QUALITATIVE REASONING AND DECISION TECHNOLOGIES
N. Piera Carreté and M.G. Singh (Eds.)

VARIABLE SELECTION AND SENSOR FUSION IN AUTOMATED HIERARCHICAL FAULT MONITORING OF LARGE SCALE SYSTEMS

Alvaro de Albornoz
Dept. Ing. de Sistemas y Automática,
Iniversidad Politécnica de Cataluña,
Av. Diagonal 647,
08028 Barcelona,
SPAIN

François E. Cellier
Dept. of Elec. and Computer Eng.,
University of Arizona,
Tucson, AZ, 85721,
U.S.A.

ABSTRACT

In this paper, the problem of human overload in monitoring the many sensors of a complex industrial plant is addressed. A hierarchical automated fault monitoring and analysis system for high-level decision making is presented. The monitoring system operates in parallel with the traditional channels. It has no effect other than being able to display its findings to the human operator, and to point out potential problems and their perceived causes. The prospects as well as difficulties in realizing such a monitoring system are analyzed by discussing a prototypical implementation of such a system on a sophisticated quantitative large-scale model of a nuclear power plant.

1 INTRODUCTION

Fault diagnosis in large-scale systems is a difficult and controversial issue. Operators of such systems usually insist on being presented with many more sensors and controls than they can safely and reasonably handle. For example, the cockpit of the Space Shuttle contains more than 3000 different sensors and controls. Similarly, the operating room of a nuclear power plant is equipped with thousends of plant status indicators (sensors), and the operators can influence the behaviour of the plant by means of hundreds of different plant set point selectors (controls). It is not reasonable to assume that, in an emergency situation, a small number of human pilots/operators would be capable of reliably monitoring all of these sensors and manually operating all of these controls properly and adequately [10].

Psychological tests have revealed that the average human, after being presented with a number of facts such as given in a news broadcast, can recall approximately 10 of these facts from short-term memory when asked to remember what had been said. While there exists a noticeable variation in individual human capabilities, psychologists tend to agree that most humans cannot reliably and safely tend to more than 10 different items at a time before they begin to misjudge some of the circumstances and make serious mistakes. Consequently, it makes little sense to provide the operator of a complex plant with hundreds of status indicators simultaneously, and expect him or her to monitor them reliably and react to them adequately. In the case of a minor problem, it is quite likely that a lone trouble indicator somewhere on the operator console will go unnoticed for quite some time. In the case of a major disaster, it is very likely that many subsystems will signal problems almost instantaneously, and it will be very difficult for the human operator to discern the true causes from their consequences, i.e., to know which subsystem experienced problems first [13].

Automated fault monitoring can be decomposed into four stages [8] and [9]:

a) *Fault detection*: During this stage, the fault monitoring system detects that the plant behaviour is abnormal.

b) *Fault diagnosis*: During this stage, the fault monitoring system traces observed symptoms back to hypothesized failures that might have caused them.

c) *Fault analysis*: During this stage, the fault monitoring system reasons about possible remedies for the previously diagnosed faults.

d) *Fault reporting*: During this stage, the fault monitoring system reports its findings back to the human operators of the plant.

Any fault monitoring system (FMS) uses a combination of knowledge-based and pattern-based approaches to achieve its goal. Stage (a) of the FMS is naturally pattern-based. It can consist of simple threshold detectors, or time-window detectors [19], or more involved demonized routines called "watchdog monitors" [4].

Stage (b) can be purely pattern-based, e.g. using statistical techniques, or purely knowledge-based, e.g. using a rule-based (expert system) diagnostic engine, or a mixture of both, e.g. using a model-based deep reasoner [13]. Model-based approaches seems to be the most powerful among them. The knowledge can be captured using either deductive techniques employing available meta-knowledge and reasoning on the basis of first principles, or inductive techniques such as neural networks or inductive reasoners.

Stage (c) of the FMS is in all likelihood predominantly knowledge-based. An automated knowledge acquisition system can be used to generate a data base that relates symptoms and failures back to previously successful repair activities [18] and [12]. Stage (d) of the FMS is usually straightforward. More refined systems may carry a model of the human operator [5] to decide on the extensiveness and explicitness of the required fault report.

