavior for the system. Whenever a sensor value crosses a threshold during operation, an alarm is raised.

The problem with this approach is that the nominal behavior of even moderately complex systems often depends on context. For example, an earth-orbiting spacecraft periodically enters and emerges from the Earth's shadow. Impingent solar radiation changes the thermal profile of the spacecraft, as does the configuration of currently active and consequently, heat-generating subsystems on board. Thresholds on temperature sensors should be adjusted accordingly. A particular temperature value may be indicative of a problem when the spacecraft is in shadow or mostly inactive, but may be within acceptable limits when the spacecraft is in sunlight or many on-board systems are operating.

Fixed alarm thresholds are useful for defining the operating limits of a physical system, such as the point of overbalance of a rover, or the temperature at which, say, the onboard computer of a spacecraft is at risk. Nonetheless, they are woefully inadequate for verifying the nominal operation of a system with many operating modes, or one which interacts with an environment. The problem is that fixed alarm thresholds are derived from an over-summarized model of the behavior of a system. If the thresholds are chosen conservatively, then false alarms occur. If they are chosen boldly, then undetected anomalies occur. What is needed is a capability for adjusting alarm thresholds dynamically. Alarm thresholds should be chosen according to expectations about the nominal behavior of a system as it changes in different operating contexts. Later on in this paper, we present our approach to dynamic alarm threshold adjustment based on causal simulation of the device.

Another issue which arises in monitoring concerns how to best utilize available sensors to efficiently and reliably, but not necessarily comprehensively, verify the nominal operation of a physical system. Just as the nominal values in a system being monitored depend on context, so do the subset of sensors which can most directly verify those values depend on context. The familiar activity of driving an automobile helps to illustrate this idea. A variety of sensors are provided to the operator of an automobile: fuel gauge, temperature gauge, speedometer, several mirrors, etc. However, the driver does not use all of these diverse sensors all of the time. The speedometer may be checked periodically, or when a speed limit sign is passed; the right-side mirror is probably only used during lane changes. There are two points to be made: one concerns relevance, the other concerns resources.

Individual sensors are appropriate for verifying only some small, localized subset of the possible behavior of a system. The choice of which sensors to sample and interpret at any particular time should be based on expectations of what is to happen in the system and, perhaps, how it is to interact with an environment. However, even after a suitable subset of the available sensors is identified, there may not be the resources available, whether human or machine, to sample all the selected sensors and interpret the data within a required response frame. What is needed is a capability for assessing the importance of predicted events, so that while it may not be possible to comprehensively verify the expected behavior of a system, still the most reliable verification within available resources can be performed.

An illustration of the need to focus attention in monitoring comes from the jet aircraft domain. Some of the recent commercial aircraft catastrophes have been attributed to insufficient thrust during the critical moments of takeoff. There are many possible indicators of low thrust available to a flight crew. For example, a low exhaust gas temperature in an engine may produce reduced thrust. Also, a low turbine fan rotation speed in an engine may imply reduced thrust, because fuel input is based partly on this parameter. The challenge is to direct the attention of the flight crew towards information useful for planning actions in real time without overwhelming them.

A monitoring strategy must take into account the reality that not all sensors should or can be checked all of the time. As the operating context of the physical system being monitored changes, the collection of sensors which provide the most immediate information on the state of the system also changes. Further on in this paper, we present our approach to sensor planning in monitoring. We describe a method for assessing the importance of predicted events in a system, based on reasoning about causal dependencies among events, and about how events relate to intended goals of the designers or operators of a system.

3. Other Work

Within NASA, there are other projects underway in which the goal is to develop a monitoring and a diagnosis capability for aerospace systems. Among these is the KATE project at the Kennedy Space Center, whose domain is the Shuttle Liquid Oxygen Loading system [Scarl *et al* 88]. In this project, causal models are used to support sensor validation, fault diagnosis, and the

planning of control actions.

