Fault Detection and Diagnosis using Qualitative Modelling and Interpretation

J. M. Vinson and L. H. Ungar

Department of Chemical Engineering, University of Pennsylvania,
Philadelphia, PA

Abstract. The Qualitative Modelling and Interpretation (QMI) system translates noisy sensor data into a qualitative description of the underlying behavior of a chemical plant and uses this information together with qualitative models to identify faults and operating regimes. Qualitative models of normal and faulty equipment are simulated to describe the range of possible behaviors in a chemical plant without the need for exact numeric models which are unavailable for many faults. Sensor data are then used to select between different models. Simultaneously using interpretations from multiple sensors reduces sensitivity to sensor noise, increasing diagnosis reliability. QMI has been implemented on a simulated propylene glycol reactor with good results.

Keywords. Failure detection, identification, leak detection, monitoring, qualitative simulation.

INTRODUCTION

Qualitative Modelling and Interpretation (QMI) is an on-line process monitoring system which combines quantitative change detection with qualitative modelling to diagnose faults in chemical plants. QMI translates sensor readings to a qualitative description where each variable is characterized by its relation to landmark values (e.g. zero, boiling point, equilibrium mole fraction) and the sign of its derivative. These sensor interpretations are compared to simulations of multiple qualitative models in order to find the model which matches the observed behavior.

Many process monitoring systems have been developed which detect and diagnose potentially costly or dangerous situations. Two common numerical methods are change detection, typically using alarm thresholds, and model selection. Because incoming data from sensors is generally noisy, alarm thresholds for single sensor detection methods must be set to balance between quick detection with frequent false alarms and infrequent false alarms with slower detection. Traditionally, alarm thresholds are set to eliminate false alarms. More sophisticated alarm threshold techniques include Shewhart and cumulative sum (CUSUM) control charts which are statistical tests for the determination of a change from expected behavior (Bowker, 1972; Box, 1970; Shewhart, 1931).

In contrast, model selection methods fit incoming data from single or multiple sensors to detailed plant models via techniques such as Kalman filtering or non-linear programming (Frank, 1990; Isermann, 1984; Willsky, 1976). The model which best fits the data is considered to be the current model of the plant, and thus the diagnosis. To achieve a wide coverage of faults and operating conditions, quantitative models must be built which describe all possible situations, including several magnitudes of the same type of faults, such as 10%, 40% and 80% leaks. Unfortunately, quantitative models are often unavailable for changes which take the plant into unknown or transient operations, such as during start-up and shutdown procedures.

Qualitative monitoring and diagnosis systems avoid the need for quantitative models by describing plant data in qualitative terms: values may be high, normal or low, and trends may be increasing, steady, or decreasing. This qualitative interpretation is used to determine the causes of changes in the plant. Rule-based or expert systems map these symptoms directly to causes via rules obtained from process experts (Ramesh, Davis and Schwenzer, 1992; Venkatasubramanian and Rich, 1988). Rule-based systems do not require a plant model, however they do require a multitude of rules to cover all possible faults in a chemical plant and have difficulties with unexpected operations or new equipment. In contrast, model-based diagnosis systems examine a qualitative plant model to determine what changes (faults) in the model could have caused the observed symptoms. Recently, several investigators have started combining numerical and qualitative methods, either by using quantitative information to constrain qualitative models (DeCoste, 1990; Dvorak and Kuipers, 1991; Berleant and Kuipers , 1990; Forbus, 1987), or by mixing qualitative and quantitative models to get a better description of the physical system (Forbus and Falkenhainer, 1990; Oyeleye, Finch and Kramer, 1990; Yu and Lee, 1991).

The Qualitative Modelling and Interpretation (QMI) system combines both detection and diagnosis into one package which borrows from both numerical and qualitative methods. Multiple sensors are analyzed and compared to known qualitative models of the plant, and that model which fits best is taken to be the current model of the plant. Tighter tolerances can be set on individual sensors because QMI filters out aphysical interpretations which may be suggested in single sensor methods such as alarm thresholds and control charts, allowing earlier diagnosis of process disturbances without causing more false alarms. Use of qualitative models also permits diagnosing a range of faults and behaviors for which exact equations are not known. QMI has been tested on a simulated continuous stirred tank reactor and gives relatively rapid and accurate diagnoses.