One of the major problems in FMS design is the possibility of the occurrence of unforeseen faults in a system. A fail-safe FMS must be able to cope with incomplete knowledge [6] and [7]. On these grounds, an inductive reasoning approach is most promising. Inductive reasoners are not yet widely used for such purposes, but they have some striking properties that may make them quite attractive for use in stages (a) and (b) of the FMS [17].

In this paper, the use of a decentralized hierarchical fuzzy inductive reasoning architecture for fault discovery and diagnosis will be presented by demonstrating a prototypical implementation of such a system applied to a sophisticated quantitative large-scale model of a Boiling Water Nuclear Reactor.

2 THE LARGE-SCALE MODEL

The quantitative (numerical) nuclear reactor model is a complex differential equation model containing approximately 500 variables, a model that is very detailed in the nuclear kinetics and in the core thermohydraulics. It includes point kinetics with main feedback mechanisms (Voids, Doppler, and Scram) and with six delayed neutron precursors. The heat generation process is represented by four radial nodes modeled by finite difference approximations

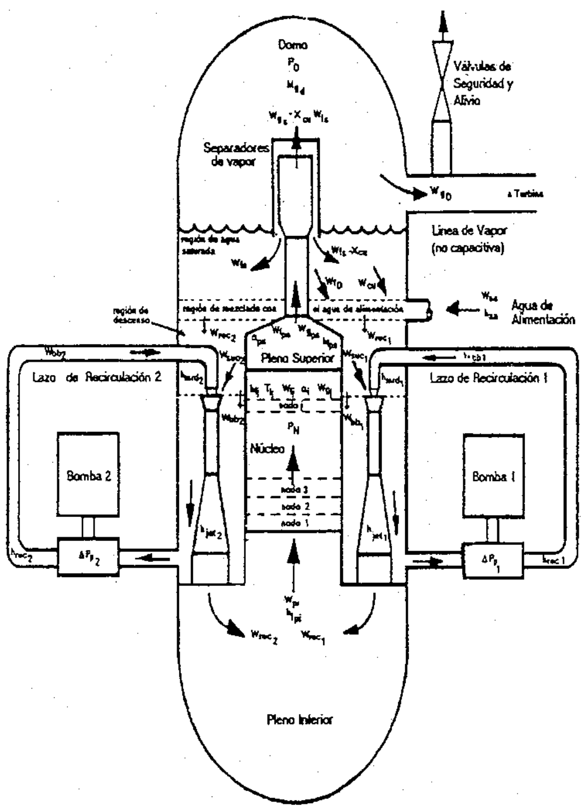

Figure 1: Simplified scheme of the nuclear reactor model

including the residual heat. The thermohydraulics are simulated using a unidimensional axial model that takes into account all boiling phases and two phase flows, and which considers one average fuel element. The heat conduction model has one axial node associated with each axial node of the thermohydraulics and is composed of two radial nodes, one for the fuel and the other for the cladding. The simulation includes the steam separators, the jet pumps, the recirculation pumps, the feedwater pumps, the steam line with all its security valves, and the reactor protection system. Figure 1 shows a simplified model of the reactor vessel.

The plant simulator can be used during all phases of plant operation (start-up, steady-state, and shutdown), and it can also be used for both normal and abnormal plant operation, i.e., during so-called "transients." Details of the nuclear plant simulator were previously published in [15], [16].

3 THE QUALITATIVE MODEL

As stated before, any fault monitoring system uses a combination of knowledge-based and pattern-based approaches to achieve its goals. In the here advocated methodology, a model-based fuzzy inductive reasoner is used. The qualitative representations of the subsystems are predominantly pattern-based models. However, the structural representation that is encoded in the so-called optimal masks represent a form of knowledge. Also, the interactions between the qualitative models of the subsystems and the propagation of information to the higher-level executive reasoners are knowledge-based. In this paper, only a brief description will