The goal of the FAULTFINDER project at Langley Research Center [Abbott 88] is to develop an inflight monitoring and diagnosis capability for jet aircraft. These investigators have explored the use of multiple representations and multiple levels of abstraction to be able to reason about diverse faults, to focus attention during reasoning, and to provide accessible information to a flight crew.

Outside of NASA, there have been a number of efforts aimed at developing knowledge-based expert systems for monitoring and control. ESCORT [Sachs *et al* 86] is a shell for developing real-time expert systems to filter and focus information during plant emergency situations. REACTOR [Nelson 82] is an expert system for monitoring nuclear power plants which detects anomalous behavior, assesses the seriousness of the situation, and recommends appropriate actions. REALM [Touchton and Casella 86] is an advisory system which detects and classifies emergencies in nuclear power plants and is able to predict further consequences of those emergencies.

Numerous other examples exist of efforts to develop monitoring and control systems. The reader is referred to Dvorak's excellent survey of the area [Dvorak 87] and to the survey of real-time knowledge-based systems in [Laffey *et al* 88].

The causal reasoning paradigm, which is at the core of our approach to the monitoring problem, is now a well-established area of investigation within Artificial Intelligence. The advantages of the causal approach, which involves modeling a system at the level of components and mechanisms, include the ability to reason about unforeseen interactions, the ability to reason about dependencies among events, and the ability to generate accessible explanations. The seminal efforts in this area include Forbus' process-centered approach [Forbus 85], de Kleer and Brown's device-centered approach [de Kleer and Brown 85], and Kuipers' qualitative mathematics approach [Kuipers 86].

In the specific area of monitoring, Dvorak's MIMIC project stands out as the most comprehensive current research effort [Dvorak and Kuipers 89]. Dvorak creates a component-connection model of a system and employs the QSIM qualitative simulator [Kuipers 86] to generate expectations about the system's nominal behavior. An inductive learning method is used to create a set of symptom-fault rules for known faults, and these rules support the formation of fault hypotheses whenever sensor data does not match predictions from the causal model. When anomalous behavior exists, several fault models can be tracked in parallel until one emerges as the hypothesis with the most explanatory power. The ability to continue tracking a faulted system is important because large, complex systems almost always contain faults and the challenge is to maintain safe operation in the presence of faults.

4. The Approach

At the center of our approach to addressing the two issues of dynamic alarm thresholds and sensor selection is a causal model of the system being monitored and possibly, its environment. Simulation of this model directly solves the problem of alarm threshold adjustment. Predicted values and their time tags indicate how and when to alter the alarm thresholds associated with sensors so that they reflect expectations about the nominal operation of the system in changing contexts.

Another result of simulation is information about causal dependencies among predicted events of a system. This information is used to assess the importance of individual events. Briefly, the most important events are taken to be those which either cause or are caused by the greatest number of other events. An ordering on predicted events reflecting this causal notion of importance serves as the basis for allocating sensor resources to selectively verify the expected behavior of a system [Doyle *et al* 87].

In the remainder of this section, we describe (1) the architecture of our predictive monitoring system, called PREMON, (2) what our causal models of physical systems look like, and how they are simulated, and finally, (3) our approach to sensor planning, based on analyzing causal dependencies.

4.1 Architecture

There are three modules in the PREMON system: a causal simulator, a sensor planner, and a sensor interpreter. See Figure 1.

The causal simulator takes as input a causal model of the system to be monitored, and a set of events describing the initial state of the system and perhaps some future scheduled events. The causal simulator produces as output a set of predicted events, and a graph of causal dependencies among those events.

The sensor planner takes as input the causal dependency graph generated by the causal simulator and determines which subset of the predicted events should be verified.

These events are passed on to the sensor interpreter.