QUALITATIVE MODELS

In qualitative modelling, an entire class of faults is described with a single model, rather than by several quantitative models with different numeric values of parameters. This approach is useful when the magnitude of the fault is unknown and plant behavior may vary with the size of the fault. For example, a "small" leak from a tank may be compensated for by a controller, but a larger leak may cause the controller to saturate.

Qualitative models use qualitative equations to describe how variables change with respect to one another (e.g. in laminar flow, flow rate is proportional to pressure drop). Variables are described by the sign of their slope and relation to landmarks (e.g. mole fraction has the landmarks of zero, one and perhaps equilibrium). The description of all the variables in the system is called the qualitative state of the system (e.g. water temperature is increasing and between freezing and boiling points, and the amount of water is steady and positive). QMI uses the QSIM qualitative simulation package (Kuipers, 1986) to "solve" the qualitative models, producing a "tree" of behaviors for the models, as shown in Fig. 1. Each behavior or path through the tree is a series of qualitative states.

For example, in Fig. 1, State 1 is the steady state operation of a level controlled, water cooled CSTR and states 2 and 3 are induced changes to the system with the subsequent states being the reaction to these changes. State 2 is the onset of a tank leak, where temperature and concentration are still steady, but tank level and outlet flow have begun to decrease in response to the leak; and state 3 is a change in the inlet flow rate. State 4 describes the qualitative state where temperature and concentration have begun to react to the change in reactor level, with temperature decreasing and reactant concentration increasing. State 7 is the qualitative state where the controller has begun to react to the level change, but the level and outlet flow rates are still decreasing. Behaviors branching from state 7 include return to set point or saturation of the controller.

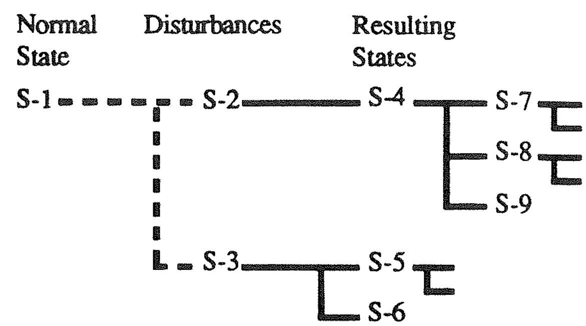

Fig. 1. Sample behavior tree

Qualitative modelling of chemical plants is an area of active research. For example, Dalle Molle and colleagues (1988, 1989) have modelled chemical reaction systems and proportional-integral level control using QSIM. Vinson, Grantham and Ungar (1992) have modelled a chemical plant consisting of heat exchangers, a reactor, a stripper and a condenser using the Qualitative Process Theory framework (Forbus, 1984), and Oyeleye and colleagues (1990) have modelled a jacketed CSTR for application in MIDAS, a fault diagnosis system. When building models of more complex systems the number of solutions can get quite large, thus care must be taken to limit the number of solutions without eliminating real behaviors (Kuipers and Chiu, 1987; Catino, 1991).

QMI ALGORITHM

The QMI algorithm conducts two operations on-line: 1) It translates noisy sensor data into a qualitative description, and 2) it compares this interpretation to the behaviors predicted by the qualitative models. Qualitative simulation of the models is done beforehand, so the behavior tree is always available to QMI. The algorithm is shown in Fig. 2.