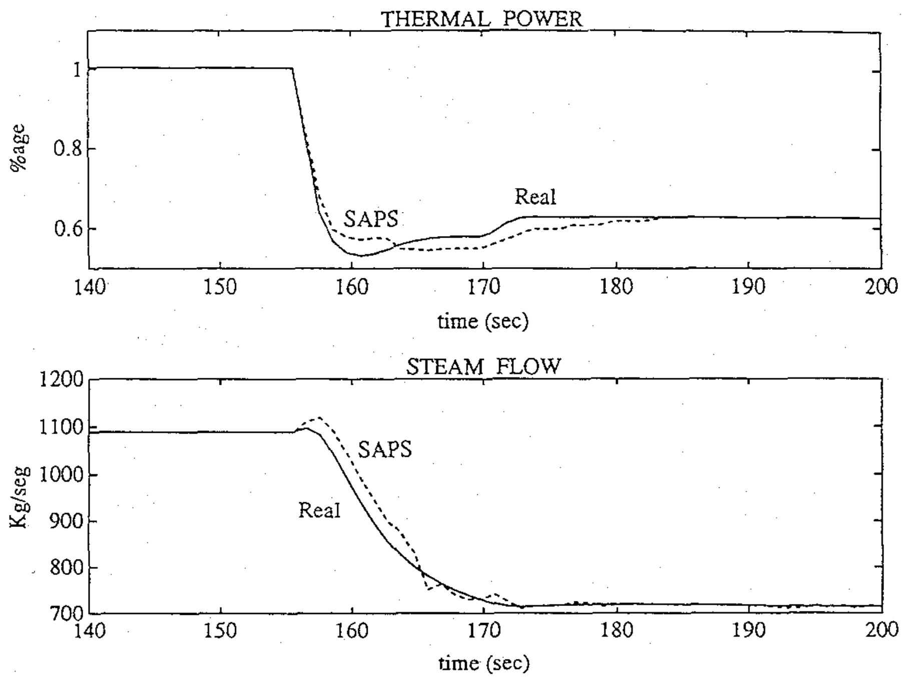

Figure 2: Forecast data using SAPS vs. real data

be given of how the fuzzy inductive reasoners work. Some additional details as well as a complete set of related references are given in a companion paper presented by the same researchers and published in the same volume. The fuzzy inductive reasoning methodology is quite involved, and space considerations do not permit to repeat a detailed description of the technology in every paper published about the subject matter.

An inductive reasoning model, constructed on the basis of measurement data, qualitatively represents the input/output behaviour of the modeled device in the vicinity of an operating point or operating trajectory. A fuzzy inductive reasoning model preserves furthermore numerical information about the plant in the form of fuzzy membership values that can be used to regenerate pseudo-continuous output signals.

Inductive reasoners, like all other qualitative reasoners, base their reasoning on discrete (qualitative) variables. To this end, it is necessary to discretize continuous input signals into discrete (class) values. This process is called "recoding" in the inductive reasoning literature, "classification" in the statistical literature, and "fuzzification" in the fuzzy systems literature. The inductive reasoner then uses the class values of the input variables to infer class values of the output variables. This inference can be performed efficiently, since the search for optimal inference rules is limited to a discrete search space. The fuzzy forecasting algorithm furthermore infers fuzzy membership values for the output variables from the fuzzy membership values of the inputs. The class values of the output variables together with their fuzzy membership values can then be used in a process of defuzzification to regenerate pseudo-continuous output signals.

Figure 2 shows a comparison between the numerical model of a subsystem of the nuclear reactor and the forecast behaviour obtained with the qualitative model during a recirculation pump slowdown transient.

The previously mentioned limitation on the number of input variables that a human can simultaneously process is shared by most automated reasoning algorithms. A single sequential reasoning algorithm turns slow and unwieldy when being requested to cope with too many facts at the same time. Inductive reasoners (like neural networks) may be quite efficient once they are properly trained because they are inherently parallel in nature, but their re-learning abilities degenerate quickly as the number of input variables (i.e., the dimension of their reasoning space) grows. It is therefore very important to select a minimum set of variables that meaningfully represent the system to be reasoned about, and that can be handled by the inductive reasoner in an efficient manner.

3.1 Variable Selection.

Inductive reasoners cannot deal simultaneously with the large number of variables that a large-scale system includes. To solve this problem, a variable selection technique based on optimal mask analysis [2] has being implemented to identify clusters of related variables that can be isolated as subsystems. Structural knowledge of the physical system can also help in this endeavor, but the subsystems that are identified using optimal masks do not necessarily coincide with physical subsystems. The variables making up one subsystem are selected on the basis of similarities in their frequency characteristics and statistical correlation rather than on the basis of geometric topology. Optimal mask analysis allows to identify the behaviour of each subsystem in qualitative terms. The purpose of the optimal masks may simply be to connect the most important of the variables from the input to the output, whereby the selection of the most important variable may be context dependent, i.e., may depend on the qualitative behaviour of any or all of the input variables. If no dominant variable can be identified in a given situation, several inputs may be fused into one output that carries the most important characteristics of all the fused inputs, but filters out some of the less important characteristics. The most important variables of each subsystem can then be propagated to the next stage of the reasoner to determine the qualitative behaviour of the composite system located at the next higher hierarchical level.