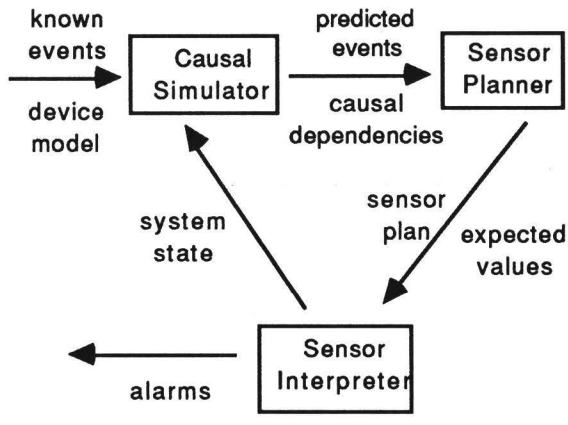

Figure 1. Architecture of PREMON.

The sensor interpreter compares expected values as predicted by the causal simulator with actual values from sensors. Alarms are raised here when there are discrepancies. Finally, the most recent sensed data is passed back to the causal simulator to contribute to another predict-plan-sense cycle of monitoring.

4.2 Causal Models and Causal Simulation

We represent physical systems as a collection of *quantities* and *mechanisms*. Quantities are continuous parameters such as temperature, position, and amount-of-stuff. Quantities are specified by a *physical object*, a *type*, and an *order*. Examples of quantities are {HEATER TEMPERATURE RATE} and {SWITCH POSITION AMOUNT}.

Events describe discontinuous changes in the value of a quantity. Events are specified by a quantity, a *value*, and a *moment*. Examples of events are {HEATER TEMPERATURE RATE POSITIVE 61} and {VALVE-17 POSITION AMOUNT OPEN 0}.

Mechanisms capture causal relations between quantities. More specifically, they describe how a change in one quantity results in a change in another quantity. Examples of mechanisms are HEAT FLOW, THERMAL EXPANSION, LATCH, and GRAVITY. A mechanism is specified by a *time constant*, a *distance*, a *sign*, an *efficiency*, a *bias*, an *alignment*, and a *medium*. Figure 2 shows the representation of a HEAT FLOW mechanism.

A *causal model* then, consists of a set of quantities and a set of mechanisms between those quantities. A causal model can be represented by a graph where the nodes are quantities and the arcs are mechanisms. Simulation of a causal model involves predicting new events, via mechanisms, from known or previously predicted events. The simulation method outlined in the next few paragraphs is described more fully in [Doyle 88, 89].

When the quantity named in an event appears as the cause quantity in a mechanism, a new event is predicted as follows: (1) the quantity of the new event is the effect quantity of the mechanism, (2) the value of the new event is computed from the value of the given event and the sign and efficiency of the mechanism, (3) the moment of the new event is computed from the moment of the given event and the time constant and distance of the mechanism, and (4) the new event occurs only when constraints specified in the bias, alignment, and medium of the mechanism are satisfied. The bias of a mechanism specifies constraints on directions of change. For example, current through a wire can cause it to heat up, but not to cool down. The alignment of a mechanism specifies constraints expressed as inequalities. For example, heat flow is from the warmer to the cooler site. The medium of a mechanism is a physical connection such as a wire, a pipe, a linkage, etc. The predicted effect occurs only when the specified physical connection is in place.

In Figure 2, a typical event is shown, this one describing a temperature change. The HEAT FLOW mechanism is used to predict another temperature change event.

QUANTITY	Chiller Temperature Rate
VALUE	Negative
MOMENT	60

TIME CONSTANT	1.0
DISTANCE	10.0
SIGN	+
EFFICIENCY	0.95
BIAS	nil
ALIGNMENT	<
MEDIUM	{Chiller Pipe-4 Mirror}

QUANTITY	Mirror Temperature Rate
VALUE	Negative
MOMENT	70

Figure 2. A cause event, a mechanism, and an effect event.

Simulation would be straightforward if physical systems could be modeled exclusively as simple mechanism chains between input and output quantities. However, some mechanisms serve to enable or disable other mechanisms, such as a valve controlling a fluid flow, or a latch inhibiting the transmission of motion through a mechanical coupling. In these cases, the contributions of the separate mechanisms combine multiplicatively. The contributions of separate mechanisms also can combine additively, as when two fluid lines empty into the same container, or two opposed forces produce an equilibrium state. The details of our simulation method for interacting mechanisms are described in [Doyle 88,89].

