



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Team DREAM

*Diagnosis, REcommending Actions and
Modelling*

Rennes

THEME 3A

Activity
R *eport*

2003

Table of contents

1. Team	1
2. Overall Objectives	1
3. Scientific Foundations	2
3.1. Physical systems monitoring aid	2
3.2. Machine learning	4
4. Application Domains	6
4.1. Panorama	6
4.2. Telecommunication networks	6
4.3. Decision aiding in medicine and health-care	6
4.4. Environmental decision making	7
6. New Results	8
6.1. Diagnosis of large scale discrete event systems	8
6.1.1. A decentralized diagnosis approach	8
6.1.2. The use of model-checking techniques	8
6.2. Cardiac monitoring	9
6.3. Learning diagnostic rules for telecommunication networks	9
6.4. Diagnosis of a synchrotron X-ray beamline	10
6.5. Learning decision rules in the water ressource management domain	10
6.6. Action languages and default inferences	11
7. Contracts and Grants with Industry	11
7.1. Magda2: Modelling, diagnosing and supervising telecommunication networks	11
7.2. Sacadeau: Decision-aid to improve streamwater quality	12
7.3. Cepica: Conception and Evaluation of an Implantable Cardiac Device	12
8. Other Grants and Activities	12
8.1. National projects	12
8.2. International networks and working groups	13
8.3. International bilateral relations	13
9. Dissemination	13
9.1. Animation of the Scientific Community	13
9.1.1. Journal editorial board	13
9.1.2. Conference program committees	13
9.2. Faculty teaching	13
10. Bibliography	14

1. Team

Head of project

Marie-Odile Cordier [Professor, University of Rennes 1]

Administrative assistant

Marie-Noëlle Georgeault [AA Inria]

Inria staff members

Yves Moinard [Research scientist]

René Quiniou [Research scientist]

Faculty members

Véronique Masson [Associate professor, University of Rennes 1 (50 % time)]

Dominique Py [Associate professor, IUFM Rennes 1 until june 2003]

Sophie Robin [Associate professor, University of Rennes 1]

Laurence Rozé [Associate professor, Insa Rennes (on maternity leave from february to september)]

Ph. D. students

Élisa Fromont [INRIA Ph. D. student, 1st year]

Alban Grastien [MENRT Ph. D. student, 1st year]

François Portet [MENRT Ph. D. student, 1st year]

Research scientists (partners)

Philippe Besnard [DR CNRS, IRIT, Toulouse]

Christine Largouët [MC, Ensar Rennes (on leave at the University of New-Caledonia since june 2003)]

2. Overall Objectives

Key words: *diagnosis, supervision, machine learning.*

The research objectives of the Dream team are about aiding monitoring and diagnosing time evolving systems. The main issue is to help the person in charge of the system by analyzing the observations provided by sensors and giving her/him information about diagnosis hypotheses (potential anomalies or failures) and recommended actions. Qualitative model-based approaches are advocated for at least two main reasons:

- they are "white-box" approaches and consequently diagnoses and recommended actions can be explained to the user in an explicit and adequate language,
- they are flexible enough and are then adapted to quickly evolving systems such as technological systems (for instance telecommunication components).

We use a model-based approach relying on normal and faulty behavioral models. These models are discrete-event models such as (temporal) communicating automata, temporal causal graphs or chronicles.

In this context, two main research themes are developed:

1. Classical model-based diagnosis methodologies cannot be directly used for complex systems due to the intractable size of the model and the computational complexity of the process. It is especially true when on-line diagnosis is considered. Two solutions are investigated:
 - We propose a decentralized approach which relies on combining local diagnoses built from local models (or local diagnosers). Three problems are currently investigated: Which strategy should be used for an optimal merge of the local diagnoses in order to preserve the efficiency and the completeness of the process? How the process incrementality can be ensured in an on-line diagnosis context where observations are incrementally collected? How let the diagnosis process deal with reconfigurable systems the topology of which can be changed at running time.

- We propose to use model-checking techniques in order to improve the efficiency of the computation and to cut down the combinatorial state explosion. It means using adequate symbolic representations as BDD for instance and partial order reduction techniques taking advantage of the existing inherent concurrency.
2. It is well recognized that model-based approaches suffer from the difficulty to acquire the model. It is why we focus on automatically acquiring models from data with symbolic learning methods coupled with data mining methods. One of the challenges we tackle is to extend existing inductive logic programming methods (ILP) to temporal data in order to be able to deal with data coming from signals (as electrocardiograms in the medical domain) or alarm logs (in the telecommunication domain). Two problems are currently investigated: how to adapt the learning process to deal with multiple sources of information (multi-sensor learning)? how to integrate signal processing algorithms to the learning or diagnosis task when this latter relies on a qualitative description of signals?

Our application domains are the following:

- industrial applications with a focus on telecommunication networks;
- medical applications and especially cardiac monitoring, i.e the on-line analysis of cardiac signals to detect arrhythmias and the development of "intelligent" cardiac devices (pacemakers and defibrillators) having some signal analysis and diagnosis capabilities;
- environmental protection, and more precisely the development of decision support systems to help the management of agricultural plots with the objective of preserving water quality threatened by pesticide pollution is the challenge.

3. Scientific Foundations

3.1. Physical systems monitoring aid

Key words: *monitoring, diagnosis, deep model, fault model, simulation, chronicle recognition, temporal causal graph, chronicle acquisition.*

Glossary

alarm a discrete indicator emitted by a monitoring system from events, that is supposed to provoke a reaction, either a human one or an automatic one.

chronicle (or scenario) a set of time-stamped events which are related by temporal constraints characterizing a situation.

chronicle recognition a system which, from a set of chronicles describing situations (the chronicle base), allows to analyze on-line a sequence of dated observations and to recognize the situations.

Our work on monitoring and diagnosis relies on model-based approaches developed by the Artificial Intelligence community since the founding studies by R. Reiter and J. de Kleer [49][27]. Our project investigates the on-line monitoring and diagnosis of systems, which are modelled as discrete events systems, focusing more precisely on monitoring by alarms management [56]. Computational efficiency is a crucial issue for real size problems. We are developing two approaches. The first one relies on diagnosers techniques [54], for which we have proposed a decentralized and generic approach. The second one uses chronicle recognition techniques, focusing on learning chronicles.

Early work on model-based diagnosis dates back in the 70-80's by R. Reiter, the reference paper on the logical theory of diagnosis being [49][27]. In the same years was constituted the community known as DX, named after the *workshop* on the principles of diagnosis. The research in these areas are still very active and the

DX workshop gathers about fifty people in the field every year. Contrary to the expert system approach, which has been the leading approach for diagnosis (medical diagnosis for instance) before 1990, the model-based approach lies on a deep model representing the expected correct behavior of the system to be supervised or on a fault model. Instead of acquiring and representing an expertise from experts, the model-based approach uses the design models of industrial systems. The approach has been initially developed for electronic circuits repair [28], focusing on off-line diagnosis of so-called static systems. Two main approaches have been proposed then: (i) the consistency-based approach, relying on a model of the expected correct behavior, which aims at detecting the components responsible for a difference between the expected observations and the really observed ones; (ii) the abductive approach which relies on a model of the failures that can affect the system, and which identifies the failures or the faulty behavior explaining the anomalous observations. See the references [32][37] for a detailed exposition of these investigations.

Since 1990, the researchers in the field have studied the monitoring and the diagnosis of dynamic systems, which made them closer to the researchers in automatic. What characterizes the IA approach is the use of qualitative models instead of quantitative ones and the importance given to the search for the real origin of the faulty behavior. Model-based diagnosis approaches rely on qualitative simulation or on causal graphs in order to look for the causes of the observed deviations. The links between the two communities have been enforced, in particular for what concerns the work about discrete events systems and hybrid systems. The used formalisms are often similar (automata, Petri nets, ...) [21][56].

Our team focuses on monitoring and on-line diagnosis of discrete events systems and in particular on monitoring by alarm management. In this context, a human operator is generally in charge of the system monitoring and receives events (the alarms) which are time-stamped and emitted by the components themselves, in reaction to external events. These observations on the system are discrete informations, corresponding to an instantaneous event or to a property associated to a time interval. The main difficulties for analyzing this flow of alarms are the following:

- the huge number of received alarms: the supervisor may receive till several hundreds of messages per second, many of which being insignificant,
- the alarm overlapping: the order in which alarms are received may be different from the order in which alarms were emitted. Moreover, various sequences of alarms resulting from concurrent failures may overlap. The propagating delays, and sometimes the ways the alarms are transmitted, must be taken into account, not only for event reordering, but also to decide at what time all the useful messages can be considered as being received.
- the redundancy of received alarms: some alarms are only routine consequence of other alarms. This can provoke a phenomenon known as cascading alarms.
- the alarm loss or alarm masking: some alarms can be lost or masked to the supervisor when an intermediate component in charge of the transmission is faulty. The absence of an alarm must be taken into account, since it can give a useful information about the state of the system.

There are two cases focusing on very different issues. In the first one, the alarms must be dealt with, *on-line*, by the operator. In this case, alarm analysis must be done in real time. The operator must react in a very short period of time to keep the system working at the best in spite of the inputs variability and the natural evolution of the processes. Consequently, the natural system damages (components wear, slow modification of the components properties, etc.) are not directly taken into account but are corrected by tuning some parameters.

This *reactive* treatment withstands the treatment of alarms maintenance. In this second case, a deeper *off line* analysis of the system is performed, by foreseeing the possible difficulties, by planning the maintenance operations in order to minimize significantly the failures and interruptions of the system.

The major part of our work focuses on on-line monitoring aid and it is assumed that the correct behavior model or the fault models of the supervised systems are available. However, an on-line use of the models is rarely possible because of its complexity with respect to real time constraints. This is especially true when

temporal models are concerned. dimension of these models. A way to tackle this problem is to make an off-line transformation (or compilation) of the models and to extract, in an adapted way, the useful elements for diagnosis.

We study two different methods:

- In the first method, the automaton used as a model is transformed off-line into an automaton adapted to diagnosis. This automaton is called a *diagnoser*. The transitions of the automaton are only triggered by observable events and the states contain only information on the failures that happened in the system. Diagnosing the system consists in going through all the different states of the diagnoser as observable events become available. This method has been proposed by M. Sampath and colleagues [54]. We have extended this method to the communicating automata formalism [52] (see also [50]). We have also developed a more generic method which takes advantage of the symmetries in the architecture of the system [51].

Our more recent work deals with a decentralized approach [44]. This approach can be compared to R. Debouk and colleagues [26][25] and also to P. Baroni and colleagues [20][19]. Our method, unlike R. Debouk et al., relies on local models. We do not need to construct a global model. Indeed, the size of the global model would have been too important in our applications. Even if the methods are very close, P. Baroni et al. are concerned with an *a posteriori* diagnosis (off-line) whereas we propose an on-line diagnosis. Each time an alarm comes, it is analyzed and the diagnosis hypotheses are incrementally computed and given to the operator. Our main theme of study is close to E. Fabre and colleagues [35][18]. The main difference is that they propose a multi-agent approach where the diagnoses are computed locally at the component level using message exchanges, whereas we construct a global diagnosis which is given to the operator at the supervisor level.

- In the second method, the idea is to associate each failure that we want to detect with a *chronicle* (or a scenario), i.e. a set of observable events interlinked by time constraints. One way to supervise dynamic systems is to recognize those chronicles on-line. The principle is to follow the possible chronicles corresponding to a set of received failure messages until finding one or several chronicles that satisfy all the constraints. To perform this task, we have to create a chronicle base that contain all the possible chronicles. This base must be updated each time the supervised system evolves physically or structurally. An expert is needed to create the chronicle base. However, this makes the maintenance of the base very expensive. That's why we prefer to use an automatic method to learn the base. Most of the studies on chronicle recognition are French [31][48][29] and are based on C. Dousson's thesis [30]. There is quite a lot of activity around this theme in France: a workshop on chronicle has taken place in Lannion in 2002 and gathered about 20 persons. Applications generally deal with system monitoring (telecommunication network), video-surveillance (underground, bank, etc...). Our research studies do not focus directly on the development of chronicle recognition systems but on the automatic acquisition of the chronicle base. This idea is developed in the next paragraph.

3.2. Machine learning

Key words: *Machine learning, Inductive Logic Programming (ILP).*

Glossary

literal an atomic formula or a negated atomic formula

clause a disjunct of literals

definite clause a clause containing exactly one positive literal

logic program a set of definite clauses.

The techniques investigated in the group aim at acquiring and improving models automatically. They belong to the field of machine or artificial learning [24]. In this domain, the goal is the induction or the discovery of objects characterizations from their descriptions by a set of features or attributes. Our work is grounded on Inductive Logic Programming (ILP).

A learning method is supervised if samples of the objects to be classified are available and labeled by the class they belong to. Such samples are often called learning examples. If the examples cannot be classified a priori, the learning method is unsupervised. Kohonen maps, induction of association rules in data mining or reinforcement learning are typical unsupervised learning methods. From another point of view, learning methods can be symbolic, such as inductive rule or decision tree learning, or numerical, such as artificial neural networks.

We are especially interested in structural learning which aims at making explicit relations among data where such links are not known. The temporal dimension is of particular importance in applications we are dealing with, such as process monitoring in health-care, environment or telecommunications. Additionally, we consider that the comprehensibility of the learned results is of crucial importance as domain experts must be able to evaluate and assess these results. ILP is the learning technique that best meets these requirements. We use a supervised version of this technique but also intend to use the unsupervised version which is called *Relational Data Mining* [33].

ILP began in the early 80's, though not under this name, when knowledge representation paradigms coming from logic programming began to be used in the field of machine learning. Such a high level language meets the needs of relational representations for the description of structured objects or true relations between objects

During the 90's, ILP has become a proper research topic at the intersection of domains such as machine learning, logic programming and automated deduction. The main goal of ILP is the induction of classification or prediction rules from examples and from domain knowledge. The ILP research field has been extended to data mining enabling the discovery of association rules describing the correlations between data descriptors. As ILP relies on first order logic, it provides a very expressive and powerful language for representing hypotheses as well as the domain knowledge, this is its major feature.

Formally, ILP can be described as follows: given a set of positive examples P and a set of negative examples N of some concept to be learned, a logical theory B called the background knowledge and a language L_H specifying which clauses are syntactically and semantically acceptable, the goal is to discover a hypothesis H in the form of a logic program belonging to L_H such that $\forall p \in P T \cup H \models p$ and $\forall n \in N T \cup H \not\models n$. This definition can be extended to multi class learning. From a computational point of view, the learning process consists in searching the hypothesis space, either top-down by refining clauses that are too general (that cover negative examples) by adding literals to clause body or bottom-up by generalizing clauses that are too specific (that do not cover enough positive examples) by deleting literals or transforming constants into variables in literals. An interesting property is that the clause space has a lattice structure which enables an efficient search.

ILP is mainly used for learning classification rules. Similar techniques can also be used for inducing decision trees as well as for first order regression. The goal of regression is to predict the value of a real variable instead of a class value. Some recent extensions deal with learning dynamic models: one such extension uses a representation coming from the qualitative simulator QSIM [38], another enables the discovery of differential equations from examples describing the behavior of a dynamic system [34].

Nowadays, work in ILP is mainly concerned with improving learning robustness (dealing with noisy or incomplete data) or efficiency (improving the search space exploration by taking structural properties into account, by stochastic techniques or by parallelizing algorithms for massively parallel computers). Another research direction investigates how to associate ILP to other learning methods which are more efficient for particular kind of data or to associate different learning strategies during ILP search. Extending the language to full first-order is also investigated. In this direction, learning from temporal data is of major interest because many application domains, such as telecommunications, health-care or environment, provide huge amounts of such data. This is why we have chosen to rely upon work by C. Rouveirol and M. Sebag [55] who have shown the value of associating ILP to CLP (Constraint Logic Programming) in order to compute efficiently numerical values. D. Page [42] wrote that a final challenge for ILP is to elaborate tight collaboration schemes

between experts and ILP systems for knowledge discovery in order to avoid their complexity i) by enabling the evaluation of alternative hypotheses and not only those that maximize some heuristic function, ii) by devising tests and experiments for choosing among several hypotheses, iii) by providing non numerical justifications of the hypotheses such as belief measures or illustrative examples, iv) by consulting the expert when anomalies are detected in the data.

Our work is more concerned with the application of ILP rather than developing or improving the techniques. Nevertheless, as Page and Srinivasan have noted [42], the target application domains (such as signal processing in health-care) can benefit from the adaptation of ILP to the particular features of the application data. Thus, we investigate how to associate temporal abstraction methods to learning and to chronicle recognition. We are also interested in constraint clause induction, particularly for managing temporal aspects. In this setting, some variables are devoted to the representation of temporal phenomena and are managed by a constraint system [47] in order to deal efficiently with the associated computations (such the covering tests, for example).

4. Application Domains

4.1. Panorama

Key words: *telecommunications, health-care, environment.*

The following application domains are concerned by our work: telecommunication networks, medicine and environment.

4.2. Telecommunication networks

Key words: *telecommunications.*

Monitoring telecommunication networks is an important task and is one of the conditions to ensure a good quality of service. Given a monitoring system continuously receiving observations (alarms) sent by the system components, our purpose is to help operators to identify failures. Two classical approaches in monitoring such systems are knowledge-based techniques which directly associate a diagnosis to a set of symptoms as expert systems or chronicle recognition systems, and model-based techniques which rely on a behavioral model of the system. The main weakness of the first approach is the lack of genericity: as the system changes, a new expertise has to be acquired. It is why we focus on model-based techniques which are known to be better suited to evolving systems as telecommunication networks are. Centralized diagnostic approaches for diagnosing discrete-event systems have been developed [53][22][50]. The main drawback of centralized approaches is that they require to explicitly build the global model of the system which is unrealistic for large and complex systems as telecommunication networks.

In this context, we developed a decentralized component-oriented approach, able to incrementally compute on-line diagnoses [43]. The efficiency of the algorithm is increased by the use of model-checking techniques as partial order reduction techniques and BDD. Currently, we extended our research to reconfigurable systems, i.e systems the topology of which is changing along time, due for instance to reconfiguration actions decided to remedy upload problems.

Another important challenge for telecommunication networks is predicting the subjective quality of proposed services (as it could be felt by the user) from collected technical data. Mixing data-mining and symbolic learning techniques is the way we chose to acquire this predictive knowledge. All this research work on telecommunication networks has been done in collaboration and with the support of France-Telecom R&D.

4.3. Decision aiding in medicine and health-care

Key words: *medicine, health-care.*

Since the development of expert systems in the 70's, decision aiding tools have been widely studied and used in medicine and health-care. The ultimate goal is to help a physician to establish his diagnosis, or prognosis from observations delivered by sensors and the individual patient's data. This involves at least three tasks:

- patient monitoring: processing and abstracting signals recorded by sensors placed on patients, in order to generate alarms when a particular situation has occurred, or is about to occur. The standard context is intensive care units in hospitals where an alarm must be treated within a very short time. With the advent of telemedicine similar situations arise, but the delay to treat an alarm may be much longer. For example, a cardiac or diabetic patient may be surveyed at home and the recorded data are sent every day at some fixed hour to the care unit. If some problem is detected, the patient is urged to consult a doctor, but a long delay may occur between the time at which the problem occurred and the treatment. Time is a major feature of medical data, thus temporal abstraction associated to signal processing techniques must be used for filtering and pre-processing the raw data;
- diagnostic and prognostic reasoning: models, such as causal or probabilistic models, have supplanted expert systems for diagnosis. As the course and outcome of a disease process is dynamic, time plays also an important role in diagnostic and prognostic models. Also, treatment planning or/and the clinical context may interact with these two basic reasoning processes and particular methods have to be studied and implemented to integrate these aspects;
- modeling: though, some particular parts of the human body are known very well, e.g. the heart, deep models are generally difficult to build in medicine because knowledge is incomplete or too complex, e.g. the brain. Fortunately, huge amounts of data have been recorded and stored in medical databases. These data can be analyzed in order to discover new knowledge that may be used to construct abstract models or behavioral models, very similar to the old expert systems, but avoiding the bottleneck of expert knowledge acquisition. Processing medical data is a specific research area known as “intelligent data analysis (IDA) in medicine” [39]. An essential feature of the techniques used in IDA is that most are knowledge-based: they can use knowledge about the problem domain. Thus, a learning approach such as inductive logic programming is a tool of choice.

These three points are studied in projects involving industrial (ELA medical), medical (University Hospital of Rennes) and academic (LTSI - University of Rennes) partners, especially in the field of cardiology. Particularly, new cardiac devices and monitoring systems are investigated.

4.4. Environmental decision making

Key words: *environment.*

The need of decision support systems in the environmental domain is now well-recognized. It is especially true in the domain of water quality and a program, named Bretagne Eau Pure (<http://www.bretagne-eau-pure.org>), was launched some years ago in order to help regional managers to protect this important resource. The challenge is to preserve the water quality from pollutants as nitrates and herbicides, when these pollutants are massively used by farmers to weed their agricultural plots and improve quality and quantity of their crops. The difficulty is then to find solutions which satisfy contradictory interests and first to get a better knowledge on pollutant transfer. For instance, it is certainly true that the pesticide transfer through catchments is still not enough analyzed and poorly known.

In this context, we are developing decision support systems to help regional managers in preserving the river water quality. Two main artificial intelligence techniques are used in this area: multi-agents systems, which are suited to model multi-expert cooperation and qualitative modeling to model biophysical processes in an explicative and understandable way. The approach we advocate is the coupling of a qualitative biophysical model, able to simulate the biophysical process, and a management model able to simulate the farmer decisions.

Two main research themes are investigated in this framework: the use of qualitative spatial modeling to simulate the pollutant transfer through agricultural catchments and the use of learning/data mining techniques to discover, from model simulation results, the discriminant variables and acquire rules relating these variables. In both cases, one of the main challenges is that we are faced with spatio-temporal data.

Our partners are mainly the SAS Inra research group, located in Rennes and other Inra research groups as the BIA group in Toulouse and the LASB group in Montpellier.

6. New Results

6.1. Diagnosis of large scale discrete event systems

Participants: Marie-Odile Cordier, Alban Grastien, Christine Largouët, Laurence Rozé.

The problem we deal with is the monitoring of complex and large discrete-event systems (DES) such as telecommunication networks. Diagnosing dynamical systems represented as DES consists in finding what happened to the system from existing observations. Different terminologies can be found in the literature as histories, scenarios, narratives, consistent paths. They all rely on the idea that the diagnostic task consists in determining the trajectories (a sequence of states and events) compatible with the sequence of observations. From these trajectories, it is then easy to determine (identify and localize) the possible faults. The main difficulty is the intractable size of the model and the huge number of states and trajectories to be explored. To cope with this problem, two approaches are investigated: the use of a decentralized diagnosis approach and the use of model-checking techniques. Both were experimented in the context of the MAGDA project.

6.1.1. A decentralized diagnosis approach

The decentralized approach enables an on-line diagnosis without requiring the computation of the global model. Given a decentralized model of the system and a flow of observations, the program computes the diagnosis by combining local diagnoses built from local models (or local diagnosers). Two main problems have been investigated in the last three years ([46][45]): which strategy for an optimal merge of the local diagnoses in order to preserve the efficiency and the completeness of the process? how to ensure the incrementality of the process in an on-line diagnosis context where observations are incrementally collected? A paper describing the formal framework of this work and its experimentation on telecommunication networks has been written in 2003 and is currently under review process (submitted to Artificial Intelligence Journal).

We are currently working on the on-line diagnosis of discrete event *reconfigurable* systems, in order to make the diagnosis process able to deal with systems whose topology can be changed at running time. This is for instance the case of telecommunication networks whose topology can be changed to solve overload problems, and whose components can even be replaced by new ones. This is the subject of Alban Grastien's thesis. A first step in this direction has been to specify the formalism that will be used to describe the dynamical topology of a network, and to design diagnosis algorithms able to cope with these topological changes.

6.1.2. The use of model-checking techniques

Another way to improve the efficiency of the computation and to cut down the combinatorial state explosion is to use model-checking techniques, initially developed for automatic verification of complex real-time systems. It means using adapted symbolic representations such as BDD, for instance, or partial order reduction techniques taking advantage of the existing concurrency. Partial order reduction techniques have been used to represent in a compact way a set of trajectories in the decentralized approach described before. BDD techniques have also been proved to improve the efficiency of the process when used to code a diagnoser (see [41]).

The work currently done in this direction has been presented in [12] and [11]. A property, called invertibility, has been defined on the events (or actions) of the system model and is used to prune the search space without losing any information. Intuitively, two events are said to be invertible when, after transiting through any sequence of events including these two events, the state reached by the system is independent of the ordering of the two events. This property provides an efficient way of representing sets of trajectories and is exploited to restrict the set of behaviors to be considered in diagnosis or planning systems (both planning and diagnostic problems can be viewed as finding a path over a behavioral model). Two algorithms are proposed: the first one efficiently computes trajectories and the second one automatically computes the invertibility properties from

the automata. The invertibility property is related to the independence property between events used in partial order reduction and to the related notion of traces. The next step consists in updating the existing decentralized diagnosis algorithms to take advantage of this property.

6.2. Cardiac monitoring

Participants: Marie-Odile Cordier, Élisabeth Fromont, François Portet, René Quiniou, Sophie Robin.

Together with the LTSI (Signal Processing Lab - INSERM, University of Rennes 1), we are studying how to use chronicle recognition techniques for cardiac monitoring and diagnosis. Our goal is to analyze the signals coming from several sensors in order to detect and characterize the cardiac arrhythmias of a monitored patient. The nature of the arrhythmias, their features and their frequency can be used to propose convenient therapies, such as specific drugs or cardiac devices (pacemaker or defibrillator).

We are particularly working on two aspects: discovering chronicles by machine learning and improving event detection on signals. Concerning chronicle discovery we are studying how to adapt machine learning techniques in order to deal with multichannel aspects. Various, control policies are implemented and assessed: global learning and recognition from information provided by all the channels, independent learning on each channel and then symbolic knowledge fusion for global recognition, or independent recognition on each channel and merging all the results. We are particularly working on symbolic knowledge fusion which appears to be a hot topic in the knowledge acquisition community. This is the subject of Élisabeth Fromont's thesis which is supported by a grant from the RNTS Cepica project (cf. 7.3).

In order to improve the quality of signal processing and event detection, we are studying techniques that could implement a tight collaboration between low level signal processing algorithms on the one hand, and high level recognition algorithms such as chronicle recognition on the other hand. Precisely, partially recognized chronicles can predict which events are expected to be observed by sensors. These predictions can be used to choose the relevant algorithms for processing the signals according to the context (noise, predicted arrhythmias, cardiac devices - pacemakers or monitors -, etc.). This is the subject of François Portet's thesis. To begin with, he has completed an extended analysis of signal processing algorithms used for ECG analysis in order to assess and to model their performance in different application contexts (noise type and level, shape of ECG beats, ECG channel, etc.). This knowledge will be used in rules for guiding the signal processing module along the recognition context.