Working with a large-scale model of the characteristics shown in section 2 means that the process of properly identifying the subsystems will easily turn out to be one of the most important and difficult problems to be solved on the way of designing the Fault Monitoring System. If the chosen set of variables is not representative enough of the characterized subsystem, the resulting optimal masks will have few interactions among their variables, which in turn will lead to poor propagation of information up the hierarchical ladder. Interestingly enough, the same can be observed if the chosen variables are too strongly correlated, since, in this case, the complexity of the search space is enhanced without augmenting the amount of available information significantly. Also, the selected subsystems should complement each other in an optimal manner. Subsystems that are too independent of each other show few interactions, so that the higher hierarchical levels of the overall architecture don't contribute significantly to the reasoning process, but simply accumulate and propagate further the findings of subordinate reasoners. On the other hand, a duplication of reasoning capabilities within different subsystems located at the same hierarchical level simply enhance the complexity of the search space of the supervisory reasoner without providing it with additional information that would justify the enhancement of its complexity.

The variables and subsystems selected to represent, in a qualitative way, the nuclear reactor are the following:

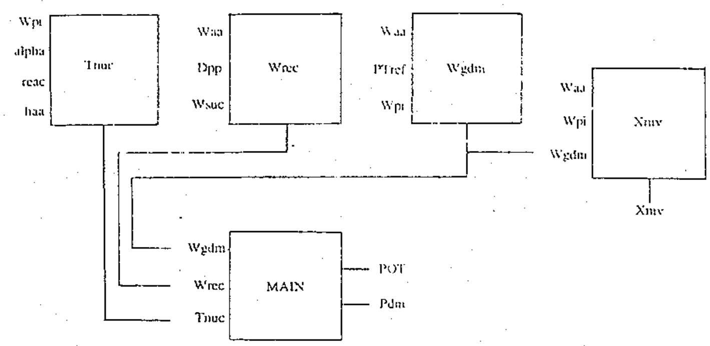

Figure 3: Hierarchical organization of identified subsystems

- **Subsystem** T_{nuc}. Its inputs are: the lower plenum flow W_{pi}, the total reactivity ρ_T, the power fraction α, and the feedwater enthalpy h_{aa}. Its output is the core temperature T_{nuc}.
- **Subsystem** W_{rec}. Its inputs are: the feedwater flow W_{aa}, the pressure drop in the recirculation pump ${}_{pp}$, and the suction flow W_{suc}. Its output is the recirculation flow W_{rec}.
- **Subsystem** W_{gdm}. Its inputs are: the feed water flow W_{aa}, the lower plenum flow W_{pi}, and the reference turbine pressure P_{Tref}. Its output is the steam flow leaving the reactor W_{gdm}.
- **Subsystem** X_{niv}. Its inputs are: the feedwater flow W_{aa}, the lower plenum flow W_{pi}, and the steam flow leaving the reactor W_{gdm}. Its output is the water level inside the reactor X_{niv}.
- **Executive** POT **and** P_{dm}. Its inputs are the outputs of the first three subsystems. Its outputs are: the total percentage power of the reactor POT, and the dome pressure P_{dm}.

Notice that the subsystems are closely related. Some of the input variables are used in more than one subsystem, and one output variable is used to feed more than one module.

3.2 Hierarchical Levels of Masks.

The Fault Monitoring System (FMS) contains a hierarchy of fuzzy inductive reasoners (FIRs). The executive FIR uses as inputs the output signals of the subsystem FIRs. Thus, the role of the subsystem FIRs is that of sensor fusion [14] and [11], i.e., to concentrate the information available through the large number of sensors to a much smaller number of signals that the executive FIR can be expected to handle correctly and efficiently. The executive FIR will report the discovered problem to the operator, while pointing out to him or her, which of the subsystem FIRs is most closely related to the problem. The operator can then turn to that FIR to receive more detailed information.

Each FIR is composed of an optimal mask that represents the behaviour of an identified subsystem. For the nuclear power plant model, four subsystems and one executive subsystem have been identified. These systems are hierarchically organized in the way depicted in figure 3.