4.3 Sensor Planning

The output of the causal simulator is a trace of predicted events and the dependencies among them. The dependencies are derived from the mechanism structure of the system. A dependency between two events is a record that there is a mechanism in the system which causally relates the events.

Analysis of the causal dependencies in a simulation trace supports decisions about which events to monitor. In our approach, the importance of events is assessed by determining how many other events are effects or causes of a given event. In other words, the importance of an event is related to the amount of subsequent activity it supports and the amount of activity which supports its occurrence. Critical events which lie on several causal paths between inputs and outputs should be verified with care, perhaps with a battery of sensors. On the other hand, events which are side effects and do not support further activity in the system may be ignored completely. See Figure 3.

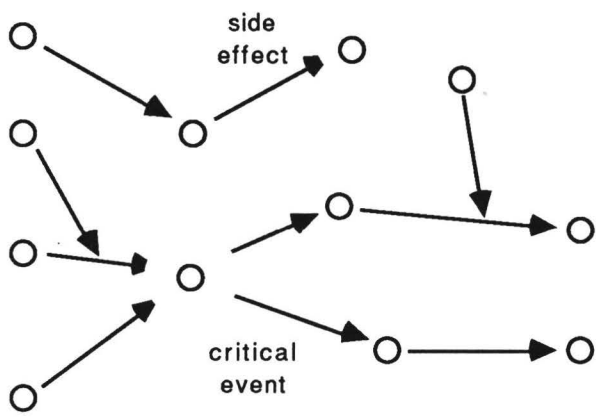

Figure 3. Assessing the importance of events.

This analysis method weights all dependencies in a causal graph equally. Several criteria might form the basis of a non-uniform weighting scheme. For example, *a priori* or empirical knowledge about probabilities of failure might bias the allocation of sensor resources towards those components in a system known to be unreliable. Similarly, parts of a system where redundancy has been built in might be given less careful attention than other parts.

Our causal analysis method for determining what subset of predicted events to monitor is similar to the minimum entropy method of [de Kleer and Williams 87] for determining the site of the most useful next measurement in troubleshooting. Their technique involves propagating observed values and failure probabilities along a causal dependency graph for a circuit.

5. An Example: The JPL Space Simulator

The JPL Space Simulator is an environmental chamber in which spacecraft and instruments can be subjected to some of the aspects of the space environment: intense cold, near vacuum, and solar radiation.

A mirror is used to direct simulated solar radiation onto the spacecraft or instrument inside the chamber. This mirror must be cooled separately from the shroud which surrounds the chamber to compensate for the additional radiation falling on it. Cold gaseous nitrogen is used as the cooling medium and is circulated by a fan. Chilling is achieved by injecting liquid nitrogen into the gaseous nitrogen. Warming is achieved through an electrical heater. A causal simulation of this cooling circuit is shown in Figure 4.

Using the causal analysis technique outlined above, the flow of gaseous nitrogen at the fan is identified as the single most critical event in the predicted nominal behavior of the circuit. This event affects gas flow around the entire circuit and indirectly, heat flow around the entire circuit. The only events unaffected by this event are the source temperature changes at the chiller and heater. This result of causal analysis captures the intuitive notion that nothing at all happens in the cooling circuit if the fan stops operating. Other important events in the predicted operation of the circuit are the temperature changes at the chiller and heater. Measurements made at these sites also provide informative feedback about the nominal operation of the circuit.