This year we have also worked on a comparison between Calicot, the system we have proposed for cardiac arrhythmia monitoring and diagnosis, and Kardio [23], a system implemented in the late 80s for learning cardiac arrhythmias classification and prediction rules. Kardio shares many goals with Calicot, this is why we had to achieve an insightful comparison. This study showed that important differences exist between the two systems, especially knowledge representation aspects. Consequently, a careful methodology to compare the two systems had to be devised. For example, Kardio has no signal processing capabilities. ECG descriptions are provided by experts. The authors chose to use high level attributes to describe ECGs, even though they were, and are still, beyond the actual capabilities of signal processing tools. Calicot relies on signal processing for providing information that will be used by the chronicle learner and recognizer. So, the used attributes are much lower level attributes. Two different comparison scenarios were performed: a direct comparison (which necessitated to describe the test ECGs in Kardio language) and a comparison with a relaxed version of Kardio, in which a part of Kardio knowledge was recoded with the representation attributes used by Calicot. The first experiment was in favor of Kardio but the second showed that the two systems performed equally well [13].

6.3. Learning diagnostic rules for telecommunication networks

Participants: Marie-Odile Cordier, René Quiniou.

Magda2 (Modélisation et Apprentissage pour une Gestion Distribuée des Alarmes De bout En boUt - Modeling and Machine Learning for Distributive Management of Alarms from End to End) is an RNRT project which aims at providing advanced solutions for managing heterogeneous networks, and for taking into account events

involving an interaction between the network and the service layers. We have been working on learning to predict the quality of service (QoS) of an application (TV on demand) running on a telecommunication network from information provided by equipments (routers) of this network. Precisely we have been working on time series abstraction by time-stamped symbolic event sequences. These abstracted time series are then processed by an inductive logic programming technique in order to learn diagnostic temporal rules predicting the QoS at the end user side. The learning data were recorded from a monitored network on which different faults were actually implemented. Though the results are not as rich as expected, because the recorded data did not show a sufficient quality, we demonstrated the feasibility of our approach to temporal learning and data mining [17].

6.4. Diagnosis of a synchrotron X-ray beamline

Participants: Marie-Odile Cordier, René Quiniou, Sophie Robin.

The "Soleil" project aims to build a 3rd generation synchrotron-radiation center close to the Orsay University Campus. We have been contacted by Soleil project researchers working in biocrystallography (which aims to identify the 3-D structure of the proteins). The objective is to supervise the operation of beamline PROXIMA 1 (under construction) so that fault conditions are diagnosed and, if necessary, corrected by automatic realignment procedures. In the first instance, we will study how far the tools developed in the scope of diagnosis can be suitable for such a monitoring and automatic control system for this beam line.

We are in the preliminary step of understanding and specifying the research topic. Regular meetings (6) were organized during last year.

6.5. Learning decision rules in the water resource management domain

Participants: Marie-Odile Cordier, Véronique Masson.

In the framework of the Sacadeau project, our aim is to build decision support systems to help catchment managers to preserve streamwater quality.

In collaboration with INRA researchers, two actions are conducted in parallel.

- The first one consists in building a qualitative model to simulate the pesticide transfer through the catchment from the time of its application by the farmers to the arrival at the stream. The model architecture relies on the coupling of two models: a biophysical transfer model and a management model able to simulate the farmer decisions in herbicide application according to the climate and the weeding strategy to cite only some of the decision criteria. Given data on the year climate, on the catchment topology and on farmer strategy, the model outputs the pesticide concentration in the stream along the year. This work is currently under development. Even if mainly in Inra hands, we actively participate to its realization.
- The second action consists in identifying some of the input variables as main pollution factors and in learning rules relating these pollution factors to the temporal distribution of the stream pesticide concentration. We chose to use Inductive Logic Programming (ILP) techniques to get easy-to-read and explicative rules. ILP algorithms build hypothetical rules generalizing ("explaining") the examples which are given to them as input. A first but important step is then to constitute the learning data base, which has to be representative of all the considered situations. This year, we analyzed which "scenarios" were representative of our application and then used the transfer model to collect significant sets of simulated data.

We are now ready to exploit this learning data base and use ILP tools to infer rules relating pollution factors and temporal distribution of the pollution.

6.6. Action languages and default inferences

Participants: Yves Moinard, Philippe Besnard.

We have continued our work on the inference by plausibility of Friedman and Halpern [36]. Default reasoning allows to make tentative conclusions which are the best conclusions that can be drawn from the present state of knowledge. This concerns rules with exceptions, which are implicit in various domains. For instance, in order to compute the result of an action, many hidden hypothesis should be made. The best conclusion is to consider that, except for the exceptional cases which can be deduced from the given knowledge, everything is normal, since, without this assumption, either no conclusion could be drawn, or no real situation could be formalized. The inference by plausibility is an interesting recent proposal which can formalize default reasoning. In particular it is well fitted for applying the methods of knowledge compilation: a plausibility is assigned to each formula, a task which can be made "off-line", then the deductions consist only in comparing these plausibilities.

As for the original proposal of inference by plausibility, which was a generalization of earlier proposals, it was not really appropriate without a slight modification. In default reasoning, it is often necessary to consider cases where the reasoning rule "(AND)" is falsified. This means that we can deduce A , and also B , but not necessarily their conjunction $A \wedge B$. This allows e.g. to translate situations involving actions with indeterministic effects. The ability for rejecting the (AND) rule was one of the features of the inference by plausibility, with respect to the earlier formalisms from which it has been designed. However, the proposed inference by plausibility satisfies a restricted kind of (AND) (when A and B are contradictory), which is a restriction so severe that in fact, in many cases, the (AND) rule is enforced. It happens that a small modification of the definition solves this problem. We have designed a new definition, and listed the main reasoning properties of the two versions. Knowing these properties, which all have an intuitive translation, is necessary for a potential user of such a formalism. Also, this study has shown that all the cases considered by the authors of the first formulation can be translated into the new one, while the converse is far from being true (see a preliminary publication in [14]).

Working directly on a language for describing actions, we are extending a recent proposal [40] in order to apply it to real situations involving time such as watching and monitoring of networks. The proposal, oriented for data mining, introduces a representational framework of sequences, SeqLog, which is a kind of DataLog where sequences take the role of deductive databases. The aim is to be able to use methods of logic programming or other methods of computational logic, in order to do serious data mining. Thus, the translation of a problem should be easy for a user, since the language used is close to formal logic, and efficient methods of computation could be used on the "machine side". We consider sequences of actions. In many situations, delays or other notions appear, which enforces the introduction of time. We are studying the introduction of an easy "time operator", which should suffice for describing various useful situations, without complicating the computation too much. We introduce, as the basic sequence concatenation operator, an interval of time: $e_i[s, f]e_{i+1}$ means e_{i+1} is the next event after e_i , and that it takes effect after a delay which is in $[s, f]$. The notion of subsumption, which is the key point of the method designed by De Raedt and Lee, is extended to these time connectors. The aim is to allow the description of real situations, and to allow a detection of rather complex patterns. Then, further work should be made on two sides: checking the efficiency of the method from a computational side, and also checking the appropriateness of the method by applying it to "real situations".

7. Contracts and Grants with Industry

7.1. Magda2: Modelling, diagnosing and supervising telecommunication networks

Participants: Marie-Odile Cordier, René Quiniou.

The RNRT Magda2 project (Modélisation et Apprentissage pour une Gestion Distribuée des Alarmes De bout En boUt - Modelling and Machine Learning for Distributive Management of Alarms from End to End) began in November 2001 and ended in December 2003. The project involves the following partners: Alcatel, France-Telecom R & D, ILOG, LIPN-University of Paris-Nord and IRISA (Sigma2, Triskell and Dream). The goal of the project is to discover advanced solutions for managing heterogenous networks and for taking into account events involving an interaction between the network and the service layers. DREAM is particularly concerned with the study and the implementation of robust algorithms for correlating network indicators and the End to End Quality of service and for diagnosing telecommunication networks. Symbolic machine learning tools were used in order to acquire the necessary knowledge and models.

7.2. Sacadeau: Decision-aid to improve streamwater quality

Participants: Marie-Odile Cordier, Véronique Masson.

The project Sacadeau (Système d'Acquisition de Connaissances pour l'Aide à la Décision pour la qualité de l'EAU - Knowledge Acquisition System for Decision-Aid to Improve Streamwater Quality) has begun in October 2002. It is funded by Inra (French institute for agronomy research) and should last three years. The project involves the following partners: three INRA research groups (SAS from Rennes, LASB from Montpellier and BIA from Toulouse) and Irisa. It also involves experts belonging to the regional administrative entities. The project aims at building a decision-aid tool to help specialists in charge of the catchment management in order to preserve the streamwater quality. The proposal relies on the building of two coupled qualitative models: a transfer model to simulate the pesticide transfer through the catchment and a management model to simulate the farmer decisions concerning the application of pesticides and the weeding strategy. The final objective is to analyze simulation results by using learning and data mining techniques, to discover the discriminant variables and to acquire rules relating the climate, the farmer strategy, the catchment topology with the pesticide concentration in the stream.

7.3. Cepica: Conception and Evaluation of an Implantable Cardiac Device

Participants: Marie-Odile Cordier, Élisabeth Fromont, François Portet, René Quiniou.

This RNTS (Réseau National Technologies pour la Santé) project has begun at the end of 2003 and will last 3 years. The partners are ELA-Medical, the department of cardiology of the Rennes University Hospital, the LTSI-University of Rennes 1 and IRISA. The project is concerned with the conception of new cardiac devices, the study of which has begun during the instigative concerted action PISE. Its main concerns are: to propose and to evaluate new sensors able to assess the hemodynamic effects of a stimulation; to develop signal processing methods devoted to the specific signals measured by the new sensors and to refine, by using machine learning methods and chronicle recognition, the scenarios that may present some risk for an individual patient; to study different stimulation protocols taking into account the device specificities and constraints; to validate these concepts in clinical situations.

8. Other Grants and Activities

8.1. National projects

Members of the Dream team are involved in the following national collaboration programs:

- IMALAIA (common working group of the GdR Automatique, GDR- PRC I3 and Afia group) which brings together researchers from automatic and artificial intelligence fields on the subject of dynamic system monitoring. M.-O. Cordier is co-chair with L. Travé-Massuyès and F. Lévy.
- RTP "information and intelligence: reasoning and decision" set up by the department STIC of the CNRS (M.-O. Cordier is a member of the steering committee).
- Café: AFIA working group (machine-learning, knowledge discovery in databases, data mining - R. Quiniou).

8.2. International networks and working groups

- Monet2 (European network of excellence on Model-based and Qualitative reasoning). DREAM is particularly involved in the Bridge task group, which attempts to integrate AI and automatic methods for diagnosis and monitoring, and the biomedical task group which attempts to forge links between the fields of Biological Research and medical qualitative reasoning research (M.-O. Cordier and R. Quiniou).

8.3. International bilateral relations

- PROCOPE project no 99027 « Foundations for the treatment of contradictions intelligent information systems » between the university of Potsdam and IRISA (Ph. Besnard, M.-O. Cordier)

9. Dissemination

9.1. Animation of the Scientific Community

9.1.1. Journal editorial board

- *AAI: Applied Artificial Intelligence* (M.-O. Cordier).
- *AICOMs: Artificial Intelligence Communications* (M.-O. Cordier).
- *JEDAI: Journal Electronique d'Intelligence Artificielle* (M.-O. Cordier).
- *Revue I3* (M.-O. Cordier).
- *Sciences et Techniques Éducatives* (D. Py).

9.1.2. Conference program committees

- DX'03, RFIA'04, IJCAI'03, KR'04, DX'04 (M.-O. Cordier).
- EIAH'2003 (D. Py).
- AIME'03 Workshop: Model-based and Qualitative Reasoning in Biomedicine (R. Quiniou).

9.2. Faculty teaching

Many members of the DREAM team are also faculty members and are actively involved in computer science teaching programs in Ifsic, INSA and ENSAR. Besides these usual teachings Dream is involved in the following programs:

- DEA of computer science (IFSIC): *RATS module: temporal and spatial reasoning* (M.-O. Cordier, Y. Moinard, R. Quiniou).
- DEA of computer science (IFSIC): *DIAG module: diagnosis* (M.-O. Cordier, L. Rozé, S. Robin).

10. Bibliography

Major publications by the team in recent years

- [1] G. CARRAULT, M.-O. CORDIER, R. QUINIOU, M. GARREAU, J.-J. BELLANGER, A. BARDOU. *A model-based approach for learning to identify cardiac arrhythmias*. in « AIMDM'99 : Artificial Intelligence in Medicine and Medical Decision Making », series Lecture Notes in Artificial Intelligence, volume 1620, Springer Verlag, W. HORN, Y. SHAHAR, G. LINDBERG, S. ANDREASSEN, J. WYATT, editors, pages 165-174, Aalborg, Denmark, 1999.
- [2] C. LARGOUËT, M.-O. CORDIER. *Improving the Landcover Classification using Domain Knowledge*. in « AI Communication special issue on Environmental Sciences and Artificial Intelligence », number 1, volume 14, 2001, pages 35-43.
- [3] Y. MOINARD. *Note about cardinality-based circumscription*. in « Artificial Intelligence », number 1-2, volume 119, May, 2000, pages 259-273, <http://www.elsevier.nl:80/inca/publications/store/5/0/5/6/0/1/>.
- [4] Y. MOINARD, R. ROLLAND. *Characterizations of Preferential Entailments*. in « Logic Journal of the IGPL », number 3, volume 10, 2002, pages 243-270, <http://www3.oup.co.uk/igpl/>.
- [5] Y. PENCOLÉ, M.-O. CORDIER, L. ROZÉ. *A decentralized model-based diagnostic tool for complex systems..* in « International Journal on Artificial Intelligence Tools », number 3, volume 11, 2002, pages 327-346.
- [6] L. ROZÉ, M.-O. CORDIER. *Diagnosing discrete-event systems : extending the “diagnoser approach” to deal with telecommunication networks*. in « Journal on Discrete-Event Dynamic Systems : Theory and Applications (JDEDS) », number 1, volume 12, 2002, pages 43-81, <http://www.inria.fr/rrrt/rr-4957.html>.
- [7] S. THIÉBAUX, M.-O. CORDIER, O. JEHL, J.-P. KRIVINE. *Supply Restoration in Power Distribution Systems - A Case Study in Integrating Model-Based Diagnosis and Repair Planning*. in « Actes de UAI-96 », pages 525-532, 1996.

Articles in referred journals and book chapters

- [8] G. CARRAULT, M.-O. CORDIER, R. QUINIOU, F. WANG. *Temporal abstraction and Inductive Logic Programming for arrhythmia recognition from electrocardiograms*. in « Artificial Intelligence in Medicine », volume 28, 2003, pages 231-263.
- [9] M.-O. CORDIER, P. DAGUE, M. DUMAS, F. LÉVY, J. MONTMAIN, M. STAROSWIECKI. *Conflicts versus Analytical Redundancy Relations*. in « IEEE Transactions on Systems, Man and Cybernetics - Part B. », 2003, to appear.
- [10] F. WANG, G. CARRAULT, R. QUINIOU, M.-O. CORDIER, P. MABO. *Fusion de méthodes d'apprentissage automatique et de traitement du signal pour la reconnaissance d'arythmies*. in « ITBM-RBM Innovation et technologie en biologie et médecine », 2003, to appear.

Publications in Conferences and Workshops

- [11] M.-O. CORDIER, A. GRASTIEN, C. LARGOUËT, Y. PENCOLÉ. *Calcul de trajectoires utilisant les propriétés d'interversibilité*. in « RJCIA'03 (Rencontres jeunes chercheurs en Intelligence Artificielle) », pages 15–28, Laval, 2003.
- [12] M.-O. CORDIER, A. GRASTIEN, C. LARGOUËT, Y. PENCOLÉ. *Efficient trajectories computing exploiting inversibility properties*. in « DX'03 (International workshop on principles of Diagnosis) », pages 93–98, Washington, 2003.
- [13] É. FROMONT, M.-O. CORDIER, R. QUINIOU, A. HERNANDEZ. *Kardio and Calicot: a comparison of two cardiac arrhythmia classifiers*. in « AIME'03 Workshop: Qualitative and Model-based Reasoning in Biomedicine », Protaras, Cyprus, 2003.
- [14] Y. MOINARD. *De l'usage de la mesure de plausibilité pour le raisonnement par défaut*. in « JNMR », pages 185 - 200, Paris, 2003, <http://www.cril.univ-artois.fr/~marquis/articlesJNMR-03.html>.
- [15] F. WANG, G. CARRAULT, R. QUINIOU, M.-O. CORDIER, L. LUO. *Temporal reasoning based automatic arrhythmias recognition*. in « IEEE International Conference on RISSP 2003 : Robotics, Intelligent Systems and Signal Processing », Changsha, China, 2003.
- [16] F. WANG, G. CARRAULT, R. QUINIOU, M.-O. CORDIER. *Reconnaissance d'événements structurés temporellement dans un signal par raisonnement temporel*. in « 19e colloque GRETSI'03 sur le traitement du signal et des images », Paris, 2003.

Internal Reports

- [17] R. QUINIOU, M.-O. CORDIER. *Apprentissage de règles de prédiction de la qualité de service d'un réseau de télécommunication*. Technical report, IRISA, 2003.

Bibliography in notes

- [18] A. AGHASARYAN, E. FABRE, A. BENVENISTE, R. BOUBOUR, C. JARD. *Fault detection and diagnosis in distributed systems : an approach by partially stochastic Petri nets*. in « Discrete Event Dynamic Systems », number 2, volume 8, Juin, 1998, pages 203-231, Special issue on Hybrid Systems.
- [19] P. BARONI, G. LAMPERTI, P. POGLIANO, M. ZANELLA. *Diagnosis of a Class of Distributed Discrete-Event Systems*. in « IEEE Transactions on systems, man, and cybernetics », number 6, volume 30, Novembre, 2000, pages 731-752.
- [20] P. BARONI, G. LAMPERTI, P. POGLIANO, M. ZANELLA. *Diagnosis of large active systems*. in « Artificial Intelligence », volume 110, 1999, pages 135-183.
- [21] M. BASSEVILLE, M.-O. CORDIER. *Surveillance et diagnostic de systèmes dynamiques : approches complémentaires du traitement de signal et de l'intelligence artificielle*. Technical report, number 1004, Irisa, 1996.
- [22] S. BIBAS, M.-O. CORDIER, P. DAGUE, F. LEVY, L. ROZE. *GASPAR: a model-based system for diagno-*

sing telecommunication networks. in « IMACS-IEEE/SMC International Multiconference of Computational Engineering in Systems Applications (CESA'96) », Lille, France, 1996.

- [23] I. BRATKO, I. MOZETIC, N. LAVRAČ. *Kardicid dro: A Study in Deep and Qualitative Knowledge for Expert Systems.* MIT Press, 1989.
- [24] A. CORNUÉJOLS, L. MICLET. *Apprentissage artificiel : concepts et algorithmes.* Eyrolles, 2002.
- [25] R. DEBOUK, S. LAFORTUNE, D. TENEKETZIS. *Coordinated Decentralized Protocols for Failure Diagnosis of Discrete Event Systems.* in « Discrete Event Dynamic Systems », number 1-2, volume 10, 2000, pages 33-86.
- [26] R. DEBOUK, S. LAFORTUNE, D. TENEKETZIS. *A coordinated decentralized protocol for failure diagnosis of discrete event systems.* in « Proceedings of the Workshop on Discrete Event Systems (WODES'98) », pages 138-143, Cagliari, Italy, 1998.
- [27] J. DE KLEER, A. MACKWORTH, R. REITER. *Characterizing diagnoses and systems.* in « Artificial Intelligence », number 2-3, volume 56, 1992 (repris dans *Readings in Model-Based Diagnosis*, W. Hamscher, L. Console, J. de Kleer (dir.), Morgan Kaufmann, p. 54-65, 1992), pages 197-222.
- [28] J. DE KLEER, B. C. WILLIAMS. *Diagnosis with behavioral modes.* in « Proceedings of the 11th International Joint Conference on Artificial Intelligence IJCAI'89 », pages 1324-1330, Detroit, MI, Etats-Unis, 1989 (repris dans *Readings in Model-Based Diagnosis*, W. Hamscher, L. Console, J. de Kleer (dir.), Morgan Kaufmann, p. 124-130, 1992).
- [29] C. DOUSSON. *Extending and unifying chronicle representation with event counters.* in « Proc. of ECAI 2002 », IOS Press, F. VAN HARMELEN, editor, 2002.
- [30] C. DOUSSON. *Suivi d'évolutions et reconnaissance de chroniques.* Ph. D. Thesis, Université Paul Sabatier de Toulouse, LAAS-CNRS, Toulouse, 1994.
- [31] C. DOUSSON, P. GABORIT, M. GHALLAB. *Situation recognition: representation and algorithms.* in « Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) », pages 166-172, Chambéry, France, 1993.
- [32] B. DUBUISSON, editor, *Diagnostic, intelligence artificielle et reconnaissance des formes.* series *Traité IC2 : Information - Commande - Communication*, Hermes, 2001.
- [33] *Relational Data Mining.* S. DZEROSKI, N. LAVRAČ, editors, Springer, Berlin, 2001.
- [34] S. DZEROSKI, L. TODOROVSKI. *Discovering dynamics: from inductive logic programming to machine discovery.* in « Journal of Intelligent Information Systems », volume 4, 1995, pages 89-108.
- [35] E. FABRE, A. BENVENISTE, C. JARD, L. RICKER, M. SMITH. *Distributed State Reconstruction for Discrete Event Systems..* in « Proc. of the 2000 IEEE Control and Decision Conference (CDC) », Sydney, Australie, 2000.

- [36] N. FRIEDMAN, J. Y. HALPERN. *Plausibility measures and default reasoning*. in « Journal of the ACM », number 4, volume 48, July, 2001, pages 648-685.
- [37] W. HAMSCHER, L. CONSOLE, J. DE KLEER, editors, *Readings in Model-Based Diagnosis*. Morgan Kaufmann, San Mateo, CA, Etats-Unis, 1992.
- [38] D. T. HAU, E. W. COIERA. *Learning qualitative models of dynamic systems*. in « Machine Learning », volume 26, 1997, pages 177-211.
- [39] N. LAVRAC, E. KERAVALOU, B. ZUPAN. *Intelligent data analysis in medicine*. in « Encyclopedia of computer science and technology », number 27, volume 42, 2000, pages 113-157, Marcel Dekker, Inc., New York.
- [40] S. D. LEE, L. D. RAEDT. *Constraint Based Mining of First Order Sequences in SeqLog*. in « Workshop on Multi-Relational Data Mining, in conjunction with 8th ACM SIG KDD conf. », Edmonton, Alberta, 2002.
- [41] H. MARCHAND, L. ROZÉ. *Diagnostic de pannes sur des systèmes à événements discrets : une approche à base de modèles symboliques*. in « 13^{ème} Congrès Francophone AFRIF-AFIA de Reconnaissance de Formes et d'Intelligence Artificielle (RFIA02) », pages 191-200, Angers, France, 2002.
- [42] D. PAGE. *ILP: Just Do It*. in « Proceedings of ILP'2000 », series LNAI, volume 1866, Springer, J. CUSSENS, A. FRISCH, editors, pages 3-18, 2000.
- [43] Y. PENCOLÉ, M.-O. CORDIER, L. ROZÉ. *A decentralized model-based diagnostic tool for complex systems*. in « ICTAI'01 (IEEE international conference on tools with artificial intelligence) », IEEE computer society, pages 95-102, 2001.
- [44] Y. PENCOLÉ, M.-O. CORDIER, L. ROZÉ. *Incremental decentralized diagnosis approach for the supervision of a telecommunication network*. in « DX'01 (International Workshop on Principles of Diagnosis) », Sansicario, Italy, 2001, <http://www.irisa.fr/dream/dataFiles/ypencole/DX01.pdf>.
- [45] Y. PENCOLÉ, M.-O. CORDIER, L. ROZÉ. *Incremental decentralized diagnosis approach for the supervision of a telecommunication network*. in « CDC'02 (IEEE Conference on Decision and Control) », pages 435-440, Las Vegas, Nevada, USA, 2002.
- [46] Y. PENCOLÉ. *Diagnostic décentralisé de systèmes à événements discrets : application aux réseaux de télécommunications*. Ph. D. Thesis, Université de Rennes 1, 2002, <http://www.irisa.fr/aida/aida-new/dataFiles/ypencole/these.pdf>.
- [47] R. QUINIOU, M.-O. CORDIER, G. CARRAULT, F. WANG. *Application of ILP to cardiac arrhythmia characterization for chronicle recognition*. in « ILP'2001 », series LNAI, volume 2157, Springer-Verlag, C. ROUVEIROL, M. SEBAG, editors, pages 220-227, 2001, <http://www.irisa.fr/dream/dataFiles/quiniou/ilp01.pdf>.
- [48] N. RAMAUX, M. DOJAT, D. FONTAINE. *Temporal scenario recognition for intelligent patient monitoring*. in « Proc. of the 6th Conference on Artificial Intelligence in Medecine Europe (AIME'97) », 1997.
- [49] R. REITER. *A theory of diagnosis from first principles*. in « Artificial Intelligence », number 1, volume 32, 1987

(repris dans *Readings in Nonmonotonic Reasoning*, M. L. Ginsberg (dir.), Morgan Kaufmann, 1987 ; aussi dans *Readings in Model-Based Diagnosis*, W. Hamscher, L. Console, J. de Kleer (dir.), Morgan Kaufmann, p. 29-48, 1992), pages 57-96.

- [50] L. ROZÉ, M.-O. CORDIER. *Diagnosing discrete-event systems : extending the “diagnoser approach” to deal with telecommunication networks*. in « Journal on Discrete-Event Dynamic Systems : Theory and Applications (JDEDS) », number 1, volume 12, 2002, pages 43-81, <http://www.irisa.fr/aida/aida-new/dataFiles/cordier/jeds.pdf>.
- [51] L. ROZÉ, M.-O. CORDIER. *Diagnosing Discrete Event Systems : An experiment in Telecommunication Networks*. in « WODES98, Fourth Workshop on Discrete Event Systems », pages 130-137, Cagliari, Italy, 1998.
- [52] L. ROZÉ. *Supervision de réseaux de télécommunication : une approche à base de modèles*. Ph. D. Thesis, Université de Rennes 1, 1997.
- [53] M. SAMPATH, R. SENGUPTA, S. LAFORTUNE, K. SINNAMOHIDEEN, D. TENEKETZIS. *A discrete event systems approach to failure diagnosis*. in « Proceedings of the Fifth international workshop on principles of diagnosis (DX'94) », pages 269-277, 1994.
- [54] M. SAMPATH, R. SENGUPTA, S. LAFORTUNE, K. SINNAMOHIDEEN, D. TENEKETZIS. *Active Diagnosis of Discrete-Event Systems*. in « IEEE Transactions on Automatic Control », number 7, volume 43, 1998, pages 908-929.
- [55] M. SEBAG, C. ROUVEIROL. *Constraint Inductive Logic Programming*. L. DE RAEDT, editor, in « Advances in Inductive Logic Programming », IOS Press, 1996, pages 277-294, <http://citeseer.nj.nec.com/sebag96constraint.html>.
- [56] GROUPE ALARME. *Monitoring and alarm interpretation in industrial environments*. in « AI Communications », volume 11, 3-4, 1998, pages 139-173, S. Cauvin, M.-O. Cordier, C. Dousson, P. Laborie, F. Lévy, J. Montmain, M. Porcheron, I. Servet, L. Travé-Massuyès.