As can be seen, there are two different hierarchy levels. The executive FIR uses the outputs of the first three subsystem FIRs. The fourth module, which uses the output of the third subsystem as its input, is located at the same hierarchy level as the executive FIR. It serves as a useful additional discriminatory tool, because the reactor level is characteristic of each transient.

4 THE FAULT MONITORING SYSTEM

There are three major problems that must be addressed when building a FMS. The first problem is the previously stated variable selection and subsystem identification process.

The second problem is an extension of the first. It copes with the variable selection and subsystem identification for the post-accident conditions. Most of the so-called "operational transients" related to a nuclear power plant end with a reactor emergency procedure which may result in a reactor shutdown. Once an emergency procedure has been initiated, there occurs a dramatic change in the meaningful minimum set of variables needed to represent the system, i.e., the valid set of variables used to describe the system prior to the transient is not the same set that is needed to represent the system during and following the transient.

The third problem relates to the excitation of the numerical model needed for the identification of the optimal masks of the subsystems. It is (fortunately!) quite impossible to excite a sophisticated nuclear reactor (or even its quantitative model) in such a way that all frequencies are richly represented in the input/output behaviour of the excited subsystem for the purpose of the best possible identification of an optimal mask for the subsystem. The problem is one that has haunted for decades the researchers who are working in the identification of control systems. When identifying a subsystem within a feedback structure, it is desirable to break the system open, since otherwise, it is never fully clear whether it is really the subsystem itself that has been identified, or whether it might not be the feedback loop around the subsystem that, by its own nature, constitutes another subsystem with the same extraneous variables but exchanged inputs and outputs, or maybe a combination of both. However, inductive models are only valid within a limited range around an operating point or operating trajectory. By opening up the feedback loop, the subsystem is likely to exhibit behavioural patterns that resemble little those of the closed-loop operation. Thus, the "learned" qualitative model will be of little or no use for predicting the behaviour of the subsystem in closed-loop operation. It is thus essential that the feedback loops are kept intact when identifying the subsystems. Consequently, the modeler has to live with the aforementioned difficulties. However, reality is even more grim than that. Nuclear reactors (and their quantitative models) are built for maximum safety. Small deviations in expected behaviour will be interpreted as anomalies that could trigger an emergency process, which may eventually lead to a shutdown of the reactor. The reactor will then behave quite differently due to the transient in comparison with its behaviour under normal operating conditions. The very test signals that are needed to excite the quantitative model for the identification of its qualitative counterpart are easily interpreted as transients by the reactor simulator. The quantitative reactor model trips over its own shoe-lace, so to speak.

The idea of building a fully integrated fault monitoring system for the nuclear reactor was abandoned. Serious transients often call for a decision to initiate an emergency procedure within a few seconds. There is no way that the inductive reasoner could possibly aid the decision making process at such a time scale given the current state of the technology. Since plant safety is and must always be given highest priority, serious transients will invariably lead to the initiation of an emergency procedure. Therefore, while an on-line early warning and transient discovery system is meaningful, an on-line transient characterization system makes little sense in the current state of affairs. Therefore, the FMS was decomposed into two parts:

(1) A semi-continuous fault monitoring system with functions for early warning (some potential problems can be discovered before they become emergencies) and a semi-continuous transient discovery system for quick detection of an evolving emergency. The transient discovery system will often be able to point out which of the subsystems is causing the problem, but will not be able to analyze the precise nature of the anomaly. The on-line system operates exclusively with the optimal masks of the properly functioning plant.

(2) An off-line transient characterization and identification system used for post-mortem analysis. Once an emergency procedure has been initiated, it will take days for the reactor to completely shut down. During this time period, it would be beneficial to have a system that can probe the reactor and reason about the possible causes of the emergency in order to come up with a complete analysis of the emergency as early as possible.

4.1 The Semi-Continuous Monitoring System

Due to the reasons stated in the preceding section, we will simulate what is known in the nuclear terminology as *small operational transients*, in which the variables selected for each identified subsystem remain unchanged before, during, and after the incident.