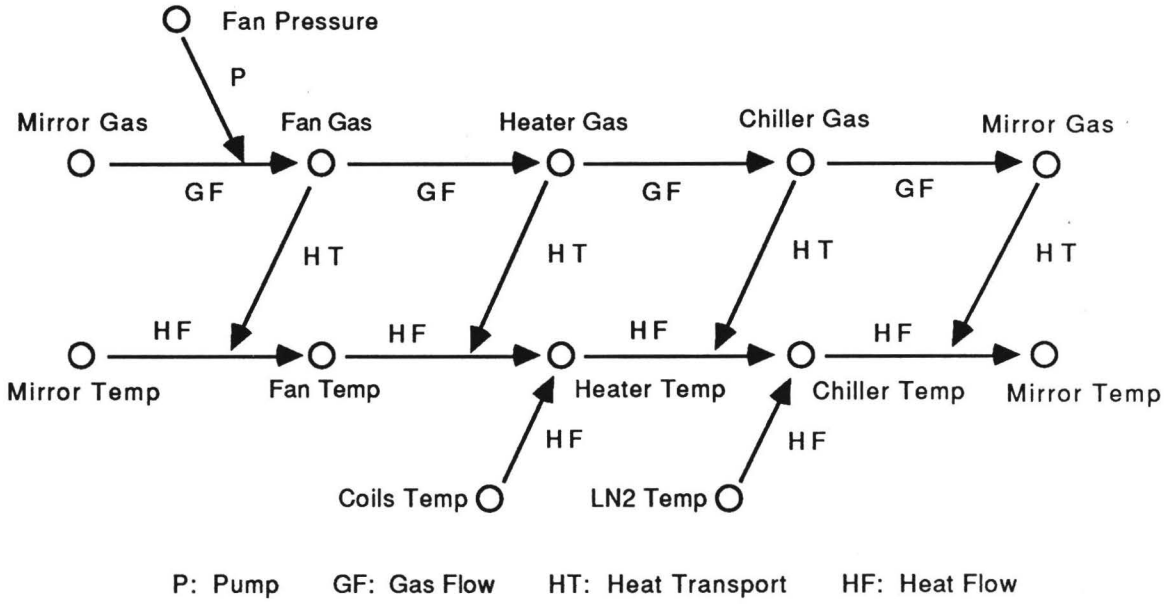

Figure 4. The mirror cooling circuit.

This example has been implemented and illustrates our current causal simulation and sensor planning capabilities. We are beginning to target our developing predictive monitoring capability to other aerospace systems existing or being designed within NASA. Potential domains include power and thermal distribution systems for planetary rovers, ground antenna control systems in the Deep Space Network, and subsystems onboard earth-orbiting spacecraft.

6. Future Work

Our investigations so far have clarified and uncovered other research issues which also need to be addressed. In this section, we briefly outline these issues.

6.1 Generation of Fault Hypotheses

Because complex systems rarely operate completely fault-free, a monitoring system should include also the capability to generate fault hypotheses and to incorporate fault models into the current model of a system so that behavior can continue to be predicted and tracked in the presence of faults. During the diagnostic process, it may be necessary to track several models in parallel, each corresponding to a different fault hypothesis.

6.2 Selective Simulation

We have already made a case for selective use of sensors. In this vein, just as there may be insufficient resources to interpret profusive sensor data, there may be also insufficient resources to conduct a comprehensive simulation of a system. Decisions about which part of a system to simulate also are likely to be context-dependent. Selective simulation interacts with sensor planning in the following way: When simulation cycles are shifted from one part of a system to another, an overhead of sensor reads is needed to determine current state, in order to bootstrap the new focus of simulation.

6.3 How Far Ahead to Predict

Another way of dealing with the problem of limited resources for simulation is to constrain the number of events predicted in a pass through the causal simulator. The tradeoff is between generating enough of a causal dependency graph to drive sensor planning and maintaining a real-time predict-plan-sense monitoring cycle. Another potential factor is ambiguity in simulation. Particularly when a system model is qualitative, there may be insufficient *a priori* information to determine which of several alternate states a system may enter. A direct way to counteract such branching is to suspend simulation and utilize explicit sensor reads to pin down ambiguous states.

