

Fault Diagnosis with Dynamic Diagnosers

Franck Cassez, Stavros Tripakis

► **To cite this version:**

Franck Cassez, Stavros Tripakis. Fault Diagnosis with Dynamic Diagnosers. 9th Workshop on Discrete-Event Systems, Jun 2008, Goteborg, Sweden. IEEE Computer Society, pp.212–217, 2008. <inria-00363035>

HAL Id: inria-00363035

<https://hal.inria.fr/inria-00363035>

Submitted on 20 Feb 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fault Diagnosis with Dynamic Observers*

Franck Cassez[†]

Stavros Tripakis[‡]

Abstract—In this paper, we review some recent results about the use of dynamic observers for fault diagnosis of discrete event systems. Dynamic observers can switch sensors on or off, thus dynamically changing the set of events they wish to observe. We study the dynamic diagnoser synthesis problem and some related optimization problems.

I. INTRODUCTION

A. Monitoring, Testing, Fault Diagnosis and Control

Many problems concerning the monitoring, testing, fault diagnosis and control of discrete event systems (DES) can be formalized using finite automata over a set of *observable* events Σ , plus a set of *unobservable* events ([4], [5]). The invisible actions can often be represented by a single unobservable event ε . Given a finite automaton over $\Sigma \cup \{\varepsilon\}$ which is a model of a *plant* (to be monitored, tested, diagnosed or controlled) and an *objective* (good behaviours, what to test for, faulty behaviours, control objective) we want to check if a monitor/tester/diagnoser/controller exists that achieves the objective, and if possible to synthesize one automatically.

The usual assumption in this setting is that the set of observable events is fixed (and this in turn, determines the set of unobservable events as well). Observing an event usually requires some detection mechanism, i.e., a *sensor* of some sort. Which sensors to use, how many of them, and where to place them are some of the design questions that are often difficult to answer, especially without knowing what these sensors are to be used for.

In this paper we review some recent results about *sensor minimization*. These results are interesting since observing an event can be costly in terms of time or energy: computation time must be spent to read and process the information provided by the sensor, and power is required to operate the sensor (as well as perform the computations). It is then essential that the sensors used really provide useful information. It is also important for the computer to discard any information given by a sensor that is not really needed.

Given a fixed set of observable events, it is not the case that all sensors *always* provide useful information and sometimes energy (used for sensor operation and computer treatment) is wasted. For example, to detect a fault f in the system

described by the automaton \mathcal{B} , Figure 1, page 3, an observer needs to watch only for event a initially, and watch for event b only after a has occurred. If the sequence $a.b$ occurs, for sure f has occurred and the observer can raise an alarm. If, on the other hand, event b is not observed after a , f has not occurred. It is then not useful to switch on sensor b before observing event a .

B. Sensor Minimization and Fault Diagnosis

We focus our attention on sensor minimization, without looking at problems related to sensor placement, choosing between different types of sensors, and so on. We also focus on a particular observation problem, that of *fault diagnosis*. We believe, however, that the results we obtain are applicable to other contexts as well.

Fault diagnosis consists in observing a plant and detecting whether a fault has occurred or not. We follow the discrete-event system (DES) setting of [6] where the behavior of the plant is known and a model of it is available as a finite-state automaton over $\Sigma \cup \{\varepsilon, f\}$ where Σ is the set of potentially observable events, ε represents the unobservable events, and f is a special unobservable event that corresponds to the faults¹. Checking *diagnosability* (whether a fault can be detected) for a given plant and a *fixed* set of observable events can be done in polynomial time [6], [7], [8]. In the general case, synthesizing a diagnoser involves determinization and thus cannot be done in polynomial time.

In this paper, we focus on *dynamic* observers. For results about sensor optimization with *static* observers, we refer the reader to [2].

In the dynamic observer framework, we assume that an observer can decide after each new observation the set of events it is going to watch. We first prove that checking diagnosability with dynamic observers that are given by finite automata can be done in polynomial time. As a second aspect, we focus on the *dynamic observer synthesis problem*. We show that computing a *dynamic observer* for a given plant, can be reduced to a *game problem*. We further investigate optimization problems for dynamic observers and define a notion of *cost* of an observer. Finally we show how to compute an optimal (cost-wise) dynamic observer.

