MONET2

Project Full Title: Network of Excellence on Model-based Systems and Qualitative Reasoning.

Contract: Concerted Action / Thematic Network

Contract No: IST-33540

Deliverable MBD5:

Report on Evaluation of Solutions to Selected Problems of Industrial Relevance

Date: 1st September 2004
Version: 2.3
Status: Release

MONET Project Office
Department of Computer Science
University of Wales
Aberystwyth
Ceredigion
SY23 3DB

Contents

Authors

Joaquim Armengol	University of Girona
Fabrice Auzanneau	Commissariat a l'Energie Atomique
Marie-Odile Cordier	Institut National de Recherche'en Informatique et en Automatique Unite de Recherche IRISA
Teresa Escobet	Universitat Politèchnica de Catalunya
Carlos Alonso Gonzalez	University of Valladolid
Robert Milne	Intelligent Applications
Stéphane Ploix	Institut National Polytechnique de Grenoble
Belarmino Pulido	University of Valladolid
Benny Raphael	Ecole Polytechnique Federale de Lausanne
Iain Russell	MONET Project Office
Marcel Staroswiecki	Universitie des Sciences et Technologies de Lille
Peter Struss	OCC'M Software GmbH and Technishe Universitaet Muenchen
Janet Thomas	MONET Project Office
Louise Travé-Massuyès	Centre Nationale de la Recherche Scientifique - Laboratoire d'Analyse et d'Architecture des Systemes
Samir Touaf	Institut National Polytechnique de Grenoble
Cecilia Zanni	LSIS - Laboratoire des Sciences de l'Information et des Systemes

1 Introduction

1.1 Purpose of this Document

This document is designed to assess the current situation with respect to the solutions to the issues raised by the industrial applications of Diagnostic technologies and also to provide perspectives on where we believe these technologies need to be driven in the future. This should lead to clearly identify lines, or areas, of investigation for the coming years; this will also feed into the update of the Bridge Task Group Technological Roadmap.

The information contained within this Document is focused on the 'Solutions' to the Industrial Reference Problems and thus it does not necessarily offer a full insight into the problems themselves. For further details of these problems please see the Deliverable MBD3 Collection of Industrial Reference Problems.

1.2 Scope

One of the major issues is that the problem space for our set of reference problems is very large. It would be useful if we had reference problems that only had one or two major features of difficulty, this would make them considerably easier to analyse and group. This is, unfortunately, not the case, as these are all 'real world' grounded and are invariably a combination of several complex problems. However, the benefit to this is that the solutions that we are suggesting have been derived from actual problems and not a simplified system without all of the noise or unexpected behaviours that are found in the real world.

The reference problems approached here represent a wide and varied set of systems and solutions; however, this is not an exhaustive list and there are characteristics that are missing. To complete the group we would need at least one pure discrete event system. Also we have had the task of characterising the systems by subject, device or domain but there is also the other dimension of characterising the task itself. We have seen, during discussion of what industry actually wants, that they often do not require a complete diagnosis, they may just want fault detection, or they want to do diagnosis for maintenance, reconfiguring control, etc. Tackling the problem from this dimension puts a different focus on the way the problem is approached.

Due to the industrial focus of the issues, we started with the reference problem and then considered the tasks that were required to solve that reference problem. This task meant that we did not try to be exhaustive, but to focus clearly on solutions that were applicable to the tasks at hand and that would give the best solution to the user in the simplest way possible. For example the Steam Generator problem solution focused on 'internal faults diagnosis', whereas 'supervision', which would have called for considering external faults as well, was not the goal.

The most important issue that has come out of this exercise is the realisation that when we are working on the analysis of any Industrial Reference problems we need to focus closely on providing 'industrially applicable' answers, or parts of answers to evaluate the solutions. We need to look at the evaluated solutions in terms of their use to industry and one way would be to use the same language as the Industrialists, to whom we are targeting the solutions and define the kind of criteria that would be understood by these industrial persons. After all, Diagnosis as a process is not just aimed at detecting faults; it is also to provide an accurate (and preferably up to date) state estimation even if there are no faults.

2 Background to Modelling Solutions

It is argued that for some approaches it is not necessary to make a distinction between internal and external faults. If you have a model that is basically a network of constraints then the important fact is whether or not you are able to characterise either the proper input or the proper output of the behaviour, by use of that set of constraints. For consistency-based diagnosis, it often does not matter what the fault is, it could be that 'the pressure is too high' or the 'valve does not open'. It is not actually important to differentiate for this approach. It all comes down to what we would consider a 'wrong' or 'bad' input for the model.

For example, there is the basic premise of the consistency-based diagnosis perspective. What we are doing is fault detection and we start the diagnosis from the point where we have a model of the correct behaviour and all the components are working. Then we have to check for consistency with the observations. This approach however, raises the question of why this is enough? After all, the result we are actually looking for, is to know whether the system (i.e. plant or device) that we are dealing with is actually fulfilling the purpose for which it was designed. The essential element of deciding whether something is going wrong therefore becomes 'the device is not (fully) fulfilling its purpose', rather than 'there are no detectable faults'. This is a very recent but very important shift in diagnosis.

Thus we need to create the model where all the components are operating correctly and then use the observations to predict the current state of the system. Then it is necessary to check whether this is consistent with the goals in the specifications of the behaviour. Previously the theory was, at least in component-based diagnosis, that if all the components work correctly that this entails that all the goals are being met. Thus it is assumed that it is a well designed system and so if the components don't fail then the goals are achieved. However, with modern diagnosis we are required to make the goals explicit, as the aforementioned belief does not necessarily hold true because we are aware of conditions where, even when all the components are working, the goals are not fulfilled.

This gives us two sources for identifying a problem. First, we have the model of the components; this tells us that a component is faulty if it exceeds a certain specification or is working outside a certain set of parameters, or indeed not working at all. The second source is the specifications themselves; this means that it may even be impossible to satisfy the goals under certain conditions even if all the components are operating optimally. A sub-set of these goals for industry could well be financial goals, which could mean that we have no interest (or will) to fix the fault, but we want to know if the system will keep working, even if this is only for a limited time frame, i.e. until the scheduled maintenance is due. The main difference is the conception of the model. The model is constructed from known behaviours of component parts, which is physics and can not be altered, however, the goals are not 'physics', they are an expression in physical terms of the expected behaviour of the device. Although this later definition does not change the physics, it imposes a restriction on what you would expect from the behaviour and thus has a different focus.

The major problem with this goal explicit approach is that it rapidly becomes very difficult to construct anything that we can check against the physical model of the components. This is also related to the problem of distinguishing between the presence of faults and deciding which of the faults discovered is most interesting at any one moment. We need, therefore, to make the distinction between component faults and specification / goal faults and utilise this focus when we analyse how we can define the bounds or the intervals for the involved parameters. Then we need to think of this in terms of where the bounds come from; do they come from the intent of the specifications or a faulty component and what does this distinction mean to the way we view the fault.

Further to this approach the environmental constraints must also be thought of as part of the process. First there is the model of the components; this is the system that has been built and this system has been built in order to achieve some goals. The system, however, can only do this by working within a set of prescribed constraints. This gives us three distinct areas to the process; Models, Constraints, and Goals. The constraints themselves are connected with the environment. This makes them distinct from the Goals, which are not connected with the environment, only with how you want the system to perform.

Goals can potentially be described within the model itself. One weakness of this type of approach from the FDI perspective is that this approach does not describe the components, but one strength is that it can describe ‘everything’. For example, we have a set of equations that describe the components and also a set of equations which describe the goals. One of these goals may be that the system should always be stable and not lose control and this type of goal implies a set of constraints. Thus the constraints can suggest goals which are actually negative. Examples of these are financial or technical constraints / goals which state that a certain variable should never exceed a certain level because it requires the utilisation of an expensive resource. With the inclusion of such ‘external’ goals and constraints the device is then dealing with ‘everything’ that you need to be concerned with. This may take away from the value of the modelled components but gives a much greater view of the process as a whole. For the model building process these two things need to be conceptually separated, although for the end user they may well appear to be the same thing. The overlap occurs where the external environment is modelled as another feeder component and, with its own set of constraints, is just another component part of the model.

3 Reference Problem Solutions

3.1 Steam Generator Process within the CHEM Project

Marcel Staroswiecki
Université des Sciences et Technologies de Lille
France

The steam generator pilot process has been used in Lille for more than ten years.

The supervision problem consists of detecting and isolating process, actuator or sensor faults as soon as possible. For this, the values of a set of “residuals” have to be calculated and the trends of these values have to be analysed, all of this, in real time. The outputs should be the trends of the residuals and the list of the detected faults.

Different approaches were tested to evaluate Fault Detection and Isolation techniques. For example, Bond Graph and/or structural models have been used to define the links between the variables and the parameters of the operating model independently of the form under which this operating model is expressed (quantitative or qualitative data, analytical or non-analytical relations). The links are represented by a graph whose structure is analysed to deduce redundancy relations and then residuals. On the other hand, Matlab-simulink models have been used to test FDI / FTC strategies through simulations.

The steam generator pilot process was recently used as a reference problem in the CHEM[1] FP5 project to evaluate FDI techniques and reconfiguration strategies.

---

[1] "Chem Project is funded by the European Community under the Competitive and Sustainable Growth programme of the Fifth RTD Framework Programme (1998-2002) under contract G1RD-CT-2001-00466."

3.1.1 The CHEM project

CHEM is a European Fifth Framework project concerned with efficient detection and repair of problems with chemical processing plants. The consortium held an open presentation of their results at the end of the project in March 2004. That meeting gave the profile of a remarkably successful European collaboration, producing significant results. This report highlights those results and speculates on why this particular project might have been so successful.

CHEM provides decision support tools to help process operators make difficult decisions in pressure situations. The CHEM toolkit includes a large number of tools to help operators in different areas:

- Trend analysis and situation assessment
- Fault diagnosis and alarm management
- Reactive planning and plant reconfiguration
- Decision support tools to underpin the other abilities in the CHEM toolkit

Several tools within the CHEM toolkit use qualitative and/or numeric models. For the two first classes such tools include:

1 - Trend Analysis and Situation Assessment:

- **SQualTrack** (UdG) and **TempBand-FDI** (Thales) tackle the problem of uncertainty using intervals.
- **FeatureX** (KCL) extracts features from trends.
- **ASCO** (IFP in collaboration with LAG and LAAS) tackles the problem of modelisation using quantitative causal graphs.
- **QualTras** and **Salsa** (UdG) classifies trends and situations.
- **iMSPC** (Computas) exploits the power of multivariate SPC.
- **MITforRD** (WUT) uses fuzzy logic and neural networks for process modelling.

2 - Fault Diagnosis and Alarm Management:

- **ModBuild** (USTL-Lail) is a functional model builder for qualitative and structural analysis.
- **TempBand-FDI** (Thales) isolates faults using a method that ensures no false alarms.
- **ExSit-M** (UPC) detects and isolates multiple faults in continuous and batch processes.
- **iFuzzy-FDI** (WUT) combine neural networks learning, HAZOP analysis and knowledge based systems to support the operator in the analysis of causes and consequences.
- **ASCO** (IFP in collaboration with LAG and LAAS) isolates faults on the components.

There were at least as many additional tools, providing a broad facility for building systems for running chemical process plants, and one message from that is qualitative reasoning is only a part (albeit an important one) of a larger system. A number of members of MONET's BRIDGE task group were involved in CHEM, and it was also good to see the level of understanding in CHEM of the differing contributions that can be made by FDI techniques and by model-based AI techniques.

Given the wide spectrum of tools available in the CHEM toolkit, it was good to see that they had applied many of the tools to the same case studies. In many instances, a number of the tools were linked together to produce a specific diagnosis or reconfiguration result. One of the case studies was the steam plant in Lille, which is our reference problem. Faults were applied to the steam plant in real time, and caught and dealt with by the CHEM tools. Another case study was a CORUS blast furnace plant in Redcar, UK. The system, using several of the CHEM tools working together, was able to identify potential instabilities inside the furnace, long before operators could, giving enough time for the operators to take action to avoid wastage within the furnace.

CHEM cost about six million Euros, with about four million Euros provided by the European Commission. For that money, the project has produced a large number of tools, integrated them, and applied them successfully to a range of realistic case studies. The extent to which they have done this is impressive, even given the size of the project. There are a number of reasons why they have achieved this, and here is an independent view from outside the project of what those reasons are:

Gestation time of technology. A number of the partners have been working on related tools on European projects for at least a decade. This effect should not be underestimated. Not all technology can be developed in a three year project. An early, promising technology can take time to develop to the stage where it is ready to be used in earnest by industry. It also takes time for different organisations to learn to trust and understand each other, and some of the relationships within CHEM are between organisations that have worked together on previous European projects.

Clear, shared view of process. CHEM used the view of the tasks involved in diagnosis and repair that has been developed by the diagnostic communities over many years, separating the tasks of monitoring, fault detection, fault identification and repair, so that it is clear which tasks are being performed by each of the tools, and how the tools can cooperate to achieve a given goal.

Shared integration technology. CHEM focused on G2 from Gensym, Jgrafchart from Lund University, and XMLBlaster as integrating technologies, with all tools communicating via a distributed architecture. Individual tools could be written using Java, C++ or even Prolog, but would communicate in common ways.

Strong, focused project management. It was clear that the project had been given strong direction by a Management Committee with a clear view of what they wanted to achieve, and they are to be congratulated on the degree to which they have achieved it. While there is still plenty of further research to be done in improving the monitoring, diagnosis and repair of chemical plants, the CHEM project has produced an important set of tools that can be used to address problems in the domain.

The CHEM project has made many of the details of the project available. If you are interested in further details, then a good first place to look is on the CHEM website at http://www.chem-dss.org/.

3.2 Evaporation Station

Carlos Alonso Gonzalez and Belarmino Pulido
University of Valladolid
Spain

The Evaporation Station was chosen as it is an important system and because it is common to several industries / industrial processes. Its main goal is to increase the concentration of a liquid or ‘juice’, by a series of evaporating tanks or boilers. The monitoring of this process faces problems because the systems are comprised of many different interconnected units each of which uses a great deal of energy. The monitoring is further complicated by the fact there are two flows that require monitoring; firstly the ‘juice’ which needs to be concentrated and secondly the steam that is used to heat it.

One of the major problems concerning this reference problem is that, unfortunately, the new solutions are not fully tested yet. So this section deals with how the problems are approached and describes the areas that have already been tested.

The industrial company (Group) who provided this reference problem can, of course, model these systems and they already have a simulation model that was built to investigate sugar factories. The testing of these models has shown that they very accurately reflect the behaviour of the industrial systems. Present system monitoring techniques do use numerically based models. The Group has developed special models which manage the plant by breaking it down into sections and modelling these sections individually.

The Group spent a lot of time working with the process engineers to identify what were the main areas / problems that they were concerned with. Luckily these were very similar to the faults that they had assumed would be the issues, i.e. leakages, blockages, etc. Mostly they are concerned with pressure in different vessels and the flow of the juice between these vessels. The process engineers want to be informed of any changes to the pressure or temperature of these units and what will be the overall output flow condensable from evaporators. Tests have shown that changes in vapour from the boilers and variations in the pressure of the output condenser are two of the most important areas in achieving the optimal output flow.

Some of the aforementioned effects can be modelled and can illustrate important areas for control, for example, if the efficiency of the thermostatic controllers in the boilers is increased by 30% then the efficiency of the juice production is increased by 50%. Another important factor is the way in which faults propagate through the whole system and minor faults in temperature or pressure of some units can propagate through and cause major reductions in the efficiency of the plant as a whole. The Group is quite familiar with these sorts of problems and their approach, which is initially at a high level, is to separate out different problems to identify the underlying fault / faults.

This system has already been tested in factories but the Group is aiming to develop it, in collaboration with the University of Valladolid, to include consistency-based diagnosis techniques and knowledge-based techniques which require formalisation of the approaches. However, the approach they are utilising is a trusted one and it has already proved its value in tests on other systems. Currently, the factory process is automated but it includes no diagnosis at all. The system currently operates with monitors and alarms but these are at a very basic level. They do not even have protocols for major incidents. The different units are monitored and it is up to the factory manager to decide when and what to do in the case of any problems, or use their personal experience to order maintenance to avoid possible problems. However, in the Group’s experience it is quite hard to convince factory management of the added value of automated diagnostic tools.

The Group has become aware of the difference between dealing with operators and dealing with engineers. They have found that you have to be mindful of who you are targeting the system at. Experiences from the Tiger Application also drew the same conclusions. Operators are usually only concerned with changing set points and maintaining the operation during their shift, they are not really concerned with what the faults are or what needs to be maintained or fixed. Thus it is usually necessary to target the system at engineers who are concerned with the system operation as a whole. One of the Bridge members had had the experience of a plant headquarters telling them that there was no point in them spending money on knowing what was wrong with the plant because even if they knew they would not spend the money in fixing it.

One issue that this group has encountered is the resistance to spending money to improve ‘efficiency’. The plant operators seem to believe that the plant operates at a certain level at a certain cost and as these factors are ‘acceptable’ it does not need to take the issue further. One counter to this argument would be that whether they know it or not they do actually care about efficiency, or at least will have to, at some point in the future. Whether it is because of the increasing price of fuel or raw materials, changes in the market or due to a new competitor, they will have to change how they view ‘efficiency’. This resistance is, however, understandable and is primarily due to the complexity of running a lot of these processes. They are complex processes and the continuous running of the plant is the prime, and sometimes only, concern.

3.3 Tiger - Gas Turbine Condition Monitoring

Rob Milne
Intelligent Applications
Edinburgh
Scotland

Tiger is an application which uses a real world data stream for diagnosing gas turbines. The purpose is continuous fault detection and diagnosis of any faults. It is installed on over 50 gas turbines worldwide.

One of the first issues to state is that compared to the original formulation of Tiger, the overall goal of the application has changed. The original primary goal of Tiger was to do fault diagnosis, i.e. to see there is a trip on the turbine and to diagnose the cause of this trip. The focus has now shifted from diagnosis towards fault detection in support of a global view of the turbine in order to show everything that is wrong with a turbine at any one moment. This enables planning of maintenance. Tiger does this through analysing 600 inputs every second, 24 hours per day for years on end; Intelligent Applications actually has over 130 years gas turbine operation data. Thus Tiger’s job is now to evaluate long streams of continuous data and build up a big picture of all the things that are wrong. The way Intelligent Applications has decided to do this is by detecting the smallest intermittent problems and then moving up to the bigger picture. This section will outline how this has been done, as well as some of the problems they have faced and choices that they have had to make.

The biggest cost for the operation of a Gas Turbine is fuel, second to that is maintenance. One factor that they have noticed is that to be able to tell a customer why the turbine shut down is worth a certain amount of money to them (this may well be large). However, the maintenance costs can run into millions of dollars (a factor of 10 greater than the ‘shut down’ cost) and so if the application can help reduce and/or plan the maintenance work then this is worth considerably more to customers and results in them saving a much greater amount of money through using the application. A key to being able to achieve this is Tiger’s continuous monitoring; so engineers can look at the turbine and see what is different in the way it is running today versus the way it was running yesterday. And, importantly, how is it different from last week or last month.

The Tiger application is not restricted to any one turbine and is designed to analyse data collected from a series of turbines to locate a fault that may be external to any single turbine itself. The aim is to pick up the smallest manifestation of a problem and the question that they have to address is “what is interesting”. This is the first process decision that was taken. They chose to pick up these small manifestations, even ones that on a global scale were not interesting. Their explicit strategy was not to tune out small elements for robustness but to go exactly the other way and pick up a lot of potential faults; this results in a global picture of all of the potential faults.

The huge number of messages generated by Tiger, although useful for long range condition monitoring, create a vast amount of data that would be too huge for a user to understand. Therefore Tiger uses 'filters' in order to only show the user what they need to know whilst at the same time still processing all of the data.

So, to describe what Tiger actually does on a day to day basis, the messages are checked for any newly developed problems. Potentially important issues are further diagnosed and problems which are now routine are 'filtered'. The route Intelligent Applications has taken is to have very sensitive fault detection to give them a big picture of what the turbine is doing during daily operation. This is done to support the goal of predicting what faults on the turbine would require maintenance or longer term care. All of this has grown out of a system that was designed to do diagnosis, i.e. something happens and it tells you what has gone wrong.

Another approach to diagnosis they looked at was the use of Neural Networks, they do not, however, use these in Tiger because the turbines are affected by inlet air conditions. Air temperature changes a lot between summer and winter and this change, over a twelve month period, means that it is problematic to train the Neural Nets due to this fluctuation over time, due to the lack of base line data. This is also compounded by the fact that they would need to have a 'standard' turbine, i.e. just post service and a range of inlet air conditions and these circumstances do not occur often enough to provide the required data.

Tiger's main task can be seen as abstracting the raw data into messages about what is interesting in the data. The diagnostics then aggregate these messages in various ways for the user to use in support of their key tasks.

The tuning of Tiger's thresholds is based on a statistical analysis of historical data, with fine tuning by the engineers. This is a relatively 'gross' and broad task and can be done quickly for each turbine. Tiger does not use tight mathematically based models because the overhead of tuning so many models for so many turbines currently requires too much man effort. The structure of each turbine, however, is very similar to each other turbine. Previous versions of Tiger had used predictive envelopes and model-based fault isolation (Sheba), but this is currently not being used due to the amount of tuning per turbine required.

Tiger now has a 'near' competitor in the US. Their approach is based on predictive models and checking residuals. So far they have worked with pumps, fans, etc. and not gas turbines. Other aspects of Tiger are now stable enough that it is planned to look again at the tuning needed to support more concise models.

The prime benefit to the customer is to be able to know everything that is going on with their turbine and when the time comes to deal with it you know what problems to focus on. The most valuable output from Tiger for the customer is a report that shows a table of all the faults that Tiger has picked up, a description of them and a suggestion of what remedial action to take. This is created by engineers using information from Tiger.

The wider issue with regard to the reference problems is the ability to look at large amounts of data that have real dynamics and real variations to them. This is important as so many applications are tested on 'Toy' models, i.e. created data, and these often do not include the level of noise that is found in real systems and the occurrence of things that should not really happen. It is also an interesting exercise in what problems are out in the real world, for instance even a simple act such as setting thresholds becomes very difficult with widely varying data. The Tiger Reference Problem also shows interesting issues with regard to how you 'filter' data and how you decide what is interesting to the user and what is not. The prime benefit is that there is too much raw data to be looked at by the operator, and with Tiger an operator or plant manager can, at any time that is convenient to them, print off the messages from Tiger and see what is wrong with their turbine.

3.4 Structural Diagnosis of Bridges

Benny Raphael
Swiss Federal Institute of Technology (EPFL)
Lausanne
Switzerland

This reference problem concerns structural diagnosis of bridges; it uses model composition and global search in order to diagnose damage to bridges. Structures are increasingly being fitted with measurement equipment so that engineers may acquire a better understanding of real behaviour during the structure's service life. These measurement systems provide data such as deflections, slopes and internal deformations. The approach that is followed involves finding models whose predicted responses closely match measured values. Through an examination of the characteristics of models that match measurement data, abnormalities can be detected.

A library consisting of fragments of models represent possible behaviours of relevant elements of bridges. For example, there are model fragments for representing support conditions, material properties, loading, geometric properties, etc. All the model fragments interact to provide the overall behaviour of the structure. The set of possible models is defined by all valid combinations of model fragments that are stored in the fragment library. The library contains fragments that represent structural components that function normally as well as those that exhibit faulty or gradually deteriorating behaviour. The objective is to identify complete models (which are combinations of fragments) whose predictions closely match measured values.

There are three important modules in this approach;

- Model Composition
- Global Search
- Measurement System Configuration

The model composition module automatically generates valid combinations of models that can then be analysed by methods such as finite element analysis. Only fragments that are compatible with possible sets of assumptions specified by users are used in model composition.

One of the problems with this approach is the huge number of possible combinations of the model fragments; therefore, an exhaustive search of all combinations is not feasible. Therefore, a global search algorithm is used for identifying models whose predictions closely match measurements.

There are a number of different combinations of fragments that might lead to the same response and the procedure could identify wrong models. In order to avoid this problem, the measurement system configuration module carefully identifies the locations and types of sensors that result in maximum separation between candidate models. Since there are situations where additional sensors do not improve model separability, it is not reliable to identify a single model as the solution. Therefore, a population of models whose predictions closely match measurements is selected. Common characteristics of these models are used to group these models into model classes. When there are many model classes containing widely varying properties, it is not possible to provide definite diagnostic assessments. In such cases, the measurement system is reconfigured in order to improve the reliability of diagnosis.

3.5 DAMADICS Actuator

Teresa Escobet
Universitat Politecnica de Catalunya
Spain

The set of faults proposed in the DAMADICS actuator benchmark reference problem are types of actuator failure modes, and the proposed objectives are fault detection, fault isolation and fault identification. The fault scenarios include both, abrupt and incipient faults, since the monitoring of the development of incipient faults is an issue not only for predicting maintenance schedules but also for monitoring the performances of the process.

An advantage of this reference problem is that it uses real process data. The industrial actuator has also been instrumented to undergo faults whilst in real process operation. This represents a significant step forward in this field. The reference problem may be applied by engineers either when implementing new diagnostic schemes or for evaluating existing schemes. For researchers, the reference problem also permits to point out those FDI algorithms that are based on unrealistic assumptions and constraints and are therefore 'not applicable'.

Most of the proposed solutions belong to the field of soft-computing. Methods that combine neural and fuzzy approaches have been developed to give solutions to the proposed problem. In general, these techniques use neuronal networks as residual generators and fuzzy reasoning for the identification of the failure modes (symptom evaluation). It is important to realise that neural networks do not require an explicit mathematical model of the system being modelled / monitored. They can also operate simultaneously on qualitative and quantitative data and they are readily applicable to multivariable systems. The papers that focus on neural networks or neuro-fuzzy approaches to the DAMADICS reference problem are: Rzepiejewski *et al.* ; Patan *et al.* ; Mrugalski *et al. ;* Uppal *et al.* ; Papadimitropoulos *et al.* ;and Calado *et al.,* 2003.

There are also a set of proposed solutions which belong to the model-based approach along the FDI line. Different modelling techniques have been tested to achieve the project goals. In general these solutions can be classified as active or passive methods. The passive methods used are: linear and non-lineal interval observers (Puig et al, 2003 and Stancu *et al.,* 2003,) and quantised modelling approaches (Lunze *et al.*2003). The active approaches (Uppal *et al.,* 2003; Witczak *et al.,* 2003) use soft computing methods to perform robust designs. Uppal *et al.* (2003) use a novel neuro-fuzzy network technique for designing multiple-model observers. Witczak *et al.* (2003) use a genetic algorithm approach to identify state-space models with a suitable structure for robust fault detection.

Only one solution, using model-free fault detection technique based on the use of a specific spectral analysis tool, know as Squared Coherency function, has been proposed (Previdi *et al.,* 2003).

Finally, a structural analysis approach has been used to determine the potentials for delectability and isolability. The possible available sensors that allow the maximum diagnosability degree, using models of normal behaviour, were derived in Travé-Massuyès et al, 2003 following this approach. These take into account models of abnormal behaviour and increase the delectability degree (Sataroswiecki et al, 2004, Frisk et al, 2003).

Performance indexes allow one to compare FDI algorithms with respect to basic features of diagnostic decisions (isolation and detection). The performance indexes are described in *Benchmark definition*, available at: *http://diag.mchtr.pw.edu.pl/damadics/*.

Detection / Isolation Indexes		**General diagnosis indexes**	
Detection / Isolation Time	t_{dt}, t_{it}	Theoretical Diagnosis Accuracy	$dacc^{i}_{t}$
Detection / Isolation Recovery Time	t_{drt}, t_{irt}	Theoretical Mean Diagnosis Accuracy	$dacc_{tm}$
Detection / Isolation Moment	t_{dm}, t_{im}	Diagnosis Accuracy	dacc
Detection Recovery Moment	t_{drm}, t_{irm}		
False Detection / Isolation Rate	r_{fd}, r_{fi}		
True Detection / Isolation Rate	r_{td}, r_{ti}		
Mismatch Isolation Rate	r_{mi}		
Fault Detection / Isolation Sensitivity Factor	fs_{d}, fs_{i}		

3.6 AOCS: Attitude and Orbit Control System

Louise Travé-Massuyès
LAAS – CNRS
Toulouse
France

The satellite Attitude and Orbit Control System reference problem focuses on the operation of sensors and actuators as this is where most of the faults are based. This description will discuss how the diagnosis is done on the satellite and what has been proposed and experimented on, with this system. One factor that makes this reference problem different from the others is that the ultimate target is that of creating an autonomous system at the end of the project. This fact puts the diagnosis 'in the loop', with the diagnosis being used to reconfigure the system and this must also be done in an autonomous way. Thus, it is not necessarily important to provide a very precise diagnosis, but what you do need is to generate the correct sequence of actions that will result in the system being in an acceptable situation.

In order to do this, with the required level of autonomy, a diagnosis must be issued at every time point but more importantly the system state must be tracked in time. Thus we are left with a sequence of observations and the problem then lies in generating a sequence of 'states' that match these observations. The solutions space scientists are currently using on the real world satellites rely on a very pragmatic approach based entirely around the setting of threshold values; however this is done in a very complex way. They take a lot of data from experts in the domain, who are aware of the underlying models, and derive the thresholds to be set; particular attention is paid to those at the functional level. They also have component level thresholds, temperature for example, which state the condition of the environment of the components. However, the functional thresholds are the major issue as they are set for every operating mode of the satellite. Presently this is all done by hand.

The solution that LAAS-CNRS (in collaboration with LIPN) suggested was a model-based approach (implemented in the KOALA system) (Benazera 2003). In order to be able to deal with different operating modes and to be able to track the system they needed to be able to define fault modes. This was therefore all represented in a discrete event level with automata. Associated with every state of the automata were the continuous models of the behaviour of the system in each particular mode (qualitative or interval models), the whole resulting in a hybrid model (Bénazéra *et al*. 2001, 2002). The main task of the group was to provide an interface between the continuous world and the discrete world so that continuous states could be projected at the discrete event level - allowing them to hence set the algorithms at the discrete event level (Bénazéra and Travé-Massuyès 2003). The actual solution was tested on a few scenarios and the algorithms that were tested were more on fault detection than on the real state tracking due to data limitations.

The work undertaken here was built upon the work done at NASA (Williams et al. 1998) and extends it to deal with hybrid models. The main characteristic of this approach was to have the double representation of uncertainties: at the discrete event level, with the uncertainties being dealt with by probabilities, i.e. the set of transitions from one mode have an associated probability distribution, and at the continuous level the uncertainty being represented by bounded (interval) models. The probabilities serve to focus the state tracking so that the diagnosis system does not require to be exhaustive.

There are other approaches which work in a pure stochastic framework (Haufbaur and Williams 2002), attaching Kalman filters to the automata states. In these, the uncertainties are unified more easily than in this group's approach. However, they do not benefit of the automatic generation of adaptive thresholds for fault detection as with interval models (Tornil *et al.* 2001). Finally, another trend goes with particle filters (Dearden and Clancy 2002) (Dearden and Hutter 2003), the difficulty being here to manage the high number of necessary particles and to propose smart reweighting procedures.

The main reason that discrete events are included in this approach is purely because they want to represent fault modes and to be able to track them. All of the control actions are taken into account but one of the difficulties of the problem is that, although all of the actions are observable, there are some mode transitions that are triggered by continuous variables and these are not necessarily observable.

4 Discussion

This section analyses the solutions that were proposed for every reference problem and draws some general conclusions about the trends of diagnosis systems.

The solutions to our reference problems can be split into two classes:

- those that are used at the industrial level (TIGER for gas turbines, Evaporation station)
- those that have been proposed by researchers and that have just been tested at an Industrial level (Steam generator, DAMADICS Actuator, Satellite AOCS, Structural diagnosis of Bridges)

The first category is far beyond the state of the art of the diagnosis technologies, what is achieved is generally done in a much more pragmatic way, putting the focus on overall performance and reliability rather than on optimised diagnosis solutions.

The reference problems that have been used as benchmarks in collaborative projects and that have been used for testing several diagnosis tools - the Actuator in the DAMADICS project, the Steam generator plant in the CHEM project - point out the difficulty of comparing different approaches on practical bases:

- the underlying assumptions are generally different
- the kind of knowledge that is necessary is also different (for example for dealing with uncertainties, statistical methods for residual evaluation need knowledge about

noises and disturbances, whereas bounded model approaches require knowledge about the model parameter acceptable deviation)

The comparative analysis should not only be performed analysing the provided diagnoses results but also, and most importantly, analysing what it takes to provide a solution. On the other hand, many of the existing approaches are complementary rather than in competition. This later point is very well illustrated by the tests performed within the CHEM project; in many instances, a number of the tools were linked together to produce a specific diagnosis result.

In summary, it is difficult to conclude that one approach is better than another and it is good to see that the level of understanding of the different contributions that can be made by FDI techniques and by model-based AI techniques is increasing in both communities. Today's most advanced solutions, such as hybrid model-based solutions, show that the solutions of tomorrow will, through necessity, be integrated; and will form the global solution with partial solutions from both sides.

5 References and Background Texts

G. Acosta, C. Alonso González, B. Pulido. Basic tasks for knowledge-based supervision in process control. Engineering Applications of Artificial Intelligence. Vol. 14, pp. 441-455. 2001.

G. Acosta Lazo, C. J. Alonso González, B. Pulido Junquera. Diagnosis basada en conocimiento de un proceso azucarero con TEKNOLID. International Sugar Journal, Vol. 103, nº 1225, pp. 44-51. 2001.

C. J. Alonso González, G. Acosta Lazo, J. Mira Mira, C. de Prada Moraga. Knowledge based process control supervision and diagnosis: the AEROLID approach. Expert Systems with Applications, Vol. 14, pp. 371-383. 1998.

C. J. Alonso González, B. Pulido Junquera, G. Acosta Lazo, C. Llamas Bello. On-line Industrial supervision and diagnosis, knowledge level description and experimental results. Expert Systems with Applications, Vol. 20, pp. 117-132. 2001.

E. Benezera. Diagnosis and reconfiguration based on hybrid concurrent models. Application to autonomous satellites. PhD Thesis of the Université Paul Sabatier, LAAS-CNRS, Toulouse France, 2003.

E. Benezera, L. Trave-Massuyes. The consistency approach to the on-line prediction of hybrid system configurations. IFAC Conference on Analysis and Design of Hybrid Systems (ADHS'03), Saint-Malo (France), 16-18 June 2003, 6p.

E. Benezera, L. Trave-Masuyes, P. Dague. State tracking of uncertain hybrid concurrent systems. 13th International Workshop on Principles of Diagnosis (DX'02), Semmering (Autriche), 2-4 May 2002, pp.106-114

E. Benezera, L. Trave-Massuyes, P. Dague. Hybrid model-based diagnosis for autonomous spacecrafts. 1st ESA Workshop on "On-board autonomy", Noordwijk (Pays-Bas), 17-19 October 2001, pp.279-286.

Calado, J.M.F., R. Louro, M.J.C. Mendes, Sá, Costa and M. Kowal (2003). Fault isolation based on HSFNN applied to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

Dearden R. and Clancy D. Particle Filters for Real Time Planetary Rovers. Proceedings of Thirteenth International Workshop on Principles of Diagnosis DX'02, Semmering, Austria, May 2-4, 2002. p. 1 – 6.

Dearden R. and Hutter F. The Gaussian Particle Filter for Diagnosis of Non-Linear Systems. Presented at Bridge Day at SafeProcess/DX'03. In Proceedings of the Fourteenth International Workshop on Principles of Diagnosis, DX'03, Washington D.C., June 2003. p. 65 - 70.

Frisk, E., D. Dustegor, M. Chrysander and V. Cocquempot (2003). Improving isolability properties by structural analysis of faulty behavior models: application to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPR OCESS '2003,* Washington, USA.

M. W. Hofbaur and B. C. Williams. Mode estimation of probabilistic hybrid sys-tems. In C.J. Tomlin and M.R. Greenstreet, editors, Hybrid Systems: Computation and Control, HSCC 2002, volume 2289 of Lecture Notes in Computer Science, pages 253–266. Springer Verlag, 2002.

Lunze, J. and P. Supavatanakul (2003). Application of timed automata to the diagnosis of the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

R. Milne, C. Nicol, M. Ghallab, L. Trave-Massuyes, K. Bousson, C. Dousson, J. Quevedo, J. Aguilar and A. Guasch, Real-Time Situation Assessment of Dynamic Systems, *Intelligent Systems Engineering Journal*. Autumn, 103-124 (1994).

R. Milne, C. Nicol, L. Travé-Massuyès, J. Quevedo. Tiger™: Knowledge Based Gas Turbine Condition Monitoring. *AI Communications Journal,* Vol 9, N^o 3, 1-17. Publishers IOS Press, 1996.

R. Milne and C. Nicol. Tiger™: Continuous Diagnosis of Gas Turbines. *In Proc. of ECAI/PAIS'00*, edited by Werner Horn, Berlin, 2000.

Mrugalski, M., E. Arinton and J. Korbicz (2003), Fault detection with dynamic GMDH neural networks: application to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

MONET Deliverable: MBD3 Collection of Industrial Reference Problems. Available on-line at http://monet.aber.ac.uk:8080/monet/docs/tg_minutes_and_reports/bridge/mbd3.pdf

Papadimitropoulos, A., G.A. Rovithakis and T. Parisini (2003). Neural approximators for fault detection of actuators in the presence of friction: the Case of the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPR OCESS '2003,* Washington, USA.

Patan, K. and T. Parisini (2003). Dynamic Neural Networks for Actuator Fault Diagnosis: Application to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

Previdi, F. and T. Parisini (2003). Model-free actuator fault detection using a spectral estimation approach: the case of the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

Puig, V., Quevedo, J., Escobet, T., Stancu, A. (2003) Robust fault detection using linear interval observers. In: preprints of 5-th IFAC Symposium SAFEPROCESS '2003, Washington, USA.

Rzepiejewski, P., M. Syfert and S. Jegorov (2003). On-line actuator diagnosis based on neural networks and fuzzy reasoning: The *DAMADICS* Benchmark Study. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

Stancu, A., V. Puig, J. Quevedo and R.J. Patton (2003). Passive robust fault detection using non-linear interval observers: Application to the *DAMADICS* Benchmark Problem, In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

Staroswiecki, M. and B. Ould Bouamama (2004). FDI of smart actuators using bond graphs application: DAMADICS valve. In: 5th Workshop on Presentation of Final Results, April 5th-7th, Łagón, Poland.

S. Tornill, T. Escobet, L. Trave-Massuyes, E. Benezera. Robust interval model based fault detection. Satellite attitude and orbit control system case study. LAAS internal report N°01469, October 2001, 6p.

L. Trave-Massuyes and R. Milne, Diagnosis of Dynamic Systems Based on Explicit and Implicit Behavioural Models: An Application to Gas Turbines in ESPRIT Project TIGER, *Applied Artificial Intelligence*. Vol 10, N° 3, 257-277. Publishers: Taylor & Francis (1996).

L. Travé-Massuyès and R. Milne. Gas Turbine Condition Monitoring Using Qualitative Model Based Diagnosis*. IEEE Expert magazine, 'Intelligent Systems & Their Applications',* Vol 12, N° 3, 21-31. Publishers: IEEE Computer Society, 1997.

Travé-Massuyès, L., Escobet, T., Spanache, S. Diagnosability analysis based on component supported analytical redundancy relations. In: preprints of 5-th IFAC Symposium SAFEPROCESS '2003, Washington, USA.

Uppal, F.J., R.J. Patton and M. Witczak (2003). A hybrid neuro-fuzzy and de-coupling approach applied to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFA C Symposium SAFEPROCESS '2003,* Washington, USA.

B. C. Williams, P. Nayak, and N. Muscettola. Remote agent: To boldly go where no AI system has gone before. Artificial Intelligence, 103(1-2):5–48, August 1998.

Witczak, M., R.J. Patton and J. Korbicz (2003). Fault detection with observers and genetic programming: Application to the *DAMADICS* Benchmark Problem. In: preprints of *5-th IFAC Symposium SAFEPROCESS '2003,* Washington, USA.

6 Document History

Version	*Date*	*Changes made to document*	*Changed by*
1.0	10th May 2004	First draft made from discussions in Grenoble.	RIR
1.1	19th May 2004	Minor amendments and proof-reading	JNT
1.2	27th May 2004	Damadics section from Teresa Escobet included	RIR
1.3	11th June 2004	Amendments incorporated from Rob	JNT

Version	*Date*	*Changes made to document*	*Changed by*
		Milne	
2.0	19th July 2004	Global Amendments and addition of CHEM section	RIR
2.1	4th August 2004	Global amendments and addition of the discussion section	LTM
2.2	5th August 2004	Addition of Bridge Diagnosis section from Benny Raphael	RIR
2.3	17th August 2004	Minor Amendments from Rob Milne	RIR
	1st September 2004	Status Set to release	RIR