7. Conclusions

Detecting anomalies in the operation of a system is a difficult problem when the behavior of the system is com-

plex or involves interaction with an environment, and when the number of sensor channels is large. Under these conditions, nominal values and the most informative sensor data change according to context. We have addressed two specific issues in monitoring: how to adjust alarm thresholds dynamically, and how to verify behavior selectively but reliably. At the center of our approach to solving both problems is the use of a causal model of the system being monitored. Simulation of a causal model serves both to generate expectations about nominal sensor values, and to provide dependency information useful in assessing the importance of predicted events and in allocating sensor resources accordingly.

The key idea in this paper is letting go of the notion of comprehensive monitoring. More likely than not, there will be insufficient resources for predicting behavior and interpreting sensor data. In the face of this limitation, our emphasis is on verifying the operation of a system efficiently and reliably, by carefully focusing computational resources to gather the most informative, if incomplete, feedback on nominal operation within changing contexts.

Acknowledgements

The content of this paper has been enhanced through discussions with Kathy Abbott, Dan Dvorak, Lorraine Fesq, David Franke, Ben Kuipers, and Ethan Scarl. The expertise of Bob Lane was crucial in developing the model of the mirror cooling circuit.

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[Abbott 88] Kathy H. Abbott, "Robust Operative Diagnosis as Problem Solving in a Hypothesis Space," *7th National Conference on Artificial Intelligence*, St. Paul, Minnesota,1988.

[de Kleer and Brown 85] Johan de Kleer and John Seely Brown, "A Qualitative Physics Based on Confluences," in *Qualitative Reasoning about Physical Systems*, D. Bobrow, ed., MIT Press, 1985.

[de Kleer and Williams 87] Johan de Kleer and Brian C. Williams, "Diagnosing Multiple Faults," *Artificial Intelligence*, **32**, 1987.

[Doyle *et al* 87] Richard J. Doyle, Suzanne M. Sellers, and David J. Atkinson, "Predictive Monitoring Based on Causal Simulation," *2nd NASA AI Forum*, 1987.

[Doyle 88] Richard J. Doyle, "Hypothesizing Device Mechanisms: Opening Up the Black Box," Report TR-1047, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1988.

[Doyle 89] Richard J. Doyle, "Reasoning About Hidden Mechanisms," *11th International Joint Conference on Artificial Intelligence*, Detroit, 1989.

[Dvorak 87] Daniel L. Dvorak, *Expert Systems for Monitoring and Control*, Report AI 87-55, Artificial Intelligence Laboratory, The University of Texas at Austin, 1987.

[Dvorak and Kuipers 89] Daniel L. Dvorak and Benjamin J. Kuipers, "Model-Based Monitoring of Dynamic Systems," *11th International Joint Conference on Artificial Intelligence*, Detroit, 1989.

[Forbus 85] Kenneth D. Forbus, "Qualitative Process Theory," in *Qualitative Reasoning about Physical Systems*, D. Bobrow, ed., MIT Press, 1985.

[Kuipers 86] Benjamin J. Kuipers, "Qualitative Simulation," *Artificial Intelligence*, **29**, 1986.

[Laffey *et al* 88] Thomas J. Laffey, Preston A. Cox, James L. Schmidt, Simon M. Kao, and Jackson Y. Read, "Real-Time Knowledge-Based Systems," *AI Magazine*, **9**, no. 1, 1988.

[Nelson 82] William R. Nelson, "Reactor: an Expert System for Diagnosis and Treatment of Nuclear Reactor Accidents," *2nd National Conference on Artificial Intelligence*, Pittsburgh, 1982.

[Sachs *et al* 86] Paul A. Sachs, Andy M. Paterson, and Michael H. M. Turner, "Escort - an Expert System for Complex Operations in Real Time," *Expert Systems*, **3**, 1986.

[Scarl *et al* 88] Ethan A. Scarl, John R. Jamieson, and Edwin New, "Deriving Fault Location and Control from a Functional Model," *3rd IEEE Symposium on Intelligent Control*, Arlington, Virginia, 1988.

[Touchton and Casella 86] Robert A. Touchton and Mike Casella, "Reactor Emergency Action Level Monitor: A Real Time Expert System," *Instrument Society of America Convention*, 1986.