PROCESS CONTROL SUPERVISION USING QUALITATIVE MODELS

R K Stobart, Cambridge Consultants Ltd, UK
N R Shadbolt, University of Nottingham, UK

Abstract

In modern process and power plant the cost of installation demands operation at peak efficiency for prolonged periods. This places significant demands on the control and monitoring systems to keep efficiency high while giving significant warning of a drop in efficiency or a component failure. We present an approach to process monitoring based on Qualitative Models which are used as a framework in which a range of monitoring techniques are located. The methods are described in the context of diagnosing faults in a heat exchanger. We present an application based on a Cogeneration scheme which highlights many of the issues in process monitoring.

1. Introduction

In the process control industry, supervision consists of monitoring the behaviour of the controller, sensors, actuators and the chemical or physical process itself. It must efficiently separate changes in condition and faults among these various components.

The "traditional" approach to diagnosis in process control has been to identify alarm conditions and present to the operator an indication of the alarm state when it arises. In certain pathological cases the flow of alarms can be sudden and in such numbers that the operators quickly become overloaded. Too much is expected of the operators who must both process the information and rely on their own understanding of the process, its fault conditions and how to correct the problems. However operators are frequently very successful if given sufficient time to assimilate the information (Carbonelle,1989).

An operator's understanding of the process may be limited to some prior experience and not the physical and chemical principles on which it operates. By resorting to a description of the process at the level of governing principles, there is a clearer opportunity to find and correct the problem. We intend to explore the techniques presented by Forbus (1984). His work is centred on the identification of a number of processes in which the subject participates.

Our interest in this field springs from previous work in the development of causal methods for predictive maintenance and the development of very able but restricted numerical models. A link between these two approaches combines the explanatory possibilities of a causal model with the very good resolution that numerical methods offer.

In the remainder of the paper we will present the framework for using qualitative models and support our ideas with a realistic example. We will present an architecture for exploiting causal models and then illustrate how, in a complex application it would be used to support monitoring and diagnosis of faults.

2. A Framework for Using Qualitative Models

Qualitative models derived from the principles of operation of a system summarise all the main causal relationships that operate between components. Increasing detail however has a severe effect on the complexity of the reasoning processes that must be used, and may not return an equivalent benefit. A recent commentary on an extension to Qualitative Reasoning by Dvorak and Kuipers(1989) describes a system in which the qualitative models are supplemented with sets of qualitative first order differential equations. This appears to add significantly to the complexity of the resulting implementation. A companion piece of work (Kuipers and Berleant, 1988), is concerned with adding incomplete numerical information to a Qualitative Model. We wish to add numerical information but in terms of supporting a simple qualitative reasoning process with tests needed to support its conclusions.

A supervision system based on simple qualitative models provides a means of monitoring and reasoning about a system at a high level. It is useful for looking at limits of behaviour, such as when a sensor fails in a closed loop control system. However, the system of processes and instances which makes up a Qualitative Model provides an excellent framework for a monitoring system in which additional sources of information can be tapped. In trying to match a set of observations with the process model, several possibilities may be proposed for system condition. To resolve the analysis we may appeal to external sources including diagnostic tests, gas analyses, or the results of front end signal processing. In an adaptive control system for example, the controller itself will maintain a numerical model of the plant under control. In most modern control and monitoring system it is possible to exploit the existence of monitoring points to support or refute conclusions.

Consider the case where an adaptive control system maintains a system model to support its design activity. The model provides an additional source of information about the process albeit indirect. While the adaptive controller continues to adapt it is continually updating an internal model of the process using recursive estimation. Usually, this is a Z domain or similar model which is used to redesign the controller at each sample. In generalised predictive control the model is used to assist in the prediction step which lies at the heart of the algorithm. McMichael (1988) has demonstrated that recursive estimation may be used for precise and versatile fault diagnosis in process systems. Unfortunately the adaptive controller may switch off adaption from time to time or its estimates may drift when the frequency content of the incoming signals falls off. Such techniques require a generic interface so that additional information can be sought during the analysis of abnormal plant conditions.