The FMS will detect that a transient is taking place because the optimal masks no longer represent correctly the behaviour of the system. An error threshold alarm matrix detects that the executive FIR's forecasting process contains too many errors. The comparison between the numerical variables and the forecast variables provide the values for the error matrix. The alarm matrix reads those values and triggers the alarm if a combination of consecutive incorrect forecasts and saturated states occurs. Once the executive FIR's alarm is triggered, the FMS proceeds downward to the next hierarchical level and checks the alarm matrices of the subsystems to try to determine which of them might have caused the transient to occur. Once a transient has been detected, the FMS will stop its regular monitoring activity, and will start to collect data for the post-mortem analysis.

If a single subsystem FIR triggers an alarm that is not picked up by the executive FIR's alarm matrix as well, the FMS is facing a small failure that eventually could cause the executive FIR to start the general alarm. This failure should be presented to the operators as an early warning to avoid the possibility of an operational transient later on, that might then lead to an unnecessary reactor shutdown.

The transient selected to demonstrate, in this paper, the detection capabilities of the FMS is a Recirculation Pump Slowdown, an incident known to be a small power transient, i.e., a transient that, at least initially, does not trigger an emergency procedure. This transient produces a reduction in the recirculated flow of aproximately 50% with an inertial time of 8 to 10 seconds, and the following effects in the subsystems variables:

$$
\begin{array}{c}
step\backslash^{var} \\
150 \\ 151 \\ 152 \\ 153 \\ 154 \\ 155 \\ 156 \\ 157 \\ 158 \\ 159 \\ 160 \\ 161 \\ 162 \\ 163 \\ 164
\end{array}
\begin{array}{c}
\begin{array}{cc} errorPOT & errorP_{dm} \end{array} \\
\begin{pmatrix}
0 & 0 \\
0 & 0 \\
0 & 1 \\
0 & 0 \\
0 & 0 \\
0 & 0 \\
0 & 0 \\
2 & 0 \\
1 & 0 \\
-1 & 2 \\
-1 & 0 \\
1 & -1 \\
2 & 1 \\
0 & -1 \\
-2 & 2
\end{pmatrix}
\end{array}
\qquad
\begin{array}{c}
alarm \\
\begin{pmatrix}
0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1
\end{pmatrix}
\end{array}
$$

Figure 4: Error and alarm matrices of the executive FIR.

- A reduction in the forced water flowing into the core, and consequently, a reduction in the steam flow leaving the reactor.
- A reduction of the reactor pressure due to a low steam flow.
- A power decrease from 100% to 65%.
- An increase of the core temperature due to a reduction in the refrigerant flow.
- Oscillations of the reactor water level.

4.2 The Transient Characterization System

The off-line transient characterization system works in similar ways as the on-line fault monitoring system. Once the transient has begun, the FMS will consult the transient library for a new set of optimal masks that describes the transient behaviour. Once found, these optimal masks will drive the qualitative model not just through the transient, but into the post-accident steady state. However, rather than operating in parallel and in real–time with the quantitative model, the transient characterization system operates on previously collected data in a post-mortem mode of operation. Real-time considerations are of no importance here, and consequently, it is possible to maintain an extensive list of transient and post-transient optimal masks in the library.

5 RESULTS

A recirculation pump slowdown was initiated, in the numerical model, at time step 155. As can be seen in figure 4, the POT (thermal power) variable of the executive FIR error matrix detects the anomaly two time steps later, however, no alarm is trigger since the other variable P_{dm} (reactor pressure) reacts very slowly in the recirculated flow. Thus, the threshold built into the construction of the alarm matrix is not immediately reached by the observed anomalv.

step\var	$errorW_{rec}$	$alarm$
150	0	0
151	0	0
152	0	0
153	0	0
154	0	0
155	0	0
156	−1	0
157	2	0
158	−1	0
159	2	1
160	−1	1
161	1	1
162	1	1
163	1	1
164	−1	1

Figure 5: Error and alarm matrices of the W_{rec} subsystem.

The transient was first detected by the W_{rec} (recirculated flow) subsystem. This subsystem is directly related to the transient through the $_{pp}$ (pressure drop) and W_{suc} (suctioned flow) variables in the recirculation circuit. In the given situation, the W_{aa} (feedwater flow) variable did not affect the reaction of the module in a significant way. Figure 5 shows that once three incorrect forecasts have been detected, i.e., three consecutive errors have occurred in the error matrix, the alarm is triggered. At this moment, the subsystem has discovered the beginning of a malfunction. If the transient would have stopped at this point, i.e., the magnitude of the reduction in the recirculated flow would have been of such a limited magnitude that no other subsystem alarm were triggered as well, the FMS would have reported its finding to the operators in the form of an early warning message.

The subsystem W_{gdm} (steam produced) detects the malfunction through the W_{pi} variable. The lower plenum flow is the sum of W_{suc} and W_{rec}. It can be observed in figure 6 that the incident is first detected at time steps 157 and 158, and then an obviously bad forecast accidentally produced a "good forecast" (a 0 error condition), still impedeing the three consecutive errors needed to start the general alarm, which is finally triggered at time step 163. At this instant, the executive FIR detects that two of its subsystem FIRs have triggered their alarms, and proceeds to trigger its own alarm.

Figure 7 shows how the alarm of subsystem T_{nuc} (fuel temperature), that includes also the W_{pi} variable, remains untriggered during the 10 seconds post-accident monitoring time. This is due to the influence of its other three input variables α, ρ_T, and haa, which are weakly coupled to the recirculation circuit. The fuel temperature does decrease, but at a very small rate. If the on-line monitoring system were allowed to continue its operation beyond the time where the general alarm is triggered, it could be seen that a T_{nuc} subsystem alarm would occur at time step 179.

The behaviour of the simulated reactor subsystems is a good aproximation of what would be observed in a real reactor under similar circumstances. In a BWR nuclear reactor confronted with such a transient, W_{rec} drops dramatically down to even negative values, i.e., reverse flow in the recirculation circuit. The power POT reacts immediately while the pressure P_{dm} and the produced steam W_{gdm} decrease much more slowly. The temperature of the core remains almost unchanged during the whole transient, a fact that is reflected by the behavior of the third subsystem in the simulation. It can be noticed that this was the last subsystem to trigger an alarm.

$$
\begin{array}{c c}
{}_{step}\backslash^{var} & errorW_{gdm} \\
\begin{array}{c}150\\151\\152\\153\\154\\155\\156\\157\\158\\159\\160\\161\\162\\163\\164\end{array} &
\left(\begin{array}{c}0\\0\\0\\0\\0\\-1\\0\\2\\-1\\0\\2\\2\\-2\\1\\1\end{array}\right)
\end{array}
\qquad
\begin{array}{c}
alarm \\
\left(\begin{array}{c}0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\1\\1\end{array}\right)
\end{array}
$$

Figure 6: Error and alarm matrices for subsystem W_{gdm}.

$$
\begin{array}{c c}
{}_{step}\backslash^{var} & errorT_{nuc} \\
\begin{array}{c}150\\151\\152\\153\\154\\155\\156\\157\\158\\159\\160\\161\\162\\163\\164\end{array} &
\left(\begin{array}{c}0\\0\\0\\1\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\end{array}\right)
\end{array}
\qquad
\begin{array}{c}
alarm \\
\left(\begin{array}{c}0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\\0\end{array}\right)
\end{array}
$$

Figure 7: Error and alarm matrices for T_{nuc}.

6 CONCLUSIONS

In this paper, a prototypical implementation of a hierarchically structured fault monitoring system for nuclear reactors based on fuzzy inductive reasoning was presented. It was indicated that such a fault monitoring system can be meaningfully used in two different modes: (i) as an on-line fault detection and early warning system, and (ii) as an off-line post-mortem fault diagnosis and analysis system. The functioning of the former of these two FMS applications was furthermore demonstrated with a concrete example.

The difficulties in constructing such an FMS were also stated. They relate to the problems of identifying the subsets of variables to be used in subsystems, and the difficulties that stem from the need to identify subsystems in a closed-loop environment. These problems have by no means all been overcome yet. While we were able to successfully identify a set of subsystems that was able to recognize the envisaged scenario of a recirculation pump slowdown, we cannot claim that we have solved the problem of building even a prototypical FMS that could be used to discover and report a wide palette of different operational transients. Moreover, the algorithm used to identify the subsystems is still utterly experimental and by no means fully automated.

The research effort will continue. While there remain many problems yet to be addressed, we are quite excited about the possibilities of the chosen approach. We believe strongly that fuzzy inductive reasoning offers great opportunities in tackling a number of "hot" issues in A.I., difficult issues that other A.I. techniques have not been able to master so far and that have hampered the progress in A.I. research.

References

1. [Albornoz and Cellier, 1993] Albornoz, A. and F.E. Cellier, "Qualitative Assessment of the Behaviour of an Aircraft Through Fuzzy Inductive Reasoning," QUARDET'93, Barcelona, Spain, 1993.

2. [Cellier, 1991] Cellier, F.E. *Continuous System Modeling*, Springer-Verlag, 1991, N.Y., U.S.A.

3. [Cellier, et al, 1992a] Cellier, F.E., A. Nebot, F. Mugica and A. Albornoz, "Combined Qualitative and Quantitative Simulation Models of Continuous-Time Processes Using Fuzzy Inductive Reasoning Techniques," in: Proceedings of the IFAC SICICA'92 Symposium, Málaga, Spain.

4. [Cellier, et al, 1992b] Cellier, F., A. Doser, G. Farrenkopf, J. Kim, Y.D. Pan, L.C. Schooley, B. Williams and B.P. Zeigler, "Watchdog Monitor Prevents Martian Oxygen Production Plant from Shutting Itself Down During Storm," in: Proceedings of the ISRAM'92 ASME Conference, Santa Fe, N.M., U.S.A.

5. [Cacciabue, 1990] Cacciabue, P.C., G. Mancini and U. Bersini, "A Model of Operator Behaviour for Man-Machine System Simulation," *Automatica*, **26**(6).

6. [Davis, 1982] Davis, R., "Diagnosis Based on Structure and Function," in: Proceedings of the National Conf. on Artificial Intelligence, Pittsburgh, Penn., U.S.A., pp. 137-142.

7. [Davis, 1984] Davis, R., "Diagnostic Reasoning Based on Structure and Behaviour," *Artificial Intelligence*, **24**.

8. [De Kleer and Williams, 1987] De Kleer, J. and B. Williams, "Diagnosing Multiple Faults," *Artificial Intelligence*, 32(1), pp. 97-130.

9. [De Kleer and Williams, 1989] De Kleer, J. and B. Williams, "Diagnosis with Behavioural Modes," in: Proceedings of the 11th Intl. Conf. on Artificial Intelligence, Detroit, Mich., U.S.A.

10. [Dvorak and Kuipers] Dvorak, D. and B. Kuipers, "Model-Based Monitoring of Dynamic Systems," in: Proceedings of the 11th Intl. Conf. on Artificial Intelligence, Detroit, Mich., U.S.A., pp. 1238-1243.

11. [Luo and Kay, 1989] Luo, R.C. and M.G. Kay, "Multisensor Integration and Fusion in Intelligent Systems," *IEEE Trans. Systems, Man and Cybernetics*, **19**(5).

12. [Motaabbed, 1992] Motaabbed, A., *A Knowledge Acquisition Scheme for Fault Diagnosis in Complex Manufacturing Processes*, MS Thesis, Dept. of Electrical & Computer Engineering, University of Arizona, 1992.

13. [Pan, 1986] Pan, J., "Qualitative Reasoning with Deep-Level Mechanism Models for Diagnose of Failures," in: Proceedings of the CAIA, Denver, Colo., U.S.A, pp. 295-301.

14. [Pau, 1989] Pau, L., "Knowledge Representation Approches in Sensor Fusion," *Automatica*, **25**(2).

15. [Ramos, 1991] Ramos, J.C., *Simulacion en Tiempo Real de la Termohidraulica de un Reactor de Agua Hirviente*, MS Thesis, Instituto Politecnico Nacional, Mexico.

16. [Ramos and Albornoz, 1988] Ramos J.C. and A. de Albornoz, "A One Dimensional Model for the Laguna Verde Boiling Water Nuclear Reactor," in: Proceedings 1988 SCS Multiconference on Power Plant Simulation, San Diego, Calif., U.S.A.

17. [Vesanterä and Cellier, 1989] Vesanterä, P.J. and F.E. Cellier, "Building Intelligence into an Autopilot Using Qualitative Simulation to Support Global Decision Making," *Simulation*, **52**(3), 1989.

18. [Wang, 1991] Wang, Q., "The Development of a Diagnostic Expert System for Equipment Maintenance," in: Proceedings of the AMD Technical Conference, 1991, San Francisco, Calif., U.S.A.

19. [Wang and Cellier, 1991] Wang, Q. and F.E. Cellier, "Time Windows: Automated Abstraction of Continuous-Time Models into Discrete Events Models in High Autonomy Systems," *Intl. J. General Systems*, **19**(3), pp. 241-262.