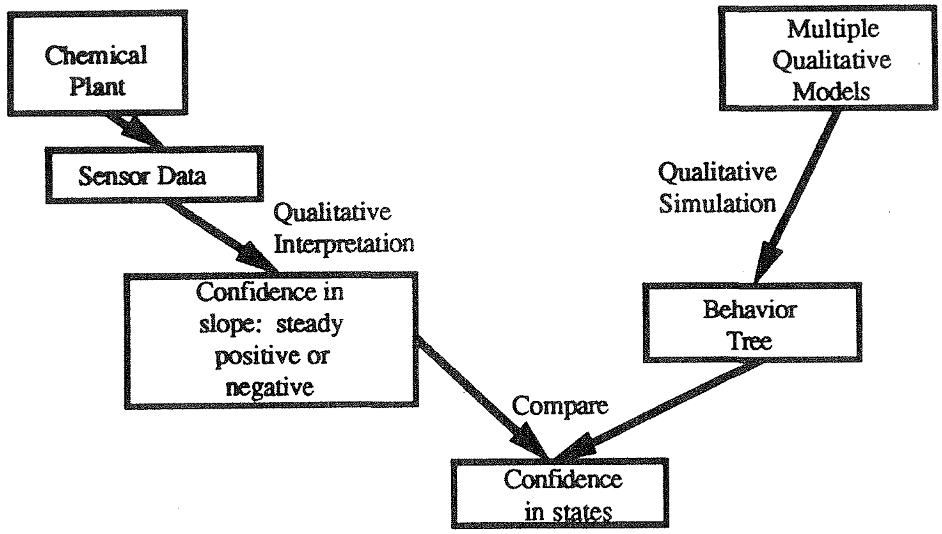

Fig. 2 Qualitative Modelling and Interpretation Algorithm

Qualitative interpretation of the incoming data is done by finding the best fit line through the incoming data with linear least squares. The estimate of the slope is then given a measure of belief (zero - one) that it is positive, zero or negative via the ad-hoc functions:[1]

$$p(\text{inc}) = \frac{1}{1 + e^{\alpha * \left(1 - \frac{\text{slope}}{\text{cutoff}}\right)}}$$

$$p(\text{dec}) = \frac{1}{1 + e^{\alpha * \left(1 + \frac{\text{slope}}{\text{cutoff}}\right)}}$$

$$p(\text{std}) = 1 - p(\text{inc}) - p(\text{dec})\,,$$

where α is used to adjust the slope of the curves and cutoff is the inflection point. Large values of α (> 10) yield on-off functions which trip at the cutoff value, and small values (~ 1) give flat curves. We use $\alpha = 5$ in our simulations. The cutoff is set from knowledge of the noise in the signal so that a value of p(std), p(inc) or p(dec) greater than 0.5 indicates that the current interpretation is probably correct.

In order to speed detection and eliminate false alarms, the beliefs in the interpretations of several sensors are multiplied to yield an overall belief in the interpretation of the state of the plant. In the reactor example above, one of the qualitative states has temperature increasing, concentration decreasing and both level and outlet flow steady. If the interpretation of these signals is p(inc,temperature) = 0.8, p(dec,concentration)=0.7, p(std,level) = 0.75 and p(std,flow) = 0.65, then the overall belief in the qualitative state is 0.273. We take a belief greater than 0.5^N to be a significant belief in the qualitative state, where N is the number of sensors. Combining the sensor interpretations allows QMI to overlook sensors which appear to be changing (due to noise) while the others remain normal. Similarly, QMI is less sensitive to sluggish sensors when the other sensors are changing as expected from a disturbance.

Initially, QMI assumes the system is behaving normally and only compares the qualitative interpretation to the normal state and the first states of the disturbances in the behavior tree. As more data comes in from the plant the qualitative interpretations of the signals change due to noise and process disturbances. If a disturbance occurs in the plant, the qualitative interpretation will change to produce a high confidence in the initial state of the disturbance. With the next set of data QMI will look for interpretations which match this state and those following it. Diagnosis is attained when one disturbance has a belief that is much higher than the others.

EXAMPLE SYSTEM AND TEST CASES

The QMI algorithm has been tested with a simulated propylene glycol production reactor pictured in Fig. 3, from Fogler (1986). Equal volumes (46.62 cfh each) of propylene oxide and methanol are mixed with water (233.1 cfh) containing 0.1 wt % H_2SO_4 which catalyzes the reaction. The PO-MeOH and water feed streams are initially at 58°F and mixing increases their temperature

---

[1] We have examined the use of statistical analysis techniques in determining the confidence in the sign of a given slope, but have found them to be limited in that there is only one hypothesis to test; whether the slope is zero. In statistics, there is no ability to test for a number being greater than zero without specifying the value which is "greater than zero."

to 75°F. The reaction is exothermic and propylene oxide is relatively low-boiling, so reactor must be cooled to keep the temperature below 125°F to prevent excessive vaporization. A proportional controller adjusts the outlet flow rate to prevent the level from changing drastically due to changes in supply or demand. The numerical simulation was written in FORTRAN and uses the LSODE numerical integration package, based on the Gear method. Possible faults in the simulated reactor are changed feed flow, tank leak, changed acid concentration, changed feed temperature, and changed reactant feed concentration.