2.1. An Example - Heat Exchange

Consider heat exchange as a supervised process. Water is passed around the jacket of the reactor vessel in order to regulate the temperature of the contents. Table 1(a) is the definition of the control process. Even in a continuous process like heat transfer the various processes such as control and heat transfer are intermittent and initiated as the result of certain preconditions.

The process description contains a number of slots. Conditions may be expressed as predicates or conditions. For example the quantity condition in the control process states that the magnitude of the error is greater than the dead-band value. The influence slot shows the direct effect of the process. For control this shows that a positive error causes a positive output.

Table 1(a) shows the precondition for control is that the control error lies outside a dead-band. Industrial controllers are usually configured with a dead band value that limits too frequent control adjustment. The influences show the direct effect of the control process. In our example a control error causes the control output to increase.

Table 1(b) shows a specification for the heat exchange process. It shows that heat exchange is only taking place when there is flow in the vessel jacket. The *relation* slot shows some dependencies which result from considering the physics of the heat exchange process. It shows that the temperature drop in the secondary fluid flow is inversely related to the flow rate and the heat transfer capability of the jacket.

The third process of interest is the heat transfer from the jacket to the fluid in the vessel. When the fluid in the vessel changes its temperature due to a chemical reaction the error between the setpoint temperature and the actual temperature increases. This creates the precondition for control activity. The control process changes its output according to the error and in turn creates the precondition for secondary fluid flow. Finally the heat transfer to the vessel is started. As the temperature in the vessel changes, the error which started the control in the first place is reduced until control is stopped and the processes are successively terminated. In a well tuned control system the processes invocations will produce a slowly changing state without oscillation.

Consider three types of deterioration,

(1) the loss of the temperature sensor,

(2) the jamming of the valve,

(3) deposition of material in the jacket leading to poor heat transfer.

If the temperature sensor fails high, the controller will attempt to cool the vessel more and more until the valve is fully open where it will stay. In using qualitative models we will consider a number of *limit hypotheses* which result from the activation of the various processes. If a large error develops, the controller will activate secondary fluid flow, this in turn will activate cooling in the process vessel. Normally the processes will form a loop which will terminate in finite time. However in this case the control output reaches its upper limit. The existence of a process at its limit should initiate an exploration of the limits to be reached by other processes. An exploration in this case will reveal that while secondary fluid flow has reached a limit (valve wide open), there is no movement in the vessel temperature. In an implementation such an exploration will consist of the presentation of hypotheses by a knowledge source which implements the qualitative model and reasons about it. We can bring some evidence to bear which deals with the two possibilities that either there is a process malfunction or that the temperature sensor has failed. If the process is very stable then we must deduce that the temperature reading is wrong.

The valve jam is similar and results in a limiting value of vessel temperature being reached. In this case we would look at each of the process elements and note that control is functioning while it appears that the secondary fluid flow is not, since the valve position is not at maximum as demanded by the controller. The break between control and the valve is indicative of either a broken connection or a valve malfunction.

The deposition of material in the jacket will cause a slow deterioration in the heat transfer performance. The temperature drop across the vessel will progressively decrease but very slowly. In exploring the limit hypotheses one will imply that ultimately there will be little heat transfer if the heat transfer coefficient rises. The deposition process will be defined to directly influence the heat transfer coefficient and will threfore be a prime suspect. There could be a number of other causes. However an appeal may be made to the process model maintained by the adaptive controller. As deposition takes place the gain of the process will fall since it will require larger and larger valve movements to bring about a given temperature change in the vessel. The adaptive controller will return this and other information and could be used to support the deposition hypothesis.

Table 1(a) Process description for Control

Process control

Individuals

C a controller, Closed-loop(C)
M a measurement

Preconditions

Measurement_connected(C)
Output_connected(C)

Quantity condition

A_m (error) > DEAD_BAND

Influences

I ± (output, error)

Table 1(b) Process description for heat exchange

Process secondary fluid flow

Individuals

S a heat exchange fluid
H a heat exchanger
V a valve

Preconditions

Connected(H)

Quantity conditions

A_m (flow-rate(S)) > 0

Relations

$\Delta T \propto_{Q^-}$ flow-rate(S)
$\Delta T \propto_{Q^-}$ heat-transfer-coefficient(H)

Influences

I ± (flow-rate(S), valve-position(V))

3. An Architecture for Exploiting Qualitative Modelling in Process Control

Figure 2 shows the architecture of the supervision process in block diagram form. Each block represents a knowledge source(KS) which in turn consists of data and rulesets to support a particular aspect of system function. The unit of data representation in the knowledge sources is the object which contains both slots for data storage and functions which respond to messages.

The knowledge sources run on a priority basis so that if a rule within a high priority KS is ready to fire, low priority KS's will be suspended while its execution completes. In this application, the interface KS would have a high priority in order to respond to external events as quickly as possible. This reflects the need to keep up with "real world" events while the Qualitative Model running in a low priority KS would occupy a background role.

Information is passed between knowledge sources using the notice board(NB). ,One knowledge source may place a piece of data in an object on the NB. An "interested" knowledge source will be monitoring this data and will run to process as its priority permits.

Such an architecture is event driven. Incoming data to the Interface KS will pass through data channels to which *demons* are attached. These are functions which note the appearance of data and will execute some action such as processing or transfer of the data. A typical sequence of events following the arrival of some new data might be as follows.

- the data arrives and triggers a demon to run a function to process the data;
- the function result is significant in that it requires some diagnostic action; it is placed in the notice board;
- the Qualitative Model KS runs and processes the data; to confirm the results of a diagnosis it refers to the Tests KS and places a request for tests on the NB;
- the Interface KS runs when the test details appear and executes the required actions; the results will be placed back on the NB ready to be processed by the Qualitative Model KS.

The purpose of separate KS's is to modularise the application and improve its development and maintenance. It is likely for example that the KS which defines what tests are available will need to be updated quite frequently while the remainder of the application will remain fairly static.

The Qualitative Models (QM) are maintained in one knowledge source which contains the process descriptions and instances for the monitored system. This knowledge source has a view of the incoming data which is located in the notice-board and draws additional information from the Supporting Knowledge KS.

Tests which may be called up to confirm or otherwise the initial conclusions will be located in another knowledge source. They will be defined in terms of difficulty, cost and the results which will be returned.

During system operation, the QM KS will produce a series of assertions about the health of the system. Consider the case where the health of a gas turbine engine is being monitored. Here we might conceive of three processes, compression (through the compressor), expansion (through the turbine), and fouling (of the compressor through the ingestion of dirty air). The efficiency of the engine is dependent on the compressor and turbine efficiencies. Whenever the engine is running and ingesting air, fouling will be taking place. The salient points about the fouling process are,

- the precondition is that the engine is running at a load condition
- the effect is the reduction of compressor efficiency
- the effect is cumulative and must be recorded.

In our implementation the fouling process will cause the object representing the compressor to be updated with the degree of fouling. As part of the continuous process of checking the model the efficiency of the engine will be investigated. Efficiency is of prime economic concern and will be calculated from time to time on the basis of

- fuel used,
- electrical power generated, and
- process heat transferred.

Before requesting a diagnosis, the monitoring system would need to be certain that efficiency was declining. We suggest that within the Interface KS there is a demon monitoring the arrival of new efficiency data. This will initiate a look at the efficiency data and perhaps fit a time series model to isolate trend from noise. Once the trend was noted, the value would be placed on the notice board. A further refinement would be to use the time series to predict to the time when efficiency would be of concern. By adopting this sort of approach the system will be event driven and will absorb only as many resources as are needed to solve the problem and make a recommendation.

The gas turbine example given here assumes that the installation consists of a gas generator (compressor and turbine : equivalent of an aerospace jet engine without a nozzle), and a power turbine which converts the high temperature gas stream to rotational energy.

In quantitative terms,

$$efficiency = \frac{W_{power\ turbine}}{heat\ input}$$

The energy available to drive the power turbine is in the gas stream produced by the gas generator. In turn the efficiency of the gas generator might be loosely defined,

$$efficiency = \frac{W_{turbine} - W_{compressor}}{fuel\ flow * calorific\ value\ of\ fuel}$$

These relations are readily translated into the QP form.

$\eta \propto_{Q^+} W_{turbine}$

$\eta \propto_{Q^-} W_{compressor}$

$\eta \propto_{Q^+} CV_{fuel}$

$\eta \propto_{Q^+} W_{power\ turbine}$

Further relations concern how the compressor work ($W_{compressor}$) is influenced.

$W_{compressor} \propto_{Q^-} \eta_{compressor}$

which states that the compressor work increases as the efficiency falls.

Finally it is the fouling process itself which causes the compressor efficiency to come down and in turn reduce the overall efficiency.

Note that the check on efficiency is a generic check which will embrace a variety of sub-checks, but efficiency is the quantity with which the user is primarily concerned. When the suspected degree of compressor fouling reaches a level affecting efficiency, external data will be sought. The QM KS will search for a test in the TEST KS and will submit a request which will be placed on the agenda. One possible conclusion, is that the power turbine work has decreased. Using the supporting knowledge this would be judged unlikely since faults in power turbines are rare. A more likely conclusion is that compressor efficiency has fallen.

One possibility for the compressor efficiency test is to look at numerical data from the engine, such as a simple dynamic model which can be updated from time to time. A very different form may be a request for manual analysis of recent performance data to give a value for compressor efficiency.

At any stage, the agenda will have a series of tests on it waiting to be done. These will include

- checking a database,
- initiating an external test, or
- initiating a numerical analysis of incoming data for a feature of interest.

A second source of information will be in a supporting knowledge source. Here various items pertaining to components will be stored such as reliability, when maintenance last took place, and rulebases to support analysis. This will support the kind of analysis where the QM requests which of a list of components is most likely to have failed.

It is clear that in simple systems the QM will be used repeatedly to solve the same sort of problem. In this case, it has been valuable in forcing a design discipline and setting an agenda for the likely faults. An implementation which most suits the environment could then be developed. However, in the general case, the QM will still be best allowing for future expansion and refinement.

4. A Complex Example

Using a commonly applied industrial system we now present a detailed and comprehensive illustration of our framework.

Cogeneration is the process of simultaneously generating electrical and heat energy. In a sense all power plant does this but there is increasing interest in selling the heat which is a guaranteed by-product. There are many industrial examples of cogeneration based on steam turbines, gas turbines and diesel engines where the heating is needed for process purposes and is often more critical than the electrical energy. Increasingly Utility companies are using cogeneration, and there are severe commercial pressures on keeping efficiency and reliability very high. There is a strong motivation for an efficient and accurate supervision and diagnostics system.

4.1. A Cogeneration Scheme

We will consider a simple cogeneration scheme (Figure 3) in which a diesel engine is used to drive an electrical generator. Cooling water is passed through the engine and to a heat exchanger where heat is transferred to a heating circuit.

The diesel engine may use a range of fuel types, but the quality will generally be low. In some parts of the world the sulphur content may be high, and in general the amount of particulate matter will be high. Some heat will be recovered from the exhaust, and the heat exchanger will foul quickly. In an industrial environment the air ingested by the engine is likely to be dirty in spite of filtering.

The district heating scheme consists of a set of pumps driving fluid around a circuit consisting of the main heat exchanger, and in each heating zone there is a valve, heat exchanger and temperature sensor. Each zone will be locally controlled so that a local closed loop controller will modulate the valve to maintain the zone temperature at the required value. The pumps will be of a centrifugal type with 3 phase induction motor drives.

4.2. A Description of the Knowledge Base

The knowledge base can make use of the multiple instances of heat exchangers in the system. While the precise arrangements are slightly different, the QM KS can still use one heat exchange model multiply instanced.

We have identified the following main processes.

- air supply to the engine
- energy conversion in the engine
- primary heat exchange (from engine circuit to heating system)
- secondary heat exchange (from heating circuit to the heated zones)
- fluid transfer

The *view instance* proposed by Forbus is not a particularly useful concept where the process runs at steady conditions. We have not used it in this example where all requirements are met using processes and individuals.

Table 2 shows the description for the air supply process.

Table 2 Process description for Air Supply

Process air supply

Individuals

E an engine
T a turbine
C a compressor

Preconditions

Engine_running (E)

Quantity condition

A_m (engine_load) > 0

Relation

boost pressure $\propto_{Q^+}$ efficiency(C)
boost pressure $\propto_{Q^+}$ efficiency(T)

Influences

I + (engine power, boost pressure)

The air supply process is only of interest when the turbocharger starts to operate and this corresponds to a moderate engine load. Note that the fouling process will reduce the efficiency of the compressor through the *influence* slot.

Problems with the air supply will only be one way of accounting for the perceived loss of efficiency. Other problems detected by the QM will include a variety of engine faults. At this stage we will need to appeal to evidence which can differentiate between the turbocharger and the engine as the source of the fault. A simple numerical technique for doing this has been developed (Stobart and Eastaugh, 1989) which could be located in the Interface KS.

This system enjoys a clean interface between the engine component and the heating scheme. On the heating side of this division, fluid transfer is an important process and is defined in Table 3.

Table 2 Process description for Fluid Transfer

Process fluid transfer

Individuals

P a pump
V a valve

Preconditions

Connected (P)

Quantity condition

A_m (speed (P)) > 0

Relation

ΔP $\propto_{Q^+}$ pump efficiency
ΔP $\propto_{Q^+}$ motor efficiency
pump efficiency $\propto_{Q^+}$ cavitation

Influences

$I+(\text{fluid flow rate}(F), \Delta P)$

One conclusion which will result from low temperatures in the heating circuit will be a loss of fluid flow. Using this process description, this will ultimately lead to an investigation of the possibility of cavitation in the pump. The QM KS will invoke a test to find out if this is the case. Cavitation is the evolution of vapour bubbles due to excessive pressure drop, which in turn is due to wear or damage in the pump. One method may be to check the pump's acoustic emissions either by ear or using a sensor. A second method is to use a numerical technique to investigate the pump's pressure history, (Geiger, 1984). Both tests could be embedded in the Interface KS.

There are further possibilities for using the QM to make an initial investigation, which can then be confirmed or otherwise by an external test. What we have demonstrated is that using simple causal models supported by additional knowledge and a battery of tests and algorithms, allows us to produce enhanced diagnostic capabilities. The presence of a causal network also allows a detailed causal explanation to be offered for any conclusion reached.

5. Conclusions

The main conclusions of our study are as follows.

- Qualitative models of industrial processes are tractable for even quite complex processes.
- The model provides a structure for using a diverse range of measurement and monitoring methods. The model is used to produce a range of assessments of system condition and external tests are used to support or refute the conclusions. In this way the model can be kept simple while the test methods change and absorb much of the complexity of the system that has posed significant problems for other QP approaches.
- An architecture for the application of this approach is based on multiple knowledge sources. Generic interfaces deal with variations in the final implementation of a test.
- In an implementation, the qualitative model can provide on-line assistance while in simple cases it offers the design framework for a monitoring system.

References

Forbus K D (1984). Qualitative Process Theory. (In D.G. Bobrow (Ed.),Qualitative Reasoning about Physical Systems), North Holland, pp85-168.

Carbonelle, J (1989). "Machine Learning and Real Time Architectures for Intelligent Systems", Keynote Address at IEA/AIE-89, University of Tennessee Space Institute.

McMichael, D W (1988). "On-line Fault Detection : A System Non-Specific Approach", DPhil Thesis, University of Oxford.

Dvorak, D & Kuipers, B (1989). "Model Based Monitoring of Diagnostic Systems", International Conference on Artificial Intelligence, Detroit, MI

Kuipers, B & Berleant, D (1988). "Using Incomplete Quantitative Knowledge in Qualitative Reasoning", The Seventh National Conference on Artificial Intelligence, St Paul, MN

Geiger, G (1984). "Fault Identification of a Motor Pump System", IFAC World Congress, Budapest.

Stobart, R K & Eastaugh, P R (1989). "Signal Processing Techniques for the Diagnosis of Diesel Engine Faults", Seminar on the Benefits of Electronic Control for Internal Combustion Engines, Institution of Mechanical Engineers, London.

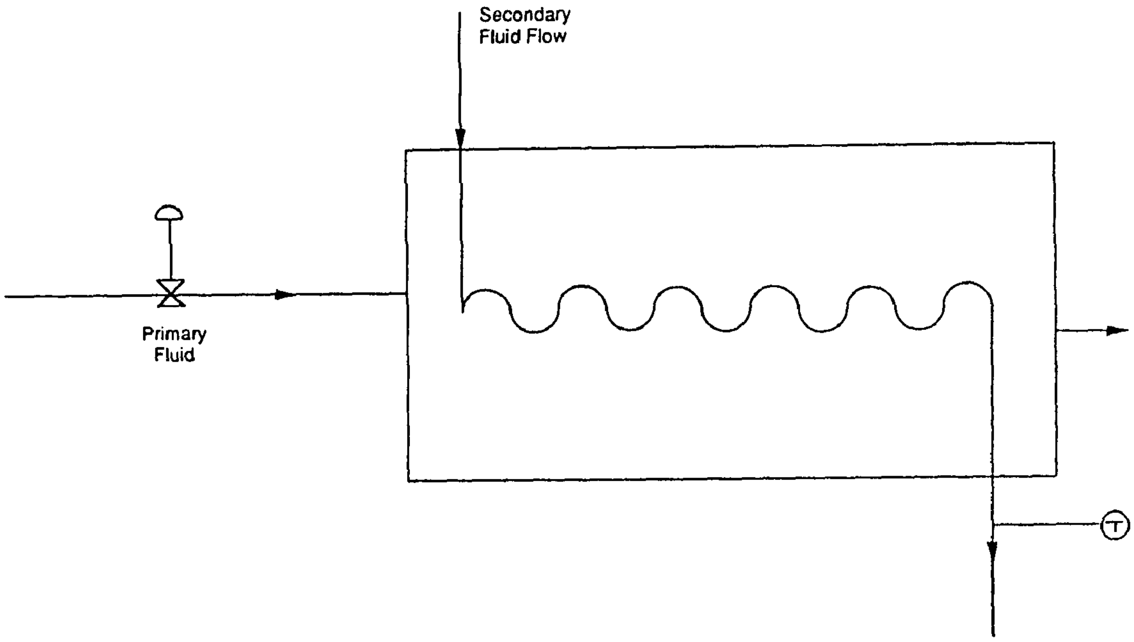

Figure 1 : A simple heat exchanger

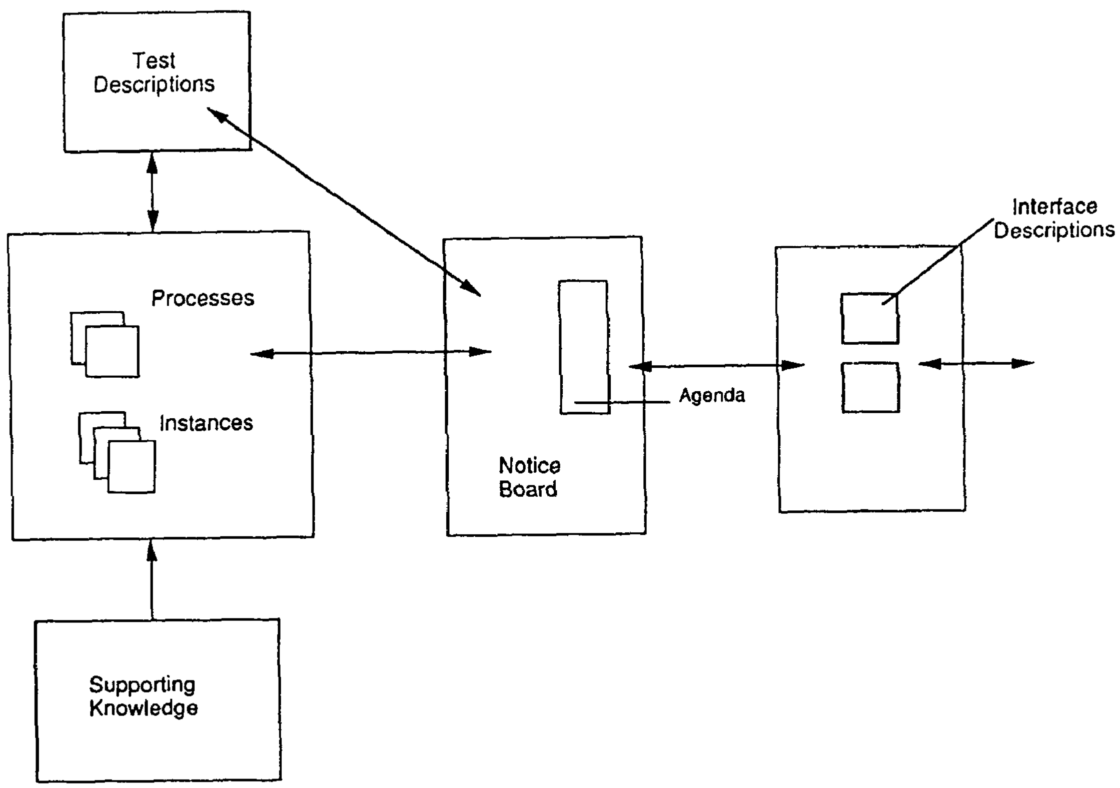

Figure 2 : An outline architecture for process monitoring

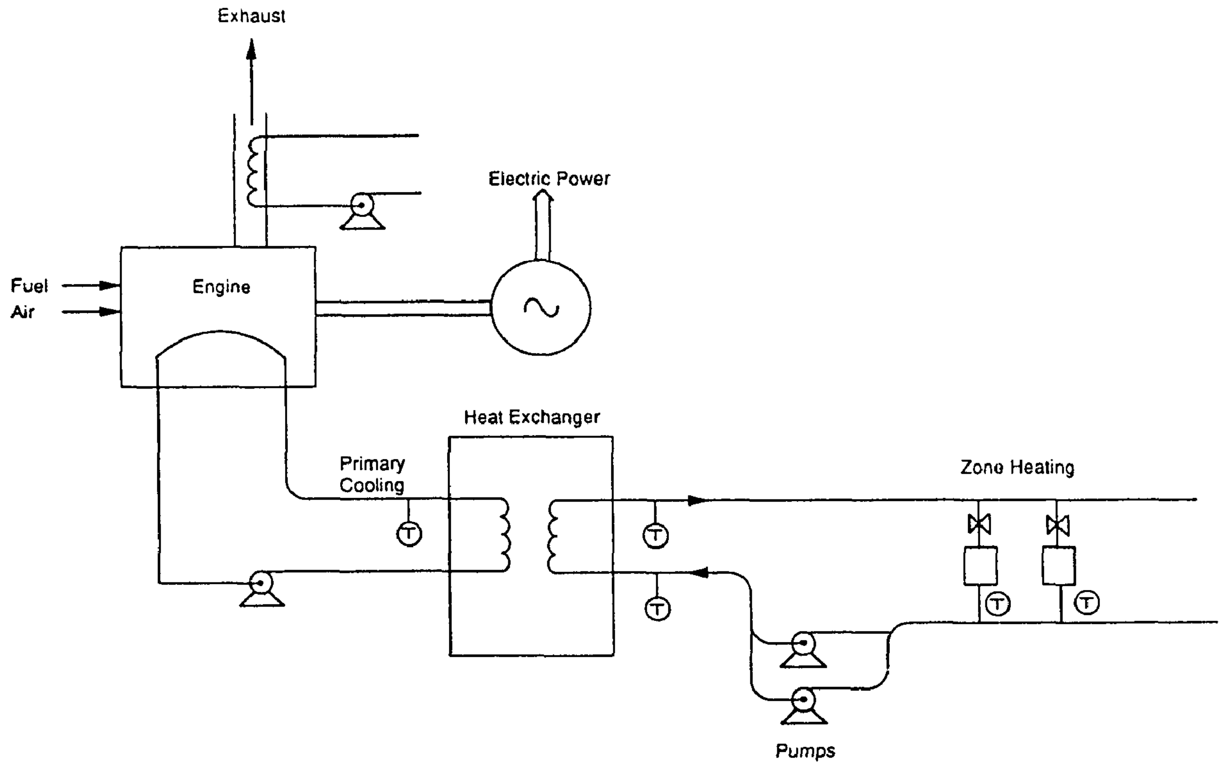

Figure 3 : A simple cogeneration scheme