C. Related Work

To our knowledge, the problems of synthesizing dynamic observers for diagnosability that we study in this paper have

¹Different types of faults could also be considered, by having different fault events f_1, f_2 , and so on. Our results can be extended in a straightforward way to deal with multiple faults. We restrict our presentation to a single fault event for the sake of simplicity.

*This paper is a digest of our previous work [1], [2] and [3].

[†] CNRS/IRCCyN, 1 rue de la Noë, BP 92101, 44321 Nantes Cedex 3, France. Email: franck.cassez@cnrs.irccyn.fr. Work supported by the French government under grant ANR-SETI-06-003.

[‡] Cadence Research Laboratories, 2150 Shattuck Avenue, 10th floor, Berkeley, CA, 94704, USA, and CNRS, Verimag Laboratory, Centre Equation, 2, avenue de Vignate, 38610 Gières, France. Email: tripakis@cadence.com.

not been addressed in the literature. Due to lack of space, we omit a discussion of previous work on related problems, and refer the reader to [1], [2], [3].

D. Organisation of the paper.

In Section II we fix notation and introduce finite automata with faults to model DES.

In Section III we introduce and study dynamic observers and show that the most permissive dynamic observer can be computed as the strategy in a safety 2-player game.

We also define a notion of cost for dynamic observers in Section IV and show that the cost of a given observer can be computed using Karp's algorithm. Finally, we define the optimal-cost observer synthesis problem and show how it can be solved using Zwick and Paterson's result on graph games.

This paper contains no proofs and the interested reader may refer to [1], [2], [3] for the details.

II. PRELIMINARIES

A. Words and Languages

Let Σ be a finite alphabet and $\Sigma^\varepsilon = \Sigma \cup \{\varepsilon\}$. Σ^* is the set of finite words over Σ and contains ε which is also the empty word and $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$. A *language* L is any subset of Σ^* . Given two words ρ, ρ' we denote $\rho.\rho'$ the concatenation of ρ and ρ' which is defined in the usual way. $|\rho|$ stands for the length of the word ρ (the length of the empty word is zero) and $|\rho|_\lambda$ with $\lambda \in \Sigma$ stands for the number of occurrences of λ in ρ . We also use the notation $|S|$ to denote the cardinality of a set S . Given $\Sigma_1 \subseteq \Sigma$, we define the *projection* operator on words, $\pi_{/\Sigma_1} : \Sigma^* \rightarrow \Sigma_1^*$, recursively as follows: $\pi_{/\Sigma_1}(\varepsilon) = \varepsilon$ and for $a \in \Sigma, \rho \in \Sigma^*$, $\pi_{/\Sigma_1}(a.\rho) = a.\pi_{/\Sigma_1}(\rho)$ if $a \in \Sigma_1$ and $\pi_{/\Sigma_1}(\rho)$ otherwise.

B. Finite Automata

Definition 1 (Finite Automaton) An automaton A is a tuple $(Q, q_0, \Sigma^\varepsilon, \delta)$ with Q a set of states², $q_0 \in Q$ is the initial state, $\delta \subseteq Q \times \Sigma^\varepsilon \times 2^Q$ is the transition relation. We write $q \xrightarrow{\lambda} q'$ if $q' \in \delta(q, \lambda)$. For $q \in Q$, $en(q)$ is the set of actions enabled at q .

