

Supervision Patterns in Discrete Event Systems Diagnosis *

Thierry Jéron, Hervé Marchand, Sophie Pinchinat, Marie-Odile Cordier
IRISA, Campus Universitaire de Beaulieu, 35042 Rennes, France
{firstName.Name}@irisa.fr

Abstract

In this paper, we are interested in the diagnosis of discrete event systems modeled by finite transition systems. We propose a model of supervision patterns general enough to capture past occurrences of particular trajectories of the system. Modeling the diagnosis objective by supervision patterns allows us to generalize the properties to be diagnosed and to render them independent of the description of the system. We first formally define the diagnosis problem in this context. We then derive techniques for the construction of a diagnoser and for the verification of the diagnosticability based on standard operations on transition systems. We show that these techniques are general enough to express and solve in a unified way a broad class of diagnosis problems found in the literature, e.g. diagnosing permanent faults, multiple faults, fault sequences and some problems of intermittent faults.

1 Introduction

Diagnosing and monitoring dynamical systems is an increasingly active research domain and model-based approaches have been proposed which differ according to the kind of models they used [13; 1; 11; 10; 3; 4]. The general diagnosis problem is to detect or identify patterns of particular events on a partially observable system. This paper focuses on discrete-event systems modeled as finite state machines. In this context, patterns usually describe the occurrence of a fault [12; 13], multiple occurrences of a fault [7], the repair of a system after the occurrence of a fault [2]. The aim of diagnosis is to decide, by means of a *diagnoser*, whether or not such a pattern occurred in the system. Even if such a decision cannot be taken immediately after the occurrence of the pattern, one requires that this decision has to be taken in a bounded delay. This property

is usually called *diagnosability*. This property can be checked *a priori* from the system model, and depends on its observability and on the kind of patterns which are looked for.

However, the approaches in the literature suffer from some deficiencies. One observes many different definitions of diagnosability and ad hoc algorithms for the construction of the diagnoser, as well as for the verification of diagnosability. As a consequence, all these results are difficult to reuse for new but similar diagnosis problems. We believe that the reason comes from an absence of a clear definition of the involved patterns, which would clarify the separation between the diagnosis objective and the specification of the system.

In this paper, we formally introduce the notion of supervision pattern as a means to define the diagnosis objectives: a supervision pattern is an automaton which language is the set of trajectories one wants to diagnose. The proposal is general enough to cover in an unified way an important class of diagnosis objectives, including detection of permanent faults, but also transient faults, multiple faults, repeating faults, as well as quite complex sequences of events.

We then propose a formal definition of the Diagnosis Problem in this context. The essential point is a clear definition of the set of trajectories *compatible* with an observed trace. Now, the Diagnosis Problem is expressed as the problem of synthesizing a function over traces, the *diagnoser*, which decrees on the possible/certain occurrence of the pattern on trajectories compatible with the trace. The diagnoser is required to fulfil two fundamental properties: *correctness* and *bounded diagnosability*. *Correctness* expresses that the diagnoser answers accurately and *Bounded Diagnosability* guarantees that only a bounded number of observations is needed to eventually answer with certainty that the pattern has occurred. *Bounded Diagnosability* is formally defined as the Ω -diagnosability of the system (where Ω is the supervision pattern), which compares to standard diagnosability by [13]. Relying on the formal frame-

*This paper is available as Inria report n° XXXX. A short version has also been submitted to the Workshop on Discrete Event System (Wodes'06).

work we have developed, we then propose algorithms for both the diagnoser's synthesis, and the verification of Ω -diagnosability. We believe that these generic algorithms as well as their correctness proofs are a lot more simple than the ones proposed in the literature.

The paper is organized as follows. In section 2, we recall standard definitions and notations on labeled transition systems, as well as the notion of compatible trajectories of an observable trace. Supervision patterns are introduced in section 3. The diagnosis problem and the Ω -diagnosability are then defined. Section 4 is dedicated to algorithms and their associated proofs, for the construction of a correct diagnoser as well as the verification of Ω -diagnosability. Finally, Section 5 illustrates the approach with an example.

2 Labelled Transition Systems and Related Notions

We start first by recalling useful standard notations: We assume given an alphabet Σ , that is a finite set $\{\sigma, \sigma_1, \dots\}$. The set of finite sequences over Σ is denoted by Σ^* , with ϵ for the empty sequence. In the paper, typical elements of Σ^* are s, t, u, \dots . For each $s, t \in \Sigma^*$ of the form $s = \sigma_1 \dots \sigma_n$ and $t = \sigma'_1 \dots \sigma'_m$ ($n, m \in \mathbb{N}$), the *concatenation of s and t* is still a sequence defined by $s.t = \sigma_1 \dots \sigma_n \sigma'_1 \dots \sigma'_m$. The *length* of $s \in \Sigma^*$ is denoted $\|s\|$.

We now come to the models of systems:

Definition 1 (LTS) An LTS over Σ is defined by a 4-tuple $M = (Q, \Sigma, \rightarrow, q_0)$ where Q is a finite set of states with a distinguished element q_0 called the initial state, Σ is the set of events of M , $q_0 \in Q$ is the initial state, and $\rightarrow \subseteq Q \times \Sigma \times Q$ is the partial transition relation.

In the rest of the section, we assume given an LTS $M = (Q, \Sigma, \rightarrow, q_0)$,

- we write $q \xrightarrow{\sigma} q'$ for $(q, \sigma, q') \in \rightarrow$. We extend \rightarrow to arbitrary sequences by setting: $q \xrightarrow{\epsilon} q$ always holds, and $q \xrightarrow{s\sigma} q'$ whenever $q \xrightarrow{s} q''$ and $q'' \xrightarrow{\sigma} q'$, for some $q'' \in Q$.
- Let $q \xrightarrow{s}$ mean that $q \xrightarrow{s} q'$ for some $q' \in Q$. The *event set* of a state $q \in Q$ is $\Sigma(q) \triangleq \{\sigma \in \Sigma \mid q \xrightarrow{\sigma}\}$.
- A state q is *reachable* if $\exists s \in \Sigma^*, q_0 \xrightarrow{s} q$.
- We set $\Delta_M(q, s) \triangleq \{q' \in Q \mid q \xrightarrow{s} q'\}$. In particular $\Delta_M(q, \epsilon) \triangleq \{q\}$. By abuse of notation, for any language $L \subseteq \Sigma^*$ $\Delta_M(q, L) \triangleq \{q' \in Q \mid q \xrightarrow{s} q' \text{ for some } s \in L\}$, and for any $Q' \subseteq Q$, $\Delta_M(Q', L) = \bigcup_{q \in Q'} \Delta_M(q, L)$.
- A subset $Q' \subseteq Q$ is *stable* whenever $\Delta_M(Q', \Sigma) \subseteq Q'$.
- M is *alive* if $\Sigma(q) \neq \emptyset$, for each $q \in Q$. It is *complete* whenever $\Sigma(q) = \Sigma$, for each $q \in Q$.

- We say that M is *deterministic* if whenever $q \xrightarrow{\sigma} q'$ and $q \xrightarrow{\sigma} q''$, then $q' = q''$, for each $q \in Q$ and each $\sigma \in \Sigma$.

The *language generated* by the system M is the set $\mathcal{L}(M) \triangleq \{s \in \Sigma^*, q_0 \xrightarrow{s}\}$ which elements are called *trajectories* of M . Given a trajectory $s \in \mathcal{L}(M)$, we write

$$\mathcal{L}(M)/s \triangleq \{t \in \Sigma^* \mid s.t \in \mathcal{L}(M)\}$$

for the set of trajectories that extend s in M .

Rapidly in the paper, we will need to distinguish a subset $Q_m \subseteq Q$ to denote final states. The notions above are extended in this setting by letting $\mathcal{L}_{Q_m}(M) = \{\sigma \in \Sigma^* \mid \Delta_M(q_0, \sigma) \subseteq Q_m\}$.

A useful operation on LTS is the synchronous product that allows to intersect languages of two LTSs.

Definition 2 Let $M^i = (Q^i, q_0^i, \Sigma, \rightarrow_i)$, $i = 1, 2$, be two LTSs. Their synchronous product is $M^1 \times M^2 = (Q^1 \times Q^2, (q_0^1, q_0^2), \Sigma_1 \cup \Sigma_2, \rightarrow)$, where $\rightarrow \subseteq Q^1 \times Q^2$ satisfies $(q^1, q^2) \xrightarrow{\sigma} (q'^1, q'^2)$ whenever $q^1 \xrightarrow{\sigma_1} q'^1$ and $q^2 \xrightarrow{\sigma_2} q'^2$.

Clearly $\mathcal{L}(M^1 \times M^2) = \mathcal{L}(M^1) \cap \mathcal{L}(M^2)$. and for $Q_1 \subseteq Q^1$ and $Q_2 \subseteq Q^2$, we also have $\mathcal{L}(M^1 \times M^2, Q_1 \times Q_2) = \mathcal{L}_{Q_1}(M^1) \cap \mathcal{L}_{Q_2}(M^2)$. Also, if two sets $Q_1 \subseteq Q^1$ and $Q_2 \subseteq Q^2$ are stable, $Q_1 \times Q_2$ is stable in $M^1 \times M^2$.

As we are interested in diagnosing systems - this will be formalized in the next section -, partial observation plays a central rôle. In this regard, the set of events Σ is partitioned into Σ_o and Σ_{uo}

$$\Sigma = \Sigma_o \cup \Sigma_{uo}, \text{ and } \Sigma_o \cap \Sigma_{uo} = \emptyset,$$

where Σ_o represents the set of *observable* events - elements of Σ_{uo} are then *unobservable* events. Typical elements of Σ_o^* will be denoted by μ, μ' .

We say that M is Σ_o -*alive* if $\forall q \in Q, \exists s \in \Sigma_o^*. q \xrightarrow{s}$, meaning that there is no terminal loop of unobservable events. Notice that when M has no loop of unobservable events, M is alive if and only if M is Σ_o -alive.

Let $P : \Sigma^* \rightarrow \Sigma_o^*$ be the natural *projection* of trajectories onto Σ_o^* defined by: $P(\epsilon) = \epsilon$ and $P(s\sigma) = P(s).\sigma$ if $\sigma \in \Sigma_o$, and $P(s)$ otherwise. The projection P simply erases the unobservable events from a trajectory. P extends to languages by defining, for $L \subseteq \Sigma^*$, $P(L) = \{P(s) \mid s \in L\}$. The inverse projection of L is defined by $P^{-1}(L) = \{s \in \Sigma^* \mid P(s) \in L\}$.

Now, the *language of traces* of M is

$$\text{Traces}(M) \triangleq P(\mathcal{L}(M))$$

It is the set of observable sequences of its trajectories.

From the projection P , we derive an equivalence relation between trajectories of M , written \equiv_M , called the *Delay-Observation equivalence* in reference to the delay-bisimulation of [9]:

Definition 3 (Delay-Observation Equivalence, \equiv_M) Let $\equiv_M \subseteq \mathcal{L}(M) \times \mathcal{L}(M)$ be the binary relation defined by $s \equiv_M s'$ whenever

- $P(s) = P(s')$ and
- $s \in \Sigma^* \Sigma_o$ if and only if $s' \in \Sigma^* \Sigma_o$.

One easily verifies that \equiv_M is an equivalence relation, and we take the convention to write $[s]$ for the equivalence class of s .

Given $s \in \mathcal{L}(M)$, s naturally maps onto a trace of M , namely $P(s)$. Now, given a non empty trace μ of M , μ does not uniquely determine a Delay-Observation equivalence class as in general μ can be brought back in M in two different manners:

- μ can be associated with the class $[s]$ with $P(s) = \mu$ and $s \in \Sigma^* \Sigma_o$,
- or μ can be associated with the class $[s']$ with $P(s') = \mu$ and $s' \in \Sigma^* \Sigma_{uo}$

Notice that by Definition 3, $[s]$ and $[s']$ are different. Henceforth, we take the convention that the equivalence class denoted by a trace μ is

$$\llbracket \mu \rrbracket_M \triangleq \begin{cases} P^{-1}(\mu) \cap \mathcal{L}(M) \cap \Sigma^* \Sigma_o & \text{if } \mu \neq \epsilon \\ \{\epsilon\} & \text{otherwise.} \end{cases}$$

We say that $\llbracket \mu \rrbracket_M$ is the set of trajectories *compatible* with the trace μ . When clear from the context, we will use $\llbracket \mu \rrbracket$ for $\llbracket \mu \rrbracket_M$. This notion of compatible trajectory will be a central notion for diagnosis as the aim will be to infer properties on the set of trajectories $\llbracket \mu \rrbracket_M$ compatible with the observation of the trace μ . The reason for choosing this definition of $\llbracket \mu \rrbracket$ is that in the case of online diagnosis, it is natural to assume that the diagnoser is reactive to an observable move of the system.

3 Supervision Patterns and the Diagnosis Problem

In this section, we introduce the notion of *supervision patterns*, which are means to define languages we are interested in for diagnosis purpose. We then give some examples of such patterns. Finally, we introduce the diagnosis problem for such patterns.

3.1 Supervision Patterns

Supervision patterns are represented by particular LTSs:

Definition 4 A supervision pattern is a 5-tuple $\Omega = (Q_\Omega, \Sigma, \rightarrow_\Omega, q_{0_\Omega}, Q_F)$, where $(Q_\Omega, \Sigma, \rightarrow_\Omega, q_{0_\Omega})$ is a deterministic and complete LTS, and $Q_F \subseteq Q$ is a distinguished stable subset of states.

As Ω is complete we get $\mathcal{L}(\Omega) = \Sigma^*$. Also notice that the assumption that Q_F is stable means that its accepted language is “extension-closed”, i.e. satisfies $\mathcal{L}_{Q_F}(\Omega) \cdot \Sigma^* = \mathcal{L}_{Q_F}(\Omega)$. Otherwise said, $\mathcal{L}_{Q_F}(\Omega)$ is a language violating a safety property. This choice is natural since we want to diagnose whether all trajectories compatible with an observed trace have a prefix recognized by the pattern.

In the next subsection we will give some examples of supervision patterns which rephrase classic properties one is interested in for diagnosis purpose.

3.2 Examples of supervision patterns

Occurrence of one fault Let $f \in \Sigma$ be a fault and consider that we are interested in diagnosing the occurrence of this fault. A trajectory $s \in \Sigma^*$ is faulty if $s \in \Sigma^* \cdot f \cdot \Sigma^*$. The supervision pattern Ω_f of Figure 1 exactly recognizes this language, $\mathcal{L}(\Omega_f) = \Sigma^* \cdot f \cdot \Sigma^*$.





Figure 1: Supervision pattern for one fault

Occurrence of multiple faults Let f_1 and f_2 be two faults that may occur in the system. Diagnosing the occurrence of these two faults in a trajectory means deciding the membership of this trajectory in $\Sigma^* \cdot f_1 \cdot \Sigma^* \cap \Sigma^* \cdot f_2 \cdot \Sigma^* = \mathcal{L}_{F_1}(\Omega_{f_1}) \cap \mathcal{L}_{F_2}(\Omega_{f_2})$, where $\Omega_{f_i}, i \in \{1, 2\}$ are isomorphic to the supervision pattern Ω_f described in Figure 1. The supervision pattern is then the product $\Omega_{f_1} \times \Omega_{f_2}$ which accepted language in $F_1 \times F_2$ is $\mathcal{L}_{F_1 \times F_2}(\Omega_{f_1} \times \Omega_{f_2}) = \mathcal{L}_{F_1}(\Omega_{f_1}) \cap \mathcal{L}_{F_2}(\Omega_{f_2})$.




More generally, the supervision pattern for the occurrence of a set of faults $\{f_1, \dots, f_l\}$ is the product $\times_{i=1, \dots, l} \Omega_{f_i}$, considering $\times_{i=1, \dots, l} F_i$ as final state set.

Ordered occurrence of events If the diagnosis that has to be performed concerns the occurrences of different faults in a precise order, for example, f_2 after f_1 , the trajectories that have to be recognized by the supervision pattern are


$$\Sigma^* \cdot f_1 \cdot \Sigma^* \cdot f_2 \cdot \Sigma^*,$$

which corresponds to the concatenation of the two languages $\mathcal{L}_{F_1}(\Omega_{f_1}) \cdot \mathcal{L}_{F_2}(\Omega_{f_2})$ as described by the following supervision pattern




If f_1 corresponds to a fault event and f_2 to the repair of this fault in the system, then we actually diagnose the repair of the fault f_1 . With this pattern, the aim is to match the *I-diagnosability* in [2].

Multiple occurrences of the same fault Another interesting problem is to diagnose the multiple occurrences of the same fault event f , say k times. This corresponds to the supervision pattern given below which accepted language is $\mathcal{L}_F(\Omega_f)^k$. The aim is to match the *k*-diagnosability of [7].



This can be easily generalized to a pattern recognizing the occurrence of k patterns (identical or not).

Intermittent Fault The following supervision pattern describes the fact that a fault (occurrence of f) occurred twice without repair (occurrence of r).



It is worthwhile noting that this can be generalized to a pattern recognizing the occurrence of k faults (identical or not) without repair.

3.3 The Diagnosis Problem

In the remainder of the paper, we consider a system whose behavior is modeled by an LTS $G = (Q, \Sigma, \rightarrow, q_0)$. The only assumption made on G is that G is Σ_o -alive. Notice that G can be non-deterministic. We also consider a supervision pattern $\Omega = (Q_\Omega, \Sigma, \rightarrow_\Omega, q_{0_\Omega}, Q_F)$ denoting the language $\mathcal{L}_{Q_F}(\Omega)$ that we want to diagnose.

We define the *Diagnosis Problem* as the problem of defining a function Diag_Ω on traces whose intention is to answer the question whether trajectories corresponding to observed traces are recognized or not by the supervision pattern. We do require some properties for Diag_Ω : *Correctness* and *Bounded Diagnosability*.

- *Correctness* means that “Yes” and “No” answers should be accurate
- *Bounded Diagnosability* means that trajectories in $\mathcal{L}_{Q_F}(\Omega)$ should be diagnosed with finitely many observations.

The *Diagnosis problem* can be stated as follows: given an LTS G and given a supervisory pattern Ω , decide whether there exists (and compute if any) a three valued function $\text{Diag}_\Omega : \text{Traces}(G) \rightarrow \{\text{“YES”}, \text{“NO”}, \text{“?”}\}$ decreeing, for each trace μ of G , on the membership in $\mathcal{L}_{Q_F}(\Omega)$ of any trajectory in $\llbracket \mu \rrbracket$. Formally,

- (Diagnosis Correctness) The function should verify

$$\text{Diag}_\Omega(\mu) = \begin{cases} \text{“YES”} & \text{if } \llbracket \mu \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega) \\ \text{“NO”} & \text{if } \llbracket \mu \rrbracket \cap \mathcal{L}_{Q_F}(\Omega) = \emptyset \\ \text{“?”} & \text{otherwise.} \end{cases}$$
- (Bounded Diagnosability) As G is only partially observed, we expect in general situations where $\text{Diag}_\Omega(\mu) = \text{“?”}$ (as neither $\llbracket \mu \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega)$ nor $\llbracket \mu \rrbracket \cap \mathcal{L}_{Q_F}(\Omega) = \emptyset$ hold). However, we require this undetermined situation not to last in the following sense: There must exist $n \in \mathbb{N}$, the bound, such that whenever $s \in \llbracket \mu \rrbracket \cap \mathcal{L}_{Q_F}(\Omega)$, for all $t \in \mathcal{L}(G)/s \cap \Sigma^* \cdot \Sigma_o$, if $\|P(t)\| \geq n$ then $\text{Diag}_\Omega(P(s.t)) = \text{“YES”}$.

Diagnosis Correctness means that the diagnosis of a trace μ is “No” if no trajectory in its semantics $\llbracket \mu \rrbracket$ lies in $\mathcal{L}_{Q_F}(\Omega)$ while it is “Yes” if all trajectories in $\llbracket \mu \rrbracket$ lie in $\mathcal{L}_{Q_F}(\Omega)$. Bounded Diagnosability means that when observing a trajectory in $\mathcal{L}_{Q_F}(\Omega)$, a “Yes” answer should be produced after finitely many observable events (See 2 for an intuitive explanation of these notions).




Figure 2: The Ω -diagnosability for $\Omega = \Omega_f$

Now, if Diag_Ω provides a Correct Diagnosis, Bounded Diagnosability can be rephrased, by replacing $\text{Diag}_\Omega(P(s.t)) = \text{“YES”}$ with $\llbracket P(s.t) \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega)$. We obtain what we call the *Ω -diagnosability*. Notice that this is now a property of G with respect to Ω .

Definition 5 An LTS G is $\Omega(n)$ -diagnosable, where $n \in \mathbb{N}$, whenever

$$\forall s \in \mathcal{L}_{Q_F}(\Omega) \cap \mathcal{L}(G) \cap \Sigma^* \Sigma_o, \forall t \in \mathcal{L}(G)/s \cap \Sigma^* \Sigma_o, \text{ if } \|P(t)\| \geq n \text{ then } \llbracket P(s.t) \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega), \quad (1)$$

We say that G is Ω -diagnosable if G is $\Omega(n)$ -diagnosable for some $n \in \mathbb{N}$.

Ω -diagnosability says that when a trajectory s ending with an observable event is recognized by the supervision pattern Ω , for any extension t with enough observable events, any trajectory s' compatible with the observation $P(s.t)$ is also recognized by Ω .

The remark before Definition 5 is formalized by :

Proposition 1 If Diag_Ω computes a Correct Diagnosis, then G is Ω -diagnosable if and only if the Bounded Diagnosability Property holds for Diag_Ω .

As to show the unifying framework based on supervision patterns, we here consider the very particular supervision pattern Ω_f of Section 3.2, originally considered by [12; 13] with the associated notion of f -diagnosability. Let us first recall this notion.

Let G be an LTS which is alive and has no loop of unobservable event. G is f -diagnosable whenever

$$\exists N \in \mathbb{N}, \forall s \in \Sigma^* \cdot f, \forall t \in \mathcal{L}(G)/s, \text{ if } \|t\| \geq N, \text{ then } \forall u \in \mathcal{L}(G), P(u) = P(s.t) \Rightarrow u \in \Sigma^* \cdot f \cdot \Sigma^* \quad (2)$$

The following proposition relates f -diagnosability with Ω_f -diagnosability:

Proposition 2 Let G be an LTS and assume that G is alive and has no loop of internal events. Then G is f -diagnosable if and only if G is Ω_f -diagnosable.

Proof We first make the following remarks:

- (a) $u \in \Sigma^* \cdot f \cdot \Sigma^*$ is equivalent to $u \in \mathcal{L}_{Q_F}(\Omega_f)$;

(b) $s \in \Sigma^*.f$ implies $s \in \mathcal{L}_{Q_F}(\Omega_f)$;

Assume G is f -diagnosable, and that $N \in \mathbb{N}$ fulfills (2). We prove that $\Omega_f(N)$ -diagnosability holds: consider $s \in \mathcal{L}_{Q_F}(\Omega_f) \cap \Sigma^*\Sigma_o$ and let $t \in \mathcal{L}(G)/s \cap \Sigma^*\Sigma_o$ with $\|P(t)\| \geq N$; note that therefore $\|t\| \geq N$. It is easy to show that s decomposes into $s = s'.s''$ where $s' \in \Sigma^*.f$, with additionally $s''.t \in \mathcal{L}(G)/s'$. Now, $\|t\| \geq N$ implies $\|s''.t\| \geq N$, which, by (2), entails that for any $u \in \mathcal{L}(G)$ with $P(u) = P(s.t)$, we have $u \in \Sigma^*.f.\Sigma^* = \mathcal{L}_{Q_F}(\Omega_f)$. This implies in particular that $\llbracket P(s.t) \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega_f)$.

Reciprocally, assume G is $\Omega_f(n)$ -diagnosable, for some n . Let m be the length of the longest unobservable trajectory in G (which exists by assumption), and consider $N = (n + 1) * m$. Consider $s \in \Sigma^*.f$ and $t \in \mathcal{L}(G)/s$ with $\|t\| \geq N$ (thus $\|P(t)\| \geq n + 1$). We have to prove that $u \in \mathcal{L}(G)$ with $P(u) = P(s.t)$, we have $u \in \Sigma^*.f.\Sigma^*$. Let $t = t_1.t_2.t_3$ with $t_1 \in \Sigma_{uo}^*\Sigma_o$, $t_2 \in \Sigma^*\Sigma_o$, and $t_3 \in \Sigma_{uo}^*$. Let $s' = s.t_1$. As Q_F is stable and $s \in \mathcal{L}_{Q_F}(\Omega_f)$, $s' \in \mathcal{L}_{Q_F}(\Omega_f) \cap \Sigma^*\Sigma_o$. We have $t_2 \in \mathcal{L}(G)/s' \cap \Sigma^*\Sigma_o$ with $\|P(t_2)\| \geq n$. By $\Omega_f(n)$ -diagnosability, for all $u \in \mathcal{L}(G)$ with $P(u) = P(s.t)$, we have $u \in \Sigma^*.f.\Sigma^* = \mathcal{L}_{Q_F}(\Omega_f)$. \diamond

4 Algorithms for the Diagnosis Problem

We now propose algorithms for the Diagnosis Problem based on standard operations on LTSs. In a first stage we base the construction of the Diag_Ω function on the synchronous product of G and Ω and its determinisation, and prove that the function Diag_Ω computes a Correct Diagnosis. Next, we propose an algorithm allowing to check for the Ω -diagnosability of an LTS, thus ensuring the Bounded Diagnosis Property of the function Diag_Ω . Hence achieving the decision of the Diagnosis Problem.

4.1 Computing a candidate for the function Diag_Ω

We propose a computation of the function Diag_Ω : given G an LTS and a supervision pattern Ω , we first consider the synchronous product G_Ω of G and Ω (see Definition 2). Next we perform on G_Ω a second operation (see Definition 6) which associates to G_Ω a deterministic LTS written $\text{Det}(G_\Omega)$. We then show how $\text{Det}(G_\Omega)$ provides a function Diag_Ω delivering a Correct Diagnosis.

Let us first introduce a determinisation function.

Definition 6 Let $M = (Q, \Sigma, \rightarrow, q_0)$ be an LTS with $\Sigma = \Sigma_{uo} \cup \Sigma$. The determinization of M is the LTS $\text{Det}(M) = (\mathcal{X}, \Sigma_o, \rightarrow_d, X_0)$ where $\mathcal{X} = 2^Q$ (the set of subsets of Q called macro-states), $X_0 = \{q_0\}$ and $\rightarrow_d = \{(X, \sigma, \Delta_M(X, \Sigma_{uo}^*\sigma) \mid X \in \mathcal{X} \text{ and } \sigma \in \Sigma_o)\}$.

Notice that for this definition the target macro-state X' of a transition $X \xrightarrow{\sigma}_d X'$ is only composed of states q' of M which are targets of sequences of transitions $q \xrightarrow{s} q'$ ending with an observable event σ . The reason for this definition is the coherency with $\llbracket \cdot \rrbracket$. In fact, from the definition of \rightarrow_d in $\text{Det}(M)$, we infer that $\Delta_{\text{Det}(M)}(X_0, \mu) = \{\Delta_M(q_0, \llbracket \mu \rrbracket)\}$, which means that the

macro-state reached from X_0 by μ in $\text{Det}(M)$ is composed of the set of states that are reached from q_0 by trajectories of $\llbracket \mu \rrbracket$ in M .

Finally, determinisation preserves traces, so we have $\mathcal{L}(\text{Det}(M)) = \text{Traces}(\text{Det}(M)) = \text{Traces}(M)$.

We now explain the construction of the diagnoser from G and Ω . Let us first consider the synchronous product $G_\Omega = G \times \Omega$ (see Definition 2). We then get $\mathcal{L}(G_\Omega) = \mathcal{L}(G) \cap \mathcal{L}(\Omega) = \mathcal{L}(G)$ as Ω is complete (thus $\mathcal{L}(\Omega) = \Sigma^*$). We also get $\mathcal{L}(G_\Omega, Q \times Q_F) = \mathcal{L}(G) \cap \mathcal{L}_{Q_F}(\Omega)$ meaning that the trajectories of G accepted by Ω are exactly the accepted trajectories of G_Ω . Finally note that $Q \times Q_F$ is stable in G_Ω as both Q and Q_F are stable by assumption.

We now apply determinisation to G_Ω . We have $\text{Traces}(\text{Det}(G_\Omega)) = \text{Traces}(G_\Omega) = \text{Traces}(G)$ thus for all $\mu \in \text{Traces}(G)$, $\Delta_{\text{Det}(G_\Omega)}(X_0, \mu) = \{\Delta_{G_\Omega}(q_0, \llbracket \mu \rrbracket)\}$.

We now establish the following fundamental results on the construction $\text{Det}(G_\Omega)$:

Proposition 3 For any $\mu \in \text{Traces}(G) = \text{Traces}(G_\Omega)$,

$$\Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \subseteq Q \times Q_F \iff \llbracket \mu \rrbracket \subseteq \mathcal{L}_{Q_F}(\Omega) \quad (3)$$

$$\begin{aligned} \Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \cap Q \times Q_F = \emptyset \\ \iff \llbracket \mu \rrbracket \cap \mathcal{L}_{Q_F}(\Omega) = \emptyset \end{aligned} \quad (4)$$

(3) means that all trajectories compatible with a trace μ are accepted by Ω if and only if μ leads to a macro-state only composed of marked states in G_Ω .

(4) means that all trajectories compatible with μ are not accepted by Ω if and only if μ leads to a macro-state only composed of unmarked states in G_Ω .

Proof The proof of (3) is established by the following sequence of equivalences

$$\begin{aligned} \Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \subseteq Q \times Q_F &\iff \\ \{\Delta_{G_\Omega}(q_0, \llbracket \mu \rrbracket)\} \subseteq Q \times Q_F &\iff \\ \llbracket \mu \rrbracket \subseteq \mathcal{L}_{Q \times Q_F}(G_\Omega) &\iff \\ \llbracket \mu \rrbracket \subseteq \mathcal{L}(G) \cap \mathcal{L}_{Q_F}(\Omega) &\iff \end{aligned}$$

Similarly, for the proof of (4) we have

$$\begin{aligned} \Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \cap Q \times Q_F = \emptyset &\iff \\ \{\Delta_{G_\Omega}(q_0, \llbracket \mu \rrbracket)\} \cap Q \times Q_F = \emptyset &\iff \\ \llbracket \mu \rrbracket \cap \mathcal{L}_{Q \times Q_F}(G_\Omega) = \emptyset &\iff \\ \llbracket \mu \rrbracket \cap \mathcal{L}(G) \cap \mathcal{L}_{Q_F}(\Omega) = \emptyset &\iff \\ \llbracket \mu \rrbracket \cap \mathcal{L}_{Q_F}(\Omega) = \emptyset \text{ as } \llbracket \mu \rrbracket \subseteq \mathcal{L}(G) &\iff \end{aligned}$$

We have now the material to define the function Diag_Ω and to obtain the Correctness Diagnosis Property, following directly from Proposition 3.

Theorem 1 Let $\text{Det}(G_\Omega)$ be the LTS built as above, and let $\text{Diag}_\Omega(\mu)$ be:

$$\left\{ \begin{array}{ll} \text{"YES"}, \text{ if} & \Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \subseteq Q \times Q_F \\ \text{"NO"}, \text{ if} & \Delta_{\text{Det}(G_\Omega)}(X_0, \mu) \cap Q \times Q_F = \emptyset \\ \text{"?"}, \text{ otherwise.} & \end{array} \right. \quad (5)$$

Diag_Ω computes a Correct Diagnosis.

Example 1 In order to illustrate the diagnoser construction, consider the LTS G of Figure 3 (left-hand side). Assume we want to diagnose the occurrence of the fault event f . We thus use the supervision pattern Ω_f described in Figure 1 and build the product $G_\Omega = G \times \Omega_f$. In this case G_Ω is isomorphic to G with set of marked states $\{2, 3, 4\}$. The diagnoser (as well as its answers) obtained by determinization of G_Ω is also represented in Figure 3 (right-hand side).




Figure 3: G and its associated diagnoser w.r.t. Ω_f

4.2 Verifying the Bounded Diagnosis Property of Diag_Ω


As we have established the Correctness of Diag_Ω , according to Proposition 1, the Bounded Diagnosability Property of Diag_Ω is provided by the Ω -diagnosability of G . We now propose an algorithm for deciding Ω -diagnosability (Definition 5).

This algorithm is adapted from [5; 14]. The idea is that G is not Ω -diagnosable if there exists an arbitrarily long trace μ , such that two trajectories compatible with μ disagree on $\mathcal{L}_{Q_F}(\Omega)$ membership (see the above example). We first introduce the Delay-Observational-Closure $\text{OBS}(G_\Omega)$ that preserves the information about $\mathcal{L}_{Q_F}(\Omega)$ membership while abstracting away unobservable events. Next, a self-product $\text{OBS}(G_\Omega) \times \text{OBS}(G_\Omega)$ allows to extract from a trace μ pairs of trajectories of G_Ω and to check their $\mathcal{L}_{Q_F}(\Omega)$ membership agreement.

Definition 7 For an LTS $M = (Q, \Sigma, \rightarrow, q_0)$, the Delay-Observational-Closure of M is the LTS $\text{OBS}(M) = (Q, \Sigma_o, \rightarrow_o, q_0)$ where $q \xrightarrow{\sigma}_o q'$ whenever $q \xrightarrow{s\sigma} q'$ in M for some $s \in \Sigma_{u_o}^*$ and $\sigma \in \Sigma_o$.

By definition, for all $\mu \in \text{Traces}(M)$, $q_0 \xrightarrow{\mu}_o q'$ in $\text{OBS}(G_\Omega)$ if and only $\exists s \in \llbracket \mu \rrbracket$ s.t. $q \xrightarrow{s} q'$ in M .

Consider now $\text{Obs}(G_\Omega) = (Q', \Sigma_o, \rightarrow_o, q_0)$ and let $\Gamma = \text{OBS}(G_\Omega) \times \text{OBS}(G_\Omega)$ be the LTS $(Q' \times Q', \Sigma_o, \rightarrow_\Gamma, (q_0, q_0))$. By definition of OBS and synchronous product, if $\mu \in \text{Traces}(G)$ and $(q_0, q_0) \xrightarrow{\mu}_\Gamma (q, q')$ there exists $s, s' \in \llbracket \mu \rrbracket$ s.t. $q_0 \xrightarrow{s} q$ and $q_0 \xrightarrow{s'} q'$ in G_Ω .



Definition 8 Given Γ defined as above.

- We say that $(q, q') \in Q' \times Q'$ is Ω -determined whenever $q \in Q \times Q_F \iff q' \in Q \times Q_F$. Otherwise, they are called undetermined.

- A path in Γ is called an n -undetermined path if it contains $n + 1$ consecutive Ω -undetermined states (thus n events between them).
- A path in Γ is an undetermined cycle if it is a cycle which states are all undetermined.

We now show the relation between $\Omega(n)$ -diagnosability and the existence of n -undetermined paths.

Lemma 1 There is no reachable n -undetermined path in Γ , if and only if G is $\Omega(n)$ -diagnosable

Proof Suppose there is no reachable n -undetermined path. Let $s \in \mathcal{L}_{Q_F}(\Omega) \cap \Sigma^*. \Sigma_o$ and $t \in \mathcal{L}(G)/s \cap \Sigma^*. \Sigma_o$ with $\|P(t)\| \geq n$. We should prove that for all $u \in \llbracket P(s.t) \rrbracket$, $u \in \mathcal{L}_{Q_F}(\Omega)$. Let $\mu = P(s.t)$. Any path $(q_0, q_0) \xrightarrow{\mu}_\Gamma (q, q')$ in Γ can be decomposed into a path $(q_0, q_0) \xrightarrow{\mu_1}_\Gamma (r, r') \xrightarrow{\mu_2}_\Gamma (q, q')$ with $\mu_1 = P(s)$ and $\mu_2 = P(t)$. We have $\|\mu_2\| = \|P(t)\| \geq n$ thus $(r, r') \xrightarrow{\mu_2}_\Gamma (q, q')$ is a path with n events. By hypothesis, one of these states, say (q_d, q'_d) is determined, and as $s \in \mathcal{L}_{Q_F}(\Omega)$, (q_d, q'_d) is surely in $(Q \times Q_F)^2$. Now as $Q \times Q_F$ is stable, $(q, q') \in (Q \times Q_F)^2$. From this it is clear that for all $u \in \llbracket P(s.t) \rrbracket$, $u \in \mathcal{L}_{Q_F}(\Omega)$.

Conversely, suppose now that there is an n -undetermined path $p = (r, r') \xrightarrow{\mu_2}_\Gamma (q, q')$ in Γ with $\|\mu_2\| = n$ (i.e. all states on the path are undetermined). If this path is reachable there is also a path $(q_0, q_0) \xrightarrow{\mu_1}_\Gamma (r, r')$. As (r, r') is undetermined, there exists $s, s' \in \llbracket \mu_1 \rrbracket$ with $s \in \mathcal{L}_{Q_F}(\Omega) \cap \Sigma^*. \Sigma_o$, and $s' \notin \mathcal{L}_{Q_F}(\Omega)$. There also exists $t \in \mathcal{L}(G)/s \cap \Sigma^*. \Sigma_o$ with $P(t) = \mu_2$. As all states in path p are undetermined, there exists $t' \in \mathcal{L}(G)/s' \cap \Sigma^*. \Sigma_o$ with $P(t') = \mu_2$, but $s'.t' \notin \mathcal{L}_{Q_F}(\Omega)$. We thus have $s \in \mathcal{L}_{Q_F}(\Omega) \cap \Sigma^*. \Sigma_o$ and $t \in \mathcal{L}(G)/s \cap \Sigma^*. \Sigma_o$ with $\|P(t)\| \geq n$, and $s'.t' \in \llbracket P(s.t) \rrbracket$ with $s'.t' \notin \mathcal{L}_{Q_F}(\Omega)$. This proves that G is not $\Omega(n)$ -diagnosable.

Theorem 2 G is Ω -diagnosable if and only if there exists n such that Γ contains no reachable n -undetermined path.

Based on theorem 2 and on the fact that Γ is finite state, we conclude that

Corollary 1 G is not Ω -diagnosable if and only if Γ contains a reachable undetermined cycle.

Using Proposition 1 and the construction of Γ , verifying diagnosability amounts to check the existence of reachable undetermined cycles in Γ , retrieving the idea of the algorithm of [14; 5].

By Corollary 1 and Lemma 1,

Corollary 2 If G is Ω -diagnosable, then G is $\Omega(n + 1)$ -diagnosable, and not $\Omega(n)$ -diagnosable where n is the length of the longest undetermined path of Γ .

We now summarize the procedure to determine whether G is Ω -diagnosable: We perform a depth first search on Γ which either exhibits undetermined cycle or ends by having computed the length of the longest undetermined sequence. Obviously, this has linear cost in the size of Γ .

Example 2 In order to illustrate the construction of Γ , let us come back to the Example 1. $OBS(G_{\Omega_f})$ is given in Figure 4 (left-hand side). The rectangles correspond to the marked states. Now $\Gamma = OBS(G_{\Omega}) \times OBS(G_{\Omega})$ is given in Figure 4 (right-hand side).




Figure 4: $OBS(G_{\Omega})$ and Γ for the LTS of Figure 3

The tuples $\{(1, 3), (3, 1), (1, 4), (4, 1)\}$ in Γ are undetermined. Now it is easy to show that there is no undetermined cycle, which according to Corollary 1 ensures that G is Ω_f -diagnosable. Indeed, as soon as f is triggered, b is observed after the occurrence of a finite number of observable events (bounded by 3). Thus the observation of b surely indicates that f occurred in the past.

A contrario, consider the LTS G' in Figure 5. Γ' given in Figure 6 has undetermined cycles (in $(1, 3)$ and $(3, 1)$), thus, G' is not Ω_f -diagnosable. In fact, for any n , the trajectories $s_1 = a^n$ and $s_2 = f.a^n$ are both compatible with $\mu = a^n$, while $s_1 \notin \mathcal{L}_F(\Omega_f)$, whereas $s_2 \in \mathcal{L}_F(\Omega_f)$.




Figure 5: G' and its associated diagnoser w.r.t. Ω_f




Figure 6: Γ' for the LTS G' of Example 1

Incidentally, this example proves that Ω -diagnosability cannot be checked directly on the diagnoser. In fact the diagnosers for G' and Ω_f (Figure 5, right) and for G for Ω (Figure 3, right) are isomorphic, and G is Ω_f -diagnosable while G' is not. \diamond

5 Supervision Example

The example we discuss here and given in Figure 7 illustrates the approach presented above. In this example,

we simply model the movement of a person in a building composed of an office (I), a library (B), a reception (A) and a coffee-shop (C). The doors from one part of the building to another can be taken in only one direction. Transitions t_i model the crossings of the doors. Some doors are secured by access-cards (possibly allowing the observation). We assume that there exist access-cards




Figure 7: G and its supervision pattern Ω

for the doors t_1 , t_2 and t_4 , meaning that when activated, it is possible to observe the fact that one person crosses the door. We consider the supervision pattern given in Ω (Figure 7) which expresses the fact that going twice to coffee-shop without going to the library is a behaviour that has to be supervised.

Following the different steps described in the previous sections, the product $G_{\Omega} = G \times \Omega$ is used to label the states of G with respect to the supervision pattern Ω . The corresponding LTS is described in Figure 8.




Figure 8: G_{Ω}

Let us first assume that only the access-cards, corresponding to the events t_1 , t_2 are activated and thus observable i.e. $\Sigma_o = \{t_1, t_2\}$. Notice that the system then has internal events loops. The observable system (i.e. $OBS(G_{\Omega})$) is given in Figure 9.




Figure 9: $OBS(G_{\Omega})$ for $\Sigma_o = \{t_1, t_2\}$

It is easy to check that the LTS $\Gamma = OBS(G_{\Omega}) \times OBS(G_{\Omega})$ (not represented here) has an undetermined reachable cycle:

$$\begin{aligned} ((A, N), (A, N)) &\xrightarrow{t_2}_{\Gamma} ((A, N), (A, N1)) \xrightarrow{t_2}_{\Gamma} \\ &((A, N), (A, F)) \xrightarrow{t_2^*}_{\Gamma} ((A, N), (A, F)) \end{aligned}$$

thus G is not Ω -diagnosable when the set of observable events is $\Sigma_o = \{t_1, t_2\}$.

However, if the access-card t_4 is activated (i.e. $\Sigma_o = \{t_1, t_2, t_4\}$), the observable system $\text{OBS}(G_\Omega)$ is given by the LTS represented in Figure 10.




Figure 10: $\text{OBS}(G_\Omega)$ for $\Sigma_o = \{t_1, t_2, t_4\}$

One can check that $\text{OBS}(G_\Omega)$ is deterministic. Thus $\Gamma = \text{OBS}(G_\Omega) \times \text{OBS}(G_\Omega)$ is isomorphic to $\text{OBS}(G_\Omega)$, thus Γ has no undetermined cycle. Consequently, G is Ω -diagnosable for $\Sigma_o = \{t_1, t_2, t_4\}$. Note that we also have that $\text{Det}(G_\Omega) = \text{OBS}(G_\Omega)$. Thus $\text{OBS}(G_\Omega)$ actually corresponds to the diagnoser.

6 Conclusion

The present paper advocates the use of supervision patterns for the description of diagnosis objectives. A supervision pattern is an automaton, like the ones used in many different domains (verification, model-based testing, pattern matching, etc), in order to unambiguously denote a formal language.

As illustrated in the paper, the fault-occurrence diagnosis is a particular case of pattern diagnosis, but patterns are also useful to describe more general objectives, as shown in subsections 3.1 and section 5. The concept of supervision patterns is even more attractive in the sense that patterns can be composed using usual combinators inherited from language theory (union, intersection, concatenation, etc.).

We are interested in diagnosing the occurrence of trajectories violating a safety property, which by definition can be violated on a finite prefix. It is then natural to assume that patterns recognize “extension-closed” languages, in the sense that if a trajectory of the system belongs to the language, so does any extension of this trajectory. This is technically achieved by the *stability* assumption on the automaton Ω . Adopting the behavioral properties point of view on the patterns leads to the attempt to diagnose any linear time pure past temporal formulas [8]. It is clear that the properties we consider do not meet the LTL definable properties handled by [6].

In the worry of exposing a fairly general framework for diagnosis issues, the Diagnosis Problem is presented in a rather denotational spirit, as opposed to the operational spirit we find in the literature: we put the emphasis on the diagnosis function Diag_Ω with its correctness and boundedness diagnosability property. Correctness is an essential property that ensures the accuracy of the diagnosis. Moreover, verifying the diagnosability property of

the system with respect to the supervision pattern guarantees that when using Diag_Ω online, an occurrence of the pattern will eventually be diagnosed, and that this eventuality can be quantified. It is the standard notion of “Diagnosability”, but seen here as a mere mean to achieve a satisfactory diagnosis function; we are aware that this point of view differs from other classical approaches. The definition of diagnosability as proposed here is automata-based, with G and Ω , but could as well be expressed in a language-based framework.

We now turn to technical aspects of the approach. We have insisted on what the semantics of a trace is: a trace denotes the set of trajectories which project onto this trace and that necessarily end up with an observable event. Consequences of this choice are manifold in the definitions of Ω -diagnosability, $\text{Det}(G_\Omega)$ and $\text{OBS}(G)$. We could have chosen another semantics, impacting on the related definitions accordingly: for example we could have considered the set of trajectories which project onto this trace. What is mostly important is the accurate match between the semantics for traces and the other definitions: hence we avoid displeasing discrepancies to determine precisely the Diagnosability Bound, and even better, we have a clear proof for the correctness of the synthesis algorithm. However, we believe our choice is the most natural when admitting that the diagnosis function implemented online as an output verdict is reactive to an observable move of the system.

A more sophisticated diagnosis than the one explained here can be derived from our construction - this is fairly standard: for example, we can take advantage of knowing that the Diagnosability bound is exactly n . Assume that, after a trace μ , the function Diag_Ω produces “?” on $n+2$ consecutive events, then necessarily the trajectories compatible with μ cannot have met the pattern.

We terminate the discussion with future work perspectives aiming two independent objectives. The first objective is to extend the algorithms to more expressive classes of systems, such as infinite systems where data informations is exploited; this would enlarge significantly the applicability of the methods. The second objective is




Figure 11: Example of a non stable supervision pattern

to relax the stability assumption, or equivalently to turn to languages which are not “extension-closed”, intending to encompass frameworks like [2] for intermittent faults, for which a possible supervision pattern would be given by the LTS of Figure 11.

References

- [1] P. Baroni, G. Lamperti, P. Pogliano, and M. Zanel-la. Diagnosis of active systems. In *Proceedings of*

the *European Conference on Artificial Intelligence (ECAI)*, pages 274–278, 1998.

event systems. *IEEE Trans. on Automatic Control*, 47(9):1491–1495, Septembre 2002.

- [2] O. Contant, S. Lafortune, and D. Teneketzis. Diagnosis of intermittent faults. *Discrete Event Dynamic Systems: Theory and Applications*, 14(2):171–202, 2004.
- [3] R. Debouk, S. Lafortune, and D. Teneketzis. Coordinated decentralized protocols for failure diagnosis of discrete-event systems. *Discrete Event Dynamic System : Theory and Applications*, 10(1/2):33–86, Janvier 2000.
- [4] E. Fabre, A. Benveniste, C. Jard, L. Ricker, and M. Smith. Distributed state reconstruction for discrete event systems. In *IEEE Control and Decision Conference (CDC)*, Sydney, 2000.
- [5] S. Jiang, Z. Huang, V. Chandra, and R. Kumar. A polynomial time algorithm for diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 46(8):1318–1321, 2001.
- [6] S. Jiang and R. Kumar. Failure diagnosis of discrete event systems with linear-time temporal logic fault specifications. *IEEE Transactions on Automatic Control*, 49(6):934–945, 2004.
- [7] S. Jiang, R. Kumar, and H.E. Garcia. Diagnosis of repeated/intermittent failures in discrete event systems. *IEEE Transactions on Robotics and Automation*, 19(2):310–323, Avril 2003.
- [8] François Laroussinie and Philippe Schnoebelen. A hierarchy of temporal logics with past. *Theoretical Computer Science*, 148(2):303–324, September 1995.
- [9] Robin Milner. A modal characterisation of observable machine-behaviour. In *CAAP*, pages 25–34, 1981.
- [10] Y. Pencolé and M-O. Cordier. 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):121–170, 2005.
- [11] L. Rozé and M.-O. Cordier. Diagnosing discrete-event systems : an experiment in telecommunication networks. In *4th International Workshop on Discrete Event Systems*, pages 130–137, 1998.
- [12] M. Sampath, R. Sengupta, S. Lafortune, K. Sinaamohideen, and D. Teneketzis. Diagnosability of discrete event systems. *IEEE Transactions on Automatic Control*, 40(9):1555–1575, 1995.
- [13] M. Sampath, R. Sengupta, S. Lafortune, K. Sinaamohideen, and D. Teneketzis. Failure diagnosis using discrete event models. *IEEE Transactions on Control Systems Technology*, 4(2):105–124, Mars 1996.
- [14] T. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially-observed discrete-