Qualitative simulation of these faults with the qualitative model for the reactor produces a behavior tree, similar to Fig. 1, with nearly 400 histories and about 2000 qualitative states. As in the example behavior tree in the Qualitative Models section, there are many histories produced by qualitative simulation, rather than one for each fault. These histories often differ in the relative time of occurrence of of different events. For example, when temperature and concentration both have inverse response, qualitative reasoning cannot determine which will peak first.

QMI was compared to a version of QMI which has Boolean equations for beliefs in the sensor interpretations and to a basic alarm detection method using single-sensor thresholds. Twelve single-fault simulations were conducted, and these three methods were compared for diagnosis time, number and duration of false alarms, and number and duration of missed diagnoses. The diagnosis time indicates how long it takes to properly identify the fault. False alarms occur when the beliefs in incorrect qualitative states become higher than all the others and is due to noise in the signal. Missed diagnoses arise when the probability of the normal state drops below 0.5^N (number of sensors), and none of the other qualitative states matches the interpretation. This is caused by a combination of noise and slower signals, such as temperature change in a large tank.

In all the test cases, QMI observes four sensors: tank temperature, tank concentration, tank level, and outlet flow rate. Figures 4 a, b, c, and d show the time plots of each of these sensors for a 0.25% decrease in the inlet reactant concentration at 1.00 hour. In Fig. 4 a, the tank composition increases to a new steady state value, but the initial response is a decrease in the tank composition, which is difficult to detect visually but is an important characteristic of this change. Temperature decreases to a new steady state and the tank level and outlet flow rate remain constant.

As this data comes from the plant, QMI calculates beliefs that each measurement is increasing, steady or decreasing; these beliefs are shown in Fig. 5. The top graph shows the interpretation of the concentration signal: above zero is the belief that the signal is increasing and below, decreasing. This shows an initial belief that the concentration is increasing, which would cause an alarm to be sounded with Boolean alarms, but comparison with the other signals shows that temperature and level are both steady while outlet flow may be decreasing (belief 0.5). This does not correlate with any qualitative state, so the normal state is still valid, although it has a low confidence. Alarms and Boolean QMI signal a change here, although no diagnosis is proposed. The ad-hoc functions capture the inverse response of the concentration at 1.105 hours, at which time diagnosis of the fault occurs. The Boolean version of QMI does not capture the fault until the slope of the concentration crosses the 0.5 belief line at 1.140 hours.

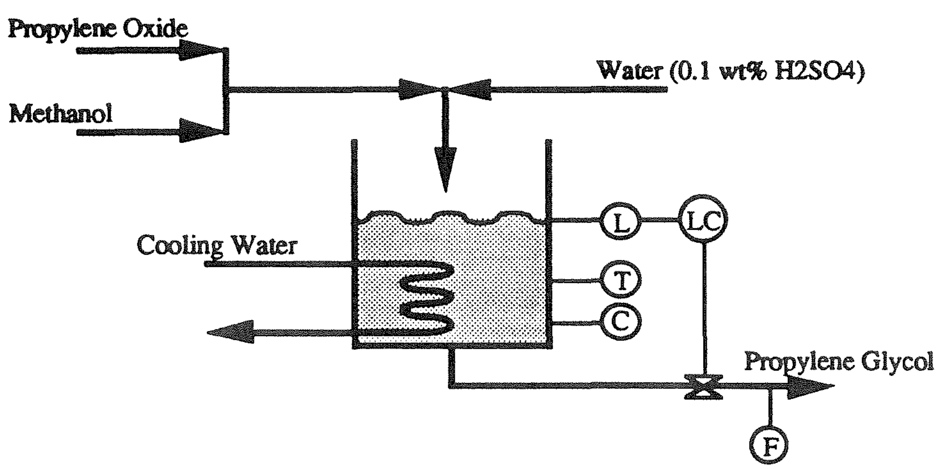

Fig. 3 Propylene Glycol Reactor

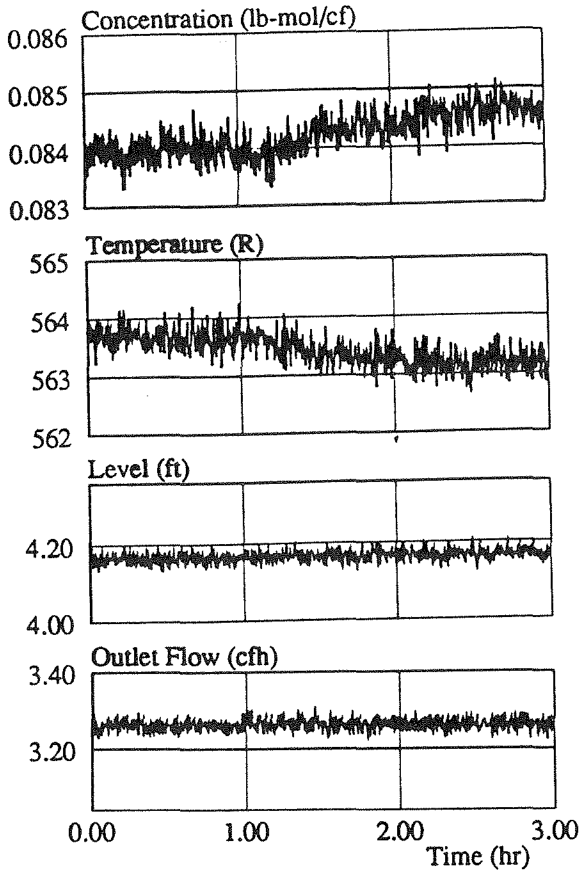

Fig. 4 a, b, c, d Concentration, Temperature, Level and Outlet Flow

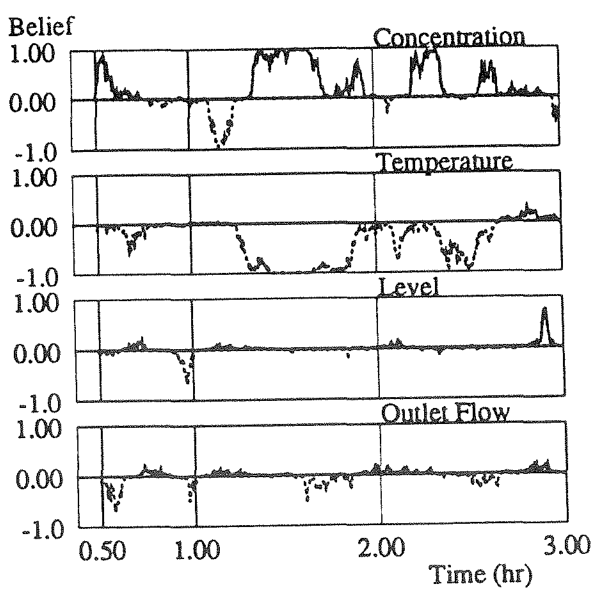

Fig. 5 Beliefs for Each Sensor

The average statistics for the test runs are presented in Table 1. The thresholds in these examples were intentionally set tightly to force false alarms in the detection and diagnosis systems. For the Boolean QMI tests, there are two cases where diagnosis is never achieved because the magnitudes of sensor noise and the change are nearly the same, so those statistics are for ten examples. QMI is able to reach a diagnosis in these cases through its use of the belief functions. Diagnosis time is reported in minutes with standard deviation in brackets. Alarms, of course, do not give diagnoses. For false alarms and missed diagnoses, Table 1 reports the average number of false alarms and their average duration. The difference between diagnosis time for QMI and Boolean QMI is negligible, but diagnosis time for the alarm threshold tests is faster than both because it is detection time and only one sensor has to change from the normal state to detect a change.

TABLE 1 Statistics for All Examples

	QMI	Boolean QMI	Alarm Thresholds
Correct Diagnosis	12	10	NA
Diagnosis Time [a] minutes	8.1 (4.5)	9.4 (5.2)	4.2 (3.8)
False Alarms	0.5	1.4	2.2
Duration [b]	0.8	3.9	21.
Missed Diagnoses	0.4	1.8	NA
Duration [b]	4.3	30.	

[a] Diagnosis time for alarm thresholds is change detection time.
[b] Duration is measured in sampling intervals.

QMI outperforms Boolean QMI by far in number of false alarms and missed diagnoses. With the excessively tight thresholds used here for illustrative purposes, QMI has one false alarm or missed diagnosis for every two examples, whereas there are three times as many with Boolean QMI. Hard thresholds also average two false alarms per test. The duration of these false alarms is also much greater for Boolean QMI and hard thresholds. With looser thresholds the number of false alarms would decrease at the cost of diagnosis speed. We expect that reasonable thresholds in QMI would still be unsatisfactory for Boolean QMI or alarm thresholds.

Related Work and Current Research

QMI is similar to two other qualitative monitoring and diagnosis systems: DATMI (DeCoste, 1990) and MIMIC (Dvorak and Kuipers, 1991). Like them, QMI attempts to use qualitative models to determine the underlying behavior of a system based on information provided by its sensors. DATMI and MIMIC are in several ways more sophisticated than QMI: both those systems are better than QMI at using that fact that instantaneous states are often not observed.

Rather than predefining all possible fault models as in QMI, MIMIC hypothesizes changes to the model when it can no longer track observations. MIMIC allows only single-change hypotheses, but subsequent discrepancies between readings and predictions may cause MIMIC to hypothesize more changes. Currently, these model changes are only changes landmarks of operating parameters with new quantitative ranges (i.e. change inlet flow from "normal" at 9 to 10 gallons per minute to "low" at 5 to 6 gpm). Each hypothesis generated in this manner is simulated and compared to the current observations, and those hypotheses that match become the set of candidate models which are tracked by MIMIC. The algorithm for MIMIC is shown in Figure 6.

QMI simulates the entire behavior tree, but MIMIC simulates behaviors on-line in a step-by-step fashion. This allows the above fault generation method to be used as well as providing an efficient way to incorporate quantitative information into the qualitative models for use with semi-quantitative simulation (Kuipers and Berleant, 1988). Semi-quantitative simulation uses ranges around the variables and envelopes around proportional constraints, providing a prediction of the possible quantitative range of values for variables as well as the usual prediction of qualitative direction provided by QSIM. This allows MIMIC to better compare predictions to observations, providing faster detection of changes in the behavior of the plant.

QMI and MIMIC both use a membership function to calculate the degree of belief that a state matches the readings. In QMI only the slope of the reading is examined, whereas MIMIC looks at both the reading and its slope. The current implementation of MIMIC is unable to handle noise because the readings are assumed to be perfect measurements of the plant.

The key feature of QMI which is lacking in DATMI and MIMIC is the explicit recognition that processes and sensors have noise and that this noise obeys statistical laws which allow one to ascribe probabilities to beliefs in interpretations of both individual sensor readings and complete qualitative states. The QMI work described above uses an ad hoc interpretation inspired by fuzzy logic and the notion of a "soft threshold" dividing different qualitative regions. We are now moving toward the use of more rigorous statistics to handle noise. This will be particularly important as we increasingly use semi-quantitative simulation to aid fault detection in dynamic systems.

This comparison has led us to form QMIMIC, a hybrid of QMI and MIMIC. The new system is called QMIMIC and is being built to handle complex physical systems such as the CSTR example of this paper. QMIMIC retains the incremental simulation of MIMIC

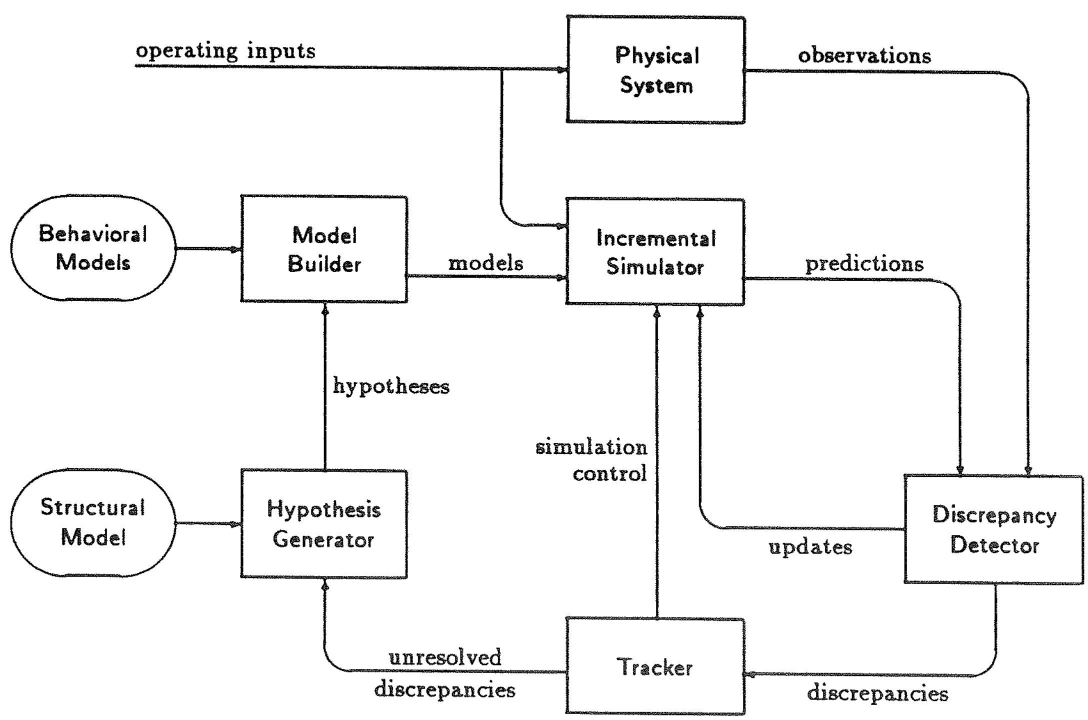

Fig. 6 Architecture of Mimic. The rectangular boxes represent processing elements and the labeled lines show information flow

for semi-quantitative prediction and fault generation purposes. From QMI it borrows and enhances the ability to deal with uncertainty in the observations. Readings are now compared to semi-quantitative predictions by statistical tests, rather than ad-hoc measures of certainty. Also, QMIMIC does not lock prematurely into a diagnoses caused by sensor noise, as does MIMIC and the Boolean versions of QMI.

CONCLUSIONS

Qualitative Modelling and Interpretation provides an interesting addition to the field of fault detection and diagnosis. Sensor signals are individually translated into a qualitative description with associated beliefs which are then combined for an overall belief in the qualitative state of the plant. Combining multiple sensors allows QMI to set tighter thresholds on single sensors because changes in sensors must correlate with each other to produce a diagnosis of faults or process changes. QMI performs better than either a single sensor alarm detection method or a version of QMI with Boolean belief functions in diagnosis time, false alarms and missed diagnoses. We are currently augmenting QMI to use probabilities of different failures.

A major goal of this work has been to integrate the qualitative information available in the real world of chemical plants and robots with qualitative models of that world. In particular, we think that it is important to explicitly model the noise that comes from disturbances to these systems, as well as the noise and inaccuracy that come from the sensors used to measure variables such as flow, temperature, concentration and pressure and to use statistical characterizations of that noise in the diagnosis process.

ACKNOWLEDGMENTS

This research has been partially supported by NSF Presidential Young Investigator award CBT 86-57899. This work originally appeared in a modified form in the Proceedings of the IFAC Symposium "On Line Fault Detection in the Chemical Process Industries."

REFERENCES

Berleant, D. and B. J. Kuipers. (1990). Qualitative-quantitative simulation with Q3. Proc. Qualitative Physics Workshop.

Bowker, A. H. and G. J. Lieberman. (1972). Engineering Statistics, 2nd Ed. Prentice-Hall, Inc., Englewood Cliffs, NJ.

Box, G. E. P. and G. M. Jenkins. (1970). Time Series Analysis. Holden-Day, San Francisco.

Catino, C. A., S. D. Grantham and L. H. Ungar. (1991). Automatic generation of qualitative models of chemical process units. Computers and Chemical Engineering, 15, No. 8, 583-599.

Dalle Molle, D. T. and T. F. Edgar. (1989). Qualitative modelling of chemical reaction systems. In M. L. Mavrovouniotis (Ed.), AI Applications in Process Engineering. Academic Press.

Dalle Molle, D. T., B. J. Kuipers and T. F. Edgar. (1988). Qualitative modelling and simulation of dynamic systems. Computers and Chemical Engineering, 12, 853-866.

Davis, R. (1984). Diagnostic reasoning based on structure and behavior Artificial Intelligence, 24, No. 1-3, 374-410.

DeCoste, D. (1990). Dynamic across-time measurement interpretation. Proc. of the Ninth National Conference on Artificial Intelligence.

de Kleer, J. and J. S. Brown. (1984). A qualitative physics based on confluences. Artificial Intelligence, 24, No. 1-3, 7-83.

Dvorak, D. and B. J. Kuipers. (1991). Process monitoring and diagnosis. IEEE Expert, 6, No. 4.

Fogler, H. S. (1986). Elements of Chemical Reaction Engineering. Prentice-Hall, Englewood Cliffs, NJ.

Forbus, K. D. (1984). Qualitative process theory. Artificial Intelligence, 24, No. 1-3, 85-168.

Forbus, K. D. (1987). Interpreting observations of physical systems. IEEE Transactions on Systems, Man and Cybernetics, 17, No. 3, 350-359.

Forbus, K. D. and B. Falkenhainer. (1990). Self-explanatory simulations: an integration of qualitative and quantitative knowledge. Proc. International Joint Conf. of Artificial Intelligence, 380-87.

Frank, P. M. (1990). Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy -- a survey and some new results. Automatica, 26, No. 3, 459-474.

Isermann, R. (1984). Process fault detection based on modeling and estimation methods - a survey. Automatica, 20, No. 4, 387-404.

Kuipers, B. J. (1986). Qualitative simulation. Artificial Intelligence, 29, 289-388.

Kuipers, B. J. and D. Berleant. (1988). Using incomplete quantitative knowledge in Qualitative Reasoning. Proc. of the Seventh National Conference on Artificial Intelligence.

Kuipers, B. J. and C. Chiu. (1987). Taming intractable branching in qualitative simulation. Proc. of International Joint Conf. on Artificial Intelligence (IJCAI-87), .

Oyeleye, O. O., F. E. Finch and M. A. Kramer. (1990). A robust event-oriented methodology for diagnosis of dynamic process systems. Computers and Chemical Engineering, 14, No. 12, 1379-1398.

Ramesh, T. S., J. F. Davis and G. M. Schwenzer. (1992). Knowledge-based diagnostic systems for continuous process operation based on the Task framework. Computers and Chemical Engineering, 16, No. 2, 109-127.

Shewhart, W. A. (1931). Economic control of quality of manufactured product. Van Nostrand Co. , New York.

Venkatasubramanian, V. and S. H. Rich. (1988). An object-oriented two-tier architecture for integrating compiled and deep-level knowledge for process diagnosis. Computers and Chemical Engineering, 12, No. 9/10, 903-922.

Vinson, J. M., S. D. Grantham and L. H. Ungar. (1992) Diagnosis using automatic rebuilding of qualitative models in chemical plants. In Print, IEEE Expert, August.

Willsky, A. S. (1976). A survey of design methods for failure detection in dynamic systems. Automatica, 12, 601-611.

Yu, C.-C. and C. Lee. (1991). Fault diagnosis based on qualitative/quantitative process knowledge. AIChE Journal, 37, 617-28.