If Q is finite, A is a finite automaton. An automaton is deterministic if for any $q \in Q$, $|\delta(q, \varepsilon)| = 0$ and for any $\lambda \neq \varepsilon$, $|\delta(q, \lambda)| \leq 1$. A labeled automaton A is a tuple $(Q, q_0, \Sigma, \delta, L)$ where (Q, q_0, Σ, δ) is an automaton and $L : Q \rightarrow P$ where P is a finite set of observations. ■

A run ρ from state s in A is a finite or infinite sequence of transitions

$$s_0 \xrightarrow{\lambda_1} s_1 \xrightarrow{\lambda_2} s_2 \cdots s_{n-1} \xrightarrow{\lambda_n} s_n \cdots$$

s.t. $\lambda_i \in \Sigma^\varepsilon$ and $s_0 = s$. If ρ is finite and ends in s_n we let $tgt(\rho) = s_n$. The set of finite runs from s in A is denoted $Runs(s, A)$ and we define $Runs(A) = Runs(q_0, A)$. The *trace* of the run ρ , denoted $tr(\rho)$, is the word obtained by concatenating the symbols λ_i appearing in ρ , for those λ_i

different from ε . A word w is *accepted* by A if $w = tr(\rho)$ for some $\rho \in Runs(A)$. The *language* $\mathcal{L}(A)$ of A is the set of words accepted by A .

Let $f \notin \Sigma^\varepsilon$ be a fresh letter that corresponds to the fault action, $\Sigma^{\varepsilon, f} = \Sigma^\varepsilon \cup \{f\}$ and $A = (Q, q_0, \Sigma^{\varepsilon, f}, \delta)$. Given $R \subseteq Runs(A)$, $Tr(R) = \{tr(\rho) \text{ for } \rho \in R\}$ is the set of traces of the runs in R . A run ρ is *k-faulty* if there is some $1 \leq i \leq n$ s.t. $\lambda_i = f$ and $n - i \geq k$. Notice that ρ can be either finite or infinite: if it is infinite, $n = \infty$ and $n - i \geq k$ always holds. $Faulty_{\geq k}(A)$ is the set of *k-faulty* runs of A . A run is *faulty* if it is *k-faulty* for some $k \in \mathbb{N}$ and $Faulty(A)$ denotes the set of faulty runs. It follows that $Faulty_{\geq k+1}(A) \subseteq Faulty_{\geq k}(A) \subseteq \cdots \subseteq Faulty_{\geq 0}(A) = Faulty(A)$. Finally, $NonFaulty(A) = Runs(A) \setminus Faulty(A)$ is the set on *non-faulty* runs of A . We let $Faulty_{\geq k}^{tr}(A) = Tr(Faulty_{\geq k}(A))$ and $NonFaulty^{tr}(A) = Tr(NonFaulty(A))$ be the sets of traces of faulty and non-faulty runs.

We assume that each faulty run of A of length n can be extended into a run of length $n + 1$. This is required for technical reasons (in order to guarantee that the set of faulty runs where sufficient time has elapsed after the fault is well-defined) and can be achieved by adding ε loop-transitions to each deadlock state of A . Notice that this transformation does not change the observations produced by the plant, thus, any observer synthesized for the transformed plant also applies to the original one.

C. Product of Automata

The product of automata with ε -transitions is defined in the usual way: the automata synchronize on common labels except for ε . Let $A_1 = (Q_1, q_0^1, \Sigma_1^\varepsilon, \rightarrow_1)$ and $A_2 = (Q_2, q_0^2, \Sigma_2^\varepsilon, \rightarrow_2)$. The *product* of A_1 and A_2 is the automaton $A_1 \times A_2 = (Q, q_0, \Sigma, \rightarrow)$ where:

- $Q = Q_1 \times Q_2$,
- $q_0 = (q_0^1, q_0^2)$,
- $\Sigma = \Sigma_1 \cup \Sigma_2$,
- $\rightarrow \subseteq Q \times \Sigma \times Q$ is defined by $(q_1, q_2) \xrightarrow{\sigma} (q'_1, q'_2)$ if:
 - either $\sigma \in \Sigma_1 \cap \Sigma_2$ and $q_k \xrightarrow{\sigma} q'_k$, for $k = 1, 2$,
 - or $\sigma \in (\Sigma_i \setminus \Sigma_{3-i}) \cup \{\varepsilon\}$ and $q_i \xrightarrow{\sigma} q'_i$ and $q'_{3-i} = q_{3-i}$, for $i = 1$ or $i = 2$.

