

Chronicles for On-line Diagnosis of Distributed Systems

Xavier Le Guillou*

Irisa – Université de Rennes 1
Campus de Beaulieu
35042 Rennes Cédex
xleguill@irisa.fr

Abstract. The formalism of chronicles has been proposed to monitor and diagnose dynamic physical systems. Even if efficient chronicle recognition algorithms exist, it is now well-known that distributed approaches are better suited to monitor actual systems. In this article, we adapt the chronicle-based approach to a distributed context and illustrate this work on the monitoring of software components.

Key words: on-line diagnosis, distributed systems, chronicle recognition

1 Introduction

Monitoring and diagnosing dynamic systems have become very active topics in research and development for a few years. Besides continuous models based on differential equations, essentially used in control theory and discrete event systems based on finite state machines (automata, Petri nets, . . .), a formalism commonly used for on-line monitoring, in particular by people from the artificial intelligence community, is the one of chronicles. This formalism, proposed in [1], has been widely used and extended [2–4]. A chronicle describes a situation that is worth identifying within the diagnosis context. It is made up a set of events and temporal constraints between those events. As a consequence, this formalism fits particularly well problems that consider a temporal dimension. The set of interesting chronicles constitutes the base of chronicles. Then, monitoring the system consists in analyzing flows of events, and recognizing on fly patterns described by the base of chronicles. Efficient algorithms exist which explain that this approach has been used for industrial applications as well as medical ones [2, 5, 6].

One of the key issues of model-based approaches for on-line monitoring is the size of the model which is generally too large when dealing with real applications. Distributed or decentralized approaches have been proposed to cope with this problem, like [7–10]. The idea is to consider the system as a set of interacting components instead of a unique entity. The behavior of the system is thus described by a set of local component models and by the synchronization constraints between the component models.

Considering chronicle-based approaches, to our knowledge, no distributed approaches exist and the contribution of this paper consists in adapting the chronicle-based approach to distributed systems.

* under the direction of M.-O. Cordier

This work has been motivated by an application that aims at monitoring the behavior of software components, and more precisely of web services within the context of the WS-DIAMOND (Web Service DIAGnosability, MONitoring and Diagnosis) European project. In this context, a request is sent to a web service which collaborates with other services to provide the adequate reply. Faults may propagate from one service to another and diagnosing them is a crucial issue, in order to react properly. We use a simplified example of an e-foodshop to illustrate our proposal.

We first recall the principles of the chronicle recognition approach and give basic definitions in Sect. 2. We introduce in Sect. 3 the simplified example that will be used all along this paper. In Sect. 4, we show how to extend the chronicle-based approach to distributed systems. We first describe the architecture of a chronicle-based distributed system (4.1). Then we extend the chronicle formalism to deal with synchronization constraints (4.2). We describe in 4.3 a push-pull algorithm able to compute a global diagnosis from the local diagnoses, computed by locally distributed chronicle recognition systems, by checking the synchronization constraints. After an illustrative example in 4.4, we compare our proposal to related work in Sect. 5 and conclude in Sect. 6.

2 Chronicle Recognition Approach

The chronicle recognition approach (first introduced in [1]) relies on a set of patterns, named chronicles, which constitutes the chronicle base. Let us recall the formalism and the chronicle recognition algorithm.

2.1 Formalism of Chronicles

A chronicle is a set of observable events which are time-constrained and is characteristic of a situation.

An **event type** defines what is observed within the system, *e.g.* the name of an activity act , the name augmented with the fact that act is starting (namely act^+) or ending (namely act^-), the name enriched with observable parameters $act(?var_1, \dots, ?var_n)$ or a combination of those possibilities. \mathcal{E} denotes the set of possible event types. An **event** is a pair $(e, ?t)$ where $e \in \mathcal{E}$ is an event type and $?t$ the occurrence date of the event.

A **chronicle** (model) \mathcal{C} is a pair $(\mathcal{S}, \mathcal{T})$ where \mathcal{S} is a set of events and \mathcal{T} a set of constraints between their occurrence dates. When its variables and its occurrence dates are instantiated, a chronicle is called a **chronicle instance**.

2.2 Chronicle Recognition

A chronicle recognition tool, called CRS (Chronicle Recognition System), has been developed by C. Dousson¹. It is in charge of analyzing the input stream of events and of identifying, on the fly, any pattern matching a situation described by a chronicle. Chronicles are compiled into temporal constraint networks which are processed by efficient graph algorithms. CRS is based on a complete forecast of the possible dates for

¹ <http://crs.elibel.tm.fr/>

each event that has not occurred yet. This set (called temporal window) is reduced by propagation of the dates of observed events through the temporal constraint network. When a new event arrives in the input stream, new instances of chronicles are generated in the set of hypotheses, which is managed as a tree. Instances are discarded as soon as possible, when constraints are violated or when temporal windows become empty.




Fig. 1. Principle of chronicle recognition

Figure 1 shows the principle of the recognition algorithm on a very simple example: a single chronicle model is defined, containing only two events: $(a, ?t_a)$ and $(b, ?t_b)$, with $?t_a + 1 \leq ?t_b \leq ?t_a + 3$. When event $(a, 1)$ is received, instance $I1$ is created, which updates the temporal window of the related node b . When a new event $(a, 3)$ occurs, a new instance $I2$ is created and the forthcoming temporal window of $I1$ is updated. When event $(b, 5)$ is received, instance $I3$ is created (from $I2$) and $I1$ is destroyed as no more event $(b, ?t_b)$ could match the temporal constraints from now on. Instance $I2$ is still waiting for another potential event $(b, ?t_b)$ before $?t_b > 6$. As all the events of $I3$ are instantiated, this instance is recognized.

3 Motivating Example

To illustrate the ideas developed in this paper, we consider an orchestration of three web services, a shop, a supplier and a warehouse, that provide e-shopping capabilities to users. This application keeps the essential properties of the applications we aim to monitor. In particular, we consider closed environments, where a workflow-like description of each web service (Fig. 2) involved in the processing of the request is supposed to be available.

A customer wants to place an order and selects items on the shop. This list of items is transferred to a supplier which sends a reservation request to a warehouse, for each item of the list. The warehouse returns an acknowledgement to the supplier for each item request and, at the end of the item list, the supplier sends a list of the available items to the shop which forwards it to the customer. The customer agreement terminates the process.

Faults may happen during this process. Figure 2.a shows two of them (represented by pentagons), related with the shop. First, when placing his order, the customer may

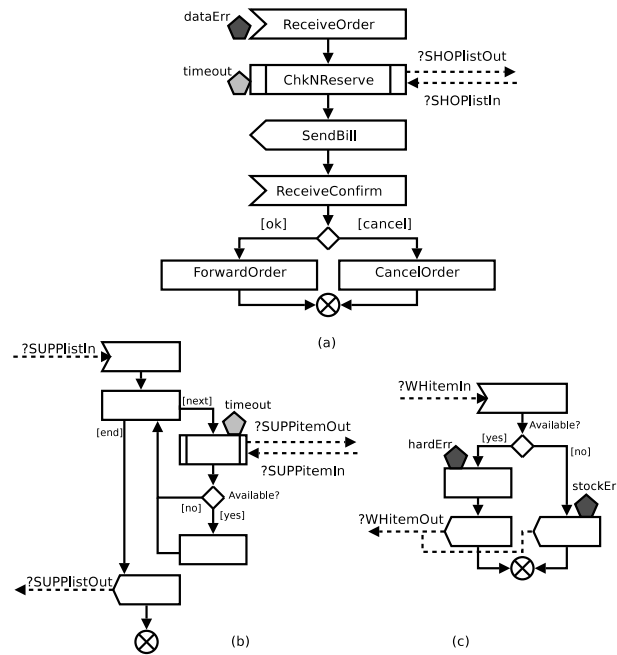


Fig. 2. (a) Workflow of the SHOP and reduced workflows of (b) the SUPP and (c) the WH

make a data acquisition error, which may result in unexpected items on his reservation list. Then, a timeout may occur when calling the supplier. We consider that only timeouts may occur on the supplier (Fig. 2.b), when calling the warehouse. On the warehouse (Fig. 2.c), things are more complicated. First, an item may be out of stock, resulting in an incomplete reservation list. Then, an internal error may happen, resulting in a denial of service.

Figure 3 presents two processes that may result in the same observation on the shop, *i.e.* a cancellation of the order due to an incorrect reservation list: (a) a data acquisition error, ordering “eggs and teak” instead of “eggs and tea”, for instance, and (b) a stock error happening on the warehouse. Here, we notice that two distinct errors that happen on two distinct services can result in the same local problem, hence the necessity of diagnosing the system globally in order to repair in an adequate way.




Fig. 3. Two scenarii that may result in a cancelled order

4 Extension to Distributed Environments

Diagnosing distributed systems thanks to chronicles requires to define a modular diagnosis architecture capable of merging diagnoses provided by local chronicle-based diagnosers and to enrich the chronicle formalism with synchronization constraints.

4.1 Architecture

Figure 4 summarizes our chronicle-based approach architecture. This decentralized system is composed of a global diagnoser (or broker) in charge of merging the local diagnoses sent by each service and sending global diagnoses to a repair module. Services are composed of the web service itself, logs generated in real time by the web service, a base of chronicles generated off-line, a local diagnoser that uses the logs to instantiate chronicles from the base.




Fig. 4. General architecture of a distributed system

4.2 Extension of the Formalism of Chronicles

As a fault occurring on a service often propagates to other services, we base our approach on the merging of local diagnoses. As a consequence, we enrich the initial formalism of chronicles with synchronization constraints that allow the broker to spot homologous chronicles and merge them.

Before defining a distributed chronicle, let us firstly define what is a synchronization point.

The **status of a variable** is a boolean that denotes if the value of a chronicle variable is normal ($\neg err$) or abnormal (err) in a given execution case. A **synchronization variable** is a pair $(?var, status)$ where $?var$ is a (non temporal) chronicle variable and $status$ the status of this variable inside a given chronicle model.

A **synchronization point** is a tuple $(e, \{vars\}, serv_{type})$ where e is an event type, $\{vars\}$ a set of synchronization variables linked with this event type and $serv_{type}$ a type of remote service the local service communicates with. An **instance of a synchronization point** is a synchronization point in which variables are instantiated, and $serv_{type}$ is instantiated as the effective address of the remote service.

A synchronization point is **incoming** if it corresponds to a $serv_{remote} \rightarrow serv_{local}$ communication, **outgoing** for the contrary (see example chronicle below).

Referring to Fig. 2.a and Sect. 3, here is one of the two synchronization points on the SHOP, which is instantiated as follows, in the execution case of an external error (see Fig. 6):

$(ChkNReserve^+, \{(?SHOPlistIn, err)\}, supplier)$.

It expresses the fact that the error is coming from the supplier, through the ?SHOPlistIn variable, which is received by the SHOP at the end of the ChkNReserve activity.

A distributed chronicle is a classical chronicle enriched with a “color” and a “synchronization” part, so that we can merge it with chronicles from adjacent services.

The **color of a chronicle** \mathcal{K} represents the degree of importance of a chronicle and its capacity to trigger a global diagnosis process. Two colors are used: *red* for faults that may trigger the broker and *green* for normal behaviors and non critical faults.

Distributed chronicle: a distributed chronicle is a tuple $\mathcal{C}_D = (\mathcal{S}, \mathcal{T}, \mathcal{O}, \mathcal{I}, \mathcal{K})$ where \mathcal{S} is a set of events, \mathcal{T} a graph of constraints between their occurrence dates, \mathcal{O} and \mathcal{I} are respectively two sets of outgoing and incoming synchronization points, and \mathcal{K} is the color of the chronicle.

Let us consider the chronicle describing the external error case. We have the distributed chronicle model $\mathcal{C}_D = (\mathcal{S}, \mathcal{T}, \mathcal{O}, \mathcal{I}, \mathcal{K})$:

$$\begin{aligned} \mathcal{S} = & \{ (ReceiveOrder, ?t_1), \\ & (ChkNReserve^-(?SHOPlistOut), ?t_2), \\ & (ChkNReserve^+(?SHOPlistIn), ?t_3), \\ & (SendBill, ?t_4), \\ & (ReceiveConfirm, ?t_5), \\ & (ForwardOrder, ?t_6) \\ & \} \\ \mathcal{T} = & \{ ?t_1 < ?t_2, ?t_2 < ?t_3, ?t_3 < ?t_4, ?t_4 < ?t_5 < ?t_6 \} \\ \mathcal{O} = & \{ (ChkNReserve^-, \{(?SHOPlistOut, -err)\}, supplier) \} \\ \mathcal{I} = & \{ (ChkNReserve^+, \{(?SHOPlistIn, err)\}, supplier) \} \\ \mathcal{K} = & red \end{aligned}$$

This chronicle triggers the broker, hence its *red* color. Having defined chronicles for each behavior of each service taking part in the foodshop orchestration, we have the tables shown in Fig. 6, in which we only give the synchronization part of the chronicles. *red* chronicles are written in bold case.

4.3 Algorithms

Our approach consists in merging local chronicles in order to compute a set of candidate global diagnoses. This set of diagnoses is represented by a diagnosis tree as explained below. There are two steps in the global diagnosis process (Fig. 5). In a first step, at “push” time, local diagnosers send recognized chronicles to the broker, which triggers the global diagnosis process. In a second step, at “pull” time, *i.e.* when the

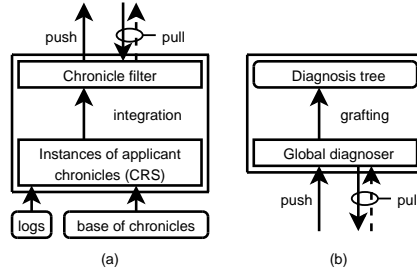


Fig. 5. Working of the (a) local and (b) global diagnosers

global diagnoser needs information, it queries local diagnosers their chronicles recognized previously or in future. This push-pull mechanism is implemented through a filter as explained below.

The computation of the local diagnosis relies on a CRS module fed by the logs of the web service and sending its recognized chronicles to the global diagnoser (Fig. 5).

In order to avoid sending useless chronicles, a filter \mathcal{M} is set for each running process. In *filter* mode, only *red* recognized chronicles are sent to the global diagnoser. Green chronicles are stored in a chronicle buffer \mathcal{C}_{buf} . Nevertheless, at “pull” time, the global diagnoser can change \mathcal{M} from *filter* to *open*, which flushes \mathcal{C}_{buf} in order to provide the global diagnoser with all the available information. In *open* mode, both *red* or *green* newly recognized chronicles will be directly sent to the global diagnoser. Algorithm 1 illustrates this operation.

```

init: mode  $\mathcal{M} := filter$ , chronicle set  $\mathcal{C}_{buf} := \emptyset$ ;
on event chronicle  $c$  recognized do
  | if  $(\mathcal{M} = filter \wedge c.color = red) \vee \mathcal{M} = open$  then
  |   | Broker.push( $c$ );
  | else
  |   |  $\mathcal{C}_{buf} := \mathcal{C}_{buf} \cup \{c\}$ ;
  | end
end
on event LocalDiagnoser.pull() do
  | foreach  $c \in \mathcal{C}_{buf}$  do Broker.push( $c$ );
  |  $\mathcal{C}_{buf} := \emptyset$ ,  $\mathcal{M} := open$ ;
end

```

Algorithm 1: Local diagnoser management

The global diagnoser algorithm relies on a diagnosis tree \mathcal{D}_t in charge of treasuring all the candidate diagnoses under the shape of partially recognized global chronicles (Fig. 5.(b)). Each candidate diagnosis is represented by a path leading to a constraintless node in \mathcal{D}_t . The global diagnoser algorithm (Algorithm 2) manages this tree and queries local diagnosers in order to make it grow and complete the pending paths.

The initial diagnosis tree only contains the *emptynode* which, being constraintless, is compatible with any recognized chronicle. When a recognized chronicle c is sent by a service s to the global diagnoser, two operations are performed. First, \mathcal{D}_t is traversed, trying to combine each node n with c thanks to the status of corresponding variables. In case of a compatibility between n and c , a child node containing c and the synchronization constraints that remain to check is grafted under n in \mathcal{D}_t . Then, the global diagnoser changes to *open* the mode of all the services mentioned in c in order to collect all the information needed for a global diagnosis (Algorithm 2).

```

init: diagnosis tree  $\mathcal{D}_t := emptynode$ ;
on event Broker.push(chronicle c) do
  | foreach node n of  $\mathcal{D}_t$  do
  | | if c compatible with n then
  | | |  $n.addChild(c)$ ;
  | | end
  | end
  | foreach service s mentioned in c do
  | |  $s.LocalDiagnoser.pull()$ ;
  | end
end

```

Algorithm 2: Global diagnoser management

When a candidate diagnosis (*i.e.* a constraintless node) is computed in \mathcal{D}_t , the broker forwards it to an external repair module and proceeds with the exhibition of other candidate diagnoses.

4.4 Illustration on the Example

The following example was tested on a distributed chronicle-based diagnosis platform called CARDECRES [11], developed during this PhD thesis.

We consider a customer placing an order on the SHOP, order which is forwarded to the SUPPLIER. For each product of the item list, the SUPPLIER calls the Warehouse so as to book the corresponding product. Unfortunately, a product is missing which provokes the recognition of the *WH:stockErr* chronicle, the color of which is *green*, because the WH doesn't consider being out of stock as an error. The broker is not triggered and the execution goes on. But when the SUPPLIER receives the negative answer of the WH, the *red* chronicle *SUPPLIER:extErr* is recognized and the SUPPLIER "pushes" this chronicle towards the broker, triggering a global diagnosis process while the service execution goes on.

\mathcal{D}_t only contains the root node, at this point. This node is compatible with the constraints of *SUPPLIER:extErr*, listed in Fig.6, and a new node containing *SUPPLIER:extErr* and its constraints is grafted under the root node. After this, the broker changes to *open* the mode of WH, "pulling" the previously recognized *WH:stockErr* chronicle towards it.

\mathcal{D}_t now contains two nodes. *WH:stockErr* is compatible with the empty root node, which results in the grafting of a child node under the root, containing *WH:stockErr*

| SHOP | ?listOut | ?listIn |
|----------------|------------|------------|
| normal | $\neg err$ | $\neg err$ |
| dataErr | err | err |
| extErr | $\neg err$ | err |
| timeout | $\neg err$ | undef |

| SUPP | ?listIn | ?itemOut | ?itemIn | ?listOut |
|----------------|------------|------------|------------|------------|
| normal | $\neg err$ | $\neg err$ | $\neg err$ | $\neg err$ |
| fwdErr | err | err | err | err |
| extErr | $\neg err$ | $\neg err$ | err | err |
| timeout | $\neg err$ | $\neg err$ | undef | undef |

| WH | ?itemIn | ?itemOut |
|----------------|------------|------------|
| normal | $\neg err$ | $\neg err$ |
| fwdErr | err | err |
| stockErr | $\neg err$ | err |
| hardErr | $\neg err$ | undef |

Fig. 6. Chronicles of the three web services

and its constraints. *WH:stockErr* is also compatible with *SUPP:extErr*, as the homologous variables have the same status: *?SUPPitemOut* and *?WHitemIn* are $\neg err$, *?SUPPitemIn* and *?WHitemOut* are *err*. This way, a child node is grafted under *SUPP:extErr*, containing *WH:stockErr* and the remaining unchecked constraints (Fig. 7).

The “pulling” process goes on, interrogating the SHOP and waiting for its recognized chronicles. At the end of the orchestration execution, \mathcal{D}_t exhibits a single constraintless node, which is then the unique candidate diagnosis: *SHOP:extErr*, *SUPP:extErr*, *WH:stockErr*.

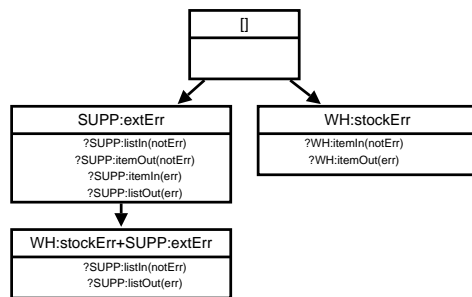


Fig. 7. Intermediate diagnosis tree

4.5 A Word About Complexity

Let us consider the complexity of such an approach. On the local side, the complexity only depends on CRS, which has already been successfully used in large scale systems.

Some basic rules about chronicle writing allow to optimize the use of CRS: PID filtering avoids the recognition of useless cross-process chronicles, delays in chronicle models flush chronicle instances automatically, etc.

On the broker side, the size of the tree only depends on the number of chronicles recognized on each service, hence a need for discriminating and exclusive chronicles. In the worst case, considering all the chronicles are compatible, we demonstrate that the maximum number of nodes in \mathcal{D}_t is

$$n_{max} = \prod_{s \in S} (|\mathcal{C}_s| + 1)$$

with S the set of implied services and \mathcal{C}_s the set of chronicles recognized on s .

5 Related Work and Discussion

Within the context of the supervision of dynamic systems, many works use the formalism of chronicles [4, 2, 3, 12, 6]. Nevertheless, few deal with using chronicles in a distributed context. The approach presented in [13] focuses on temporal aspects and proposes a distributed checking of temporal constraints (by introducing both local and global temporal constraints). In [14], the authors study the problem of acquiring chronicles from the fault model of a system, described with Petri nets. They use a method based on unfolding Petri nets. The formalism of chronicles is enriched with pre- and post-conditions on the current system state, and the recognition algorithm modified consequently. However, to our knowledge, nobody directly worked on the use of distributed chronicles, in particular on the integration of synchronization constraints between components inside the formalism and on the adaptation of the corresponding algorithm, as we propose in this article.

The way we approach the problem of monitoring dynamic systems from a distributed chronicle-based modeling of the system may be compared with works dealing with distributed approaches of monitoring discrete-event systems, such as [7–9, 15, 10, 16, 17]. In each of those works, local diagnoses computed by the different components of the system are synchronized in order to compute a diagnosis taking into account the constraints between components. For instance, the approach of [9] is not so far away from ours, as they fit parts together to build the system diagnosis, like in a puzzle. Those parts, called *tiles*, are labelled by alarms and represent pieces of trajectories. The main difference between the two approaches, apart from the Petri-net-based formalism they use, is that theirs is fully distributed and uses communications between local components to do the computations, without any supervisor. In our decentralized case, a supervisor is in charge of fitting local chronicles together after having synchronized them, so that a global chronicle could be built. [18] is also interested in software components monitoring. The components of the system are described by Petri nets and each component is associated with a local controller that monitors the evolution of this component, observing the messages exchanged between the component and its neighbors.

Concerning web services monitoring, we can cite [19], the objective of which is to acquire a model as automata that will permit to monitor components thanks to the BPEL description of their process. Closer to us, [20] proposes to use planning tools to

allow the user to express his requests thanks to a high-level language and to control the execution of his plans by interlacing execution and plan update. The authors of [21] are interested in checking on line the consistency between what a web service should do, called a contract, and its effective execution. Contracts are expressed as constraints in a constraint-oriented language, and integrated in the BPEL files under the shape of annotations. Then, monitors, implemented as web services, observe the behavior of the web services and are capable of detecting timeout problems or functional errors. In [22], a quite similar approach relies on a monitoring of plans to monitor requests and uses the KPLTL temporal logic in order to express the specifications that have to be respected.

In [23], the decentralized architecture is close to ours. Each web service is equipped with a local diagnoser generating hypotheses that are consistent with the local model and the observations. A supervisor merges local diagnoses to compute a global one, by propagating hypotheses from a local diagnoser to its neighbors. The main difference is that they rely on a static diagnosis approach: using dependencies between state variables, their approach consists in explaining the alarms that have arisen at a given time. In our case, we monitor the behavior of the components as it evolves. This allows, on the one hand, to identify problems related to alarm firing and, on the other hand, to forestall a potential problem and avoid its occurrence.

6 Conclusion

Our contribution in this paper is to propose a *distributed* chronicle-based monitoring and diagnosis approach. Even if it is now recognized that distributed approaches are the only realistic way to monitor large-scale systems, no work exist, to our knowledge, as far as chronicle-based approaches are concerned. We propose a distributed architecture in which a broker service is in charge of synchronizing the local diagnoses computed from chronicles at the component level. We extend the formalism of chronicles and introduce synchronization points that express the synchronization constraints which are checked by the broker according to a push-pull mechanism. We describe the main algorithms and illustrate them on a simplified e-shopping example. A platform has been developed and allows us to make experiments in the framework of the WS-DIAMOND European project, dedicated to the monitoring of software components.

The main perspectives are twofold. The first one is to couple the diagnosis service with a repair service (developed by a partner of ours), the goal being to ensure a good QoS, even in case of fault occurrences. The second one is to build acquisition tools to help building the set of local chronicles, starting from workflow descriptions. A first step in this direction can be found in [19].

References

1. Dousson, C., Gaborit, P., Ghallab, M.: Situation recognition: representation and algorithms. In: Proc. of the Int. Joint Conf. on Artificial Intelligence (IJCAI'93). (1993) 166–172
2. Cordier, M.O., Krivine, J., Laborie, P., Thiébaux, S.: Alarm processing and reconfiguration in power distribution systems. In: Proc. of IEA-AIE'98. (1998) 230–240

3. Dojat, M., Ramaux, N., Fontaine, D.: Scenario recognition for temporal reasoning in medical domains. *Artificial Intelligence in Medicine* **14**(1-2) (1998) 139–155
4. Cordier, M.O., Dousson, C.: Alarm driven monitoring based on chronicles. In: *Proc. of Safeprocess'2000*. (2000) 286–291
5. Pencolé, Y., Cordier, M.O., Rozé, L.: Incremental decentralized diagnosis approach for the supervision of a telecommunication network. In: *IEEE Conf. on Decision and Control (CDC'02)*. (2002) 435–440
6. Aguilar, J., Bousson, K., Dousson, C., Ghallab, M., Guasch, A., Milne, R., Nicol, C., Quevedo, J., Travé-Massuyès, L.: Tiger: real-time situation assessment of dynamic systems. Technical report (1994)
7. Baroni, P., Lamperti, G., Pogliano, P., Zanella, M.: Diagnosis of a class of distributed discrete-event systems. *IEEE Transactions on systems, man, and cybernetics* (2000) 731–752
8. Debouk, R., Lafortune, S., Teneketzis, D.: Coordinated decentralized protocols for failure diagnosis of discrete event systems. *Discrete Event Dynamic Systems* **10**(1-2) (2000) 33–86
9. Aghasaryan, A., Fabre, E., Benveniste, A., Boubour, R., Jard, C.: Fault detection and diagnosis in distributed systems : an approach by partially stochastic petri nets. *Discrete Event Dynamic Systems* **8**(2) (1998) 203–231
10. Pencolé, Y., Cordier, M.O.: A formal framework for the decentralised diagnosis of large scale discrete event systems and its application to telecommunication networks. *Artificial Intelligence Journal* **164**(1-2) (2005) 121–170
11. Le Guillou, X., Cordier, M.O., Robin, S., Roze, L.: Chronicles for on-line diagnosis of distributed systems. Internal IRISA report #1890 (2008)
12. Quiniou, R., Cordier, M.O., Carrault, G., Wang, F.: Application of ilp to cardiac arrhythmia characterization for chronicle recognition. In: *ILP'2001*. (2001) 220–227
13. Boufaied, A., Subias, A., Combaceau, M.: Distributed fault detection with delays consideration. In: *Proc. of the 15th Int. Workshop on Principles of Diagnosis (DX'04)*. (2004)
14. Guerraz, B., Dousson, C.: Chronicles construction starting from the fault model of the system to diagnose. In: *Proc. of the 15th Int. Workshop on Principles of Diagnosis (DX'04)*. (2004) 51–56
15. Jiroveanu, G., Boël, R.: Petri net model-based distributed diagnosis for large interacting systems. In: *Proc. of the 16th Int. Workshop on Principles of Diagnosis (DX'05)*. (2005)
16. Roos, N., Teije, A., Bos, A., Witteveen, C.: An analysis of multi-agent diagnosis. In: *Proc. of the 1st Int. Joint Conf. on Autonomous Agents and MultiAgent Systems (AAMAS'02)*. (2002)
17. Provan, G.: A model-based diagnosis framework for distributed systems. In: *Proc. of the 13th Int. Workshop on Principles of Diagnosis (DX'02)*. (2002) 16–25
18. Grosclaude, I.: Model-based monitoring of component-based software systems. In: *Proc. of the 15th Int. Workshop on Principles of Diagnosis (DX'04)*. (2004) 51–56
19. Yan, Y., Pencolé, Y., Cordier, M.O., Grastien, A.: Monitoring web service networks in a model-based approach. In: *3rd European Conf. on Web Services (ECOWS)*. (2005)
20. Lazovik, A., Aiello, M., Papazoglou, M.: Planning and monitoring the execution of web service requests. In: *Proc. of the 1st Int. Conf. on Service-Oriented Computing (ICSOC'03)*. Volume 2910 of *Lecture Notes in Computer Science*. (2003) 335–350
21. Baresi, L., Ghezzi, C., Guinea, S.: Smart monitors for composed services. In: *Proc. of the 2nd Int. Conf. on Service-Oriented Computing (ICSOC'04)*. (2004) 193–202
22. Barbon, F., Traverso, P., Pistore, M., Trainotti, M.: Run-time monitoring of instances and classes of web service compositions. In: *Proc. of the IEEE Int. Conf. on Web Services (ICWS'06)*. (2006) 63–71
23. Ardissono, L., Console, L., Goy, A., Petrone, G., Picardi, C., Segnan, M., Theseider Dupré, D.: Cooperative model-based diagnosis of web services. In: *Proceedings of DX'05, International Workshop on the Principles of Diagnosis*, Pacific Grove, California (2005)