III. FAULT DIAGNOSIS WITH DYNAMIC OBSERVERS

In this section we introduce *dynamic observers*. They can choose after each new observation the set of events they are going to watch for. To illustrate why dynamic observers can be useful consider the following example.

Example 1 (Dynamic Observation) Assume we want to detect faults in automaton \mathcal{B} of Figure 1. A static diagnoser that observes $\Sigma = \{a, b\}$ can detect faults. However, no proper subset of Σ can be used to detect faults in \mathcal{B} . Thus the minimum cardinality of the set of observable events for diagnosing \mathcal{B} is 2 i.e., a static observer will have to monitor two events during the execution of the DES. This means that an observer will have to be receptive to at least two inputs at each point in time to detect a fault in \mathcal{B} . One can think of

²In this paper we often use finite automata that generate prefix-closed languages, hence we do not need to use a set of final or accepting states.

being receptive as switching on a device to sense an event. This consumes energy. We can be more efficient using a dynamic observer; that only turns on sensors when needed, thus saving energy: in the beginning we only switch on the a -sensor; once an a occurs the a -sensor is switched off and the b -sensor is switched on. Compared to the previous diagnosers we use half as much energy.




Fig. 1. The automaton \mathcal{B}

A. Dynamic Observers

We formalize the above notion of dynamic observation using *observers*. The choice of the events to observe can depend on the choices the observer has made before and on the observations it has made. Moreover an observer may have *unbounded* memory.

Definition 2 (Observer) An observer Obs over Σ is a deterministic labeled automaton $\text{Obs} = (S, s_0, \Sigma, \delta, L)$, where S is a (possibly infinite) set of states, $s_0 \in S$ is the initial state, Σ is the set of observable events, $\delta : S \times \Sigma \rightarrow S$ is the transition function (a total function), and $L : S \rightarrow 2^\Sigma$ is a labeling function that specifies the set of events that the observer wishes to observe when it is at state s . We require for any state s and any $a \in \Sigma$, if $a \notin L(s)$ then $\delta(s, a) = s$: this means the observer does not change its state when an event it has chosen not to observe occurs. ■

As an observer is deterministic we use the notation $\delta(s_0, w)$ to denote the state s reached after reading the word w and $L(\delta(s_0, w))$ is the set of events Obs observes after w .

An observer implicitly defines a *transducer* that consumes an input event $a \in \Sigma$ and, depending on the current state s , either outputs a (when $a \in L(s)$) and moves to a new state $\delta(s, a)$, or outputs ε , (when $a \notin L(s)$) and remains in the same state waiting for a new event. Thus, an observer defines a mapping Obs from Σ^* to Σ^* (we use the same name “Obs” for the automaton and the mapping). Given a run ρ , $\text{Obs}(\pi_{/\Sigma}(\text{tr}(\rho)))$ is the output of the transducer on ρ . It is called the *observation* of ρ by Obs . We next provide an example of a particular case of observer which can be represented by a finite-state machine.




Fig. 2. A finite-state observer Obs

Example 2 Let Obs be the observer of Figure 2. Obs maps the following inputs as follows: $\text{Obs}(\text{baab}) = ab$,

$\text{Obs}(\text{bababbaab}) = ab$, $\text{Obs}(\text{bbbbba}) = a$ and $\text{Obs}(\text{bbaaa}) = a$. If Obs operates on the DES \mathcal{B} of Figure 1 and \mathcal{B} generates $f.a.b$, Obs will have as input $\pi_{/\Sigma}(f.a.b) = a.b$ with $\Sigma = \{a, b\}$. Consequently the observation of Obs is $\text{Obs}(\pi_{/\Sigma}(f.a.b)) = a.b$.

B. Fault Diagnosis with Dynamic Diagnosers

Definition 3 ((Obs, k)-diagnoser) Let A be a finite automaton over $\Sigma^{\varepsilon, f}$ and Obs be an observer over Σ . $D : \Sigma^* \rightarrow \{0, 1\}$ is an (Obs, k)-diagnoser for A if

- $\forall \rho \in \text{NonFaulty}(A)$, $D(\text{Obs}(\pi_{/\Sigma}(\text{tr}(\rho)))) = 0$ and
- $\forall \rho \in \text{Faulty}_{\geq k}(A)$, $D(\text{Obs}(\pi_{/\Sigma}(\text{tr}(\rho)))) = 1$. ■

A is (Obs, k)-diagnosable if there is an (Obs, k)-diagnoser for A . A is Obs-diagnosable if there is some k such that A is (Obs, k)-diagnosable.

If a diagnoser always selects Σ as the set of observable events, it is a static observer and (Obs, k)-diagnosability amounts to the standard (Σ, k) -diagnosis problem [6].

As for Σ -diagnosability, we have the following equivalence for dynamic observers: A is (Obs, k)-diagnosable iff

$$\text{Obs}(\pi_{/\Sigma}(\text{Faulty}_{\geq k}^{\text{tr}}(A))) \cap \text{Obs}(\pi_{/\Sigma}(\text{NonFaulty}^{\text{tr}}(A))) = \emptyset.$$

Problem 1 (Finite-State Obs-Diagnosability)

INPUT: A , a finite automaton and Obs a finite-state observer.

PROBLEM:

(A) Is A Obs-diagnosable?

(B) If the answer to (A) is “yes”, compute the minimum k such that A is (Obs, k)-diagnosable.

Theorem 1 Problem 1 is in P.

To prove Theorem 1 we build a *product* automaton³ $A \otimes \text{Obs}$ such that: A is (Obs, k)-diagnosable $\iff A \otimes \text{Obs}$ is (Σ, k) -diagnosable. Given two finite automata $A = (Q, q_0, \Sigma^{\varepsilon, f}, \rightarrow)$ and $\text{Obs} = (S, s_0, \Sigma, \delta, L)$, the automaton $A \otimes \text{Obs} = (Q \times S, (q_0, s_0), \Sigma^{\varepsilon, f}, \rightarrow)$ is defined as follows:

- $(q, s) \xrightarrow{\beta} (q', s')$ iff $\exists \lambda \in \Sigma$ s.t. $q \xrightarrow{\lambda} q'$, $s' = \delta(s, \lambda)$ and $\beta = \lambda$ if $\lambda \in L(s)$, $\beta = \varepsilon$ otherwise;
- $(q, s) \xrightarrow{\lambda} (q', s)$ iff $\exists \lambda \in \{\varepsilon, f\}$ s.t. $q \xrightarrow{\lambda} q'$.

The number of states of $A \otimes \text{Obs}$ is at most $|Q| \times |S|$ and the number of transitions is bounded by the number of transitions of A . Hence the size of the product is polynomial in the size of the input $|A| + |\text{Obs}|$. Checking that $A \otimes \text{Obs}$ is diagnosable can be done in polynomial time and Problem 1.(A) is in P.

Example 3 Let \mathcal{B} be the DES given in Figure 1 and Obs the observer of Figure 2. The product $\mathcal{B} \otimes \text{Obs}$ used in the above proof is given in Figure 3.

For Problem 1, we have assumed that an observer was given. It would be even better if we could *synthesize* an observer Obs such that the plant is Obs-diagnosable. Before attempting to synthesize such an observer, we should first

³We use \otimes to clearly distinguish this product from the usual synchronous product \times .




Fig. 3. The product $\mathcal{B} \otimes \text{Obs}$

check that the plant is Σ -diagnosable: if it is not, then obviously no such observer exists; if the plant is Σ -diagnosable, then the trivial observer that observes all events in Σ at all times works⁴. As a first step towards synthesizing non-trivial observers, we can attempt to compute the set of *all* valid observers, which includes the trivial one but also non-trivial ones (if they exist).

Problem 2 (Dynamic-Diagnosability)

INPUT: A .

PROBLEM: Compute the set \mathcal{O} of all observers such that A is *Obs-diagnosable* iff $\text{Obs} \in \mathcal{O}$.

We do not have a solution to the above general problem. Instead, we introduce a restricted variant:

Problem 3 (Dynamic- k -Diagnosability)

INPUT: $A, k \in \mathbb{N}$.

PROBLEM: Compute the set \mathcal{O} of all observers such that A is *(Obs, k)-diagnosable* iff $\text{Obs} \in \mathcal{O}$.

C. Problem 3 as a Game Problem

To solve Problem 3 we reduce it to a *safety* 2-player game. In short, the reduction we propose is the following:

- Player 1 chooses the set of events it wishes to observe, then it hands over to Player 2;
- Player 2 chooses an event and tries to produce a run which is the observation of a k -faulty run and a non-faulty run.

Player 2 wins if he can produce such a run. Otherwise Player 1 wins. Player 2 has complete information of Player 1's moves (i.e., it can observe the sets that Player 1 chooses to observe). Player 1, on the other hand, only has partial information of Player 2's moves because not all events are observable (details follow). Let $A = (Q, q_0, \Sigma^{\varepsilon, f}, \rightarrow)$ be a finite automaton. To define the game, we use two copies of automaton A : A_1^k and A_2 . The accepting states of A_1^k are those corresponding to runs of A which are faulty and where more than k steps occurred after the fault. A_2 is a copy of A where the f -transitions have been removed. The game we are going to play is the following (see Figure 4; Player 1 states are depicted with square boxes and Player 2 states with round shapes):

- 1) the game starts in an state (q_1, q_2) corresponding to the initial state of the product of A_1^k and A_2 . Initially, it is Player 1's turn to play. Player 1 chooses a set of events

he is going to observe i.e., a subset X of Σ and hands it over to Player 2;

- 2) assume the automata A_1^k and A_2 are in states (q_1, q_2) . Player 2 can change the state of A_1^k and A_2 by:
 - a) firing an action (like $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ in Figure 4) which is not in X . This action can be fired either in A_1^k or A_2 (no synchronization). In this case a new state (q, q') is reached and Player 2 can play again from this state;
 - b) firing an action λ in X (like σ_1, σ_2 in Figure 4): to do this both A_1^k and A_2 must be in a state where λ is enabled (synchronization); after the action is fired a new state (q'_1, q'_2) is reached: now it is Player 1's turn to play, and the game continues as in step 1 above from the new state (q'_1, q'_2) .

Player 2 wins if he can reach a state (q_1, q_2) in $A_1^k \times A_2$ where q_1 is an accepting state of A_1^k (this means that Player 1 wins if it can avoid ad infinitum this set of states). In this sense this is a safety game for Player 1 (and a reachability game for Player 2). This game can be defined formally (see [2]), as a game $G_A = (S_1 \uplus S_2, s_0, \Sigma_1 \uplus \Sigma_2, \delta)$. We can show that for any observer O s.t. A is (O, k) -diagnosable, there is a strategy $f(O)$ for Player 1 in G_A s.t. $f(O)$ is *trace-based* and winning. A *strategy* for Player 1 is a mapping $f : \text{Runs}(G_A) \rightarrow \Sigma_1$ that associates a move $f(\rho)$ in Σ_1 to each run ρ in G_A that ends in an S_1 -state. A strategy f is trace-based if given two runs ρ, ρ' , if $\text{tr}(\rho) = \text{tr}(\rho')$ then $f(\rho) = f(\rho')$. Conversely, for any trace-based winning strategy f (for Player 1) in G_A , we can build an observer $O(f)$ s.t. A is $(O(f), k)$ -diagnosable.

Let $O = (S, s_0, \Sigma, \delta, L)$ be an observer for A . We define the strategy $f(O)$ on finite runs of G_A ending in a Player 1 state by: $f(O)(\rho) = L(\delta(s_0, \pi_{/\Sigma}(\text{tr}(\rho))))$. The intuition is that we take the run ρ in G_A , take the trace of ρ (choices of Player 1 and moves of Player 2) and remove the choices of Player 1. This gives a word in Σ^* . The strategy for Player 1 after ρ is the set of events the observer O chooses to observe after reading $\pi_{/\Sigma}(\text{tr}(\rho))$ i.e., $L(\delta(s_0, \pi_{/\Sigma}(\text{tr}(\rho))))$.

Conversely, with each trace-based strategy f of the game G_A we can associate an automaton $O(f) = (S, s_0, \Sigma, \delta, L)$ defined by:

- $S = \{\pi_{/\Sigma}(\text{tr}(\rho)) \mid \rho \in \text{Out}(G_A, f) \text{ and } \text{tgt}(\rho) \in S_1\}$;
- $s_0 = \varepsilon$;
- $\delta(v, \ell) = v'$ if $v \in S, v' = v.\ell$ and there is a run $\rho \in \text{Out}(G_A, f)$ with $\rho = q_0 \xrightarrow{X_0} q_0^1 \xrightarrow{\varepsilon^*} q_0^{n_0} \xrightarrow{\lambda_1} q_1 \xrightarrow{X_1} q_1^1 \xrightarrow{\varepsilon^*} q_1^{n_1} \xrightarrow{\lambda_2} q_2 \cdots q_{k_1} \xrightarrow{\varepsilon^*} q_{k-1}^{n_{k-1}} \xrightarrow{\lambda_k} q_k$ with each $q_i \in S_1, q_i^j \in S_2, v = \pi_{/\Sigma}(\text{tr}(\rho))$, and $\rho \xrightarrow{X_k} q_k^1 \xrightarrow{\varepsilon^*} q_k^{n_k} \xrightarrow{\ell} q_{k+1}$ with $q_{k+1} \in S_1, \ell \in X_k$.
 $\delta(v, \ell) = v$ if $v \in S$ and $\ell \notin f(\rho)$;
- $L(v) = f(\rho)$ if $v = \pi_{/\Sigma}(\text{tr}(\rho))$.

Using the previous definitions and constructions we obtain the following theorems:

Theorem 2 Let O be an observer s.t. A is (O, k) -diagnosable. Then $f(O)$ is a trace-based winning strategy in G_A .

⁴Notice that this also shows that existence of an observer implies existence of a finite-state observer, since the trivial observer is finite-state.




Fig. 4. Game reduction for problem 3

Theorem 3 Let f be a trace-based winning strategy in G_A . Then $O(f)$ is an observer and A is $(O(f), k)$ -diagnosable.

Known results [9] on a game like G_A imply that, if there is a winning trace-based strategy for Player 1, then there is a most permissive strategy \mathcal{F}_A which has finite memory. It can be represented by a finite automaton $\mathcal{S}_{\mathcal{F}_A} = (W_1 \uplus W_2, s_0, \Sigma \cup 2^\Sigma, \Delta_A)$ s.t. $\Delta_A \subseteq (W_1 \times 2^\Sigma \times W_2) \cup (W_2 \times \Sigma \times W_1)$ which has size exponential in the size of G_A . For a given run $\rho \in (\Sigma \cup 2^\Sigma)^*$ ending in a W_1 -state, we have $\mathcal{F}_A(w) = en(\Delta_A(s_0, w))$. We can also prove (Cf. [2]) that:

Theorem 4 \mathcal{F}_A is the most permissive observer.

IV. OPTIMAL DYNAMIC OBSERVERS

In this section we define a notion of *cost* for observers. This will allow us to compare observers w.r.t. to this criterion and later on to synthesize an optimal observer. The notion of cost we are going to use is inspired by *weighted automata*.

A. Cost of a Dynamic Observer

Let $\text{Obs} = (S, s_0, \Sigma, \delta, L)$ be an observer and $A = (Q, q_0, \Sigma^{\varepsilon, f}, \rightarrow)$. We would like to define a notion of *cost* for observers in order to select an optimal one among all of those which are valid, i.e., s.t. A is (Obs, k) -diagnosable. Intuitively this notion of cost should imply that the more events we observe at each step, the more expensive it is.

There is not one way of defining a notion of cost and the reader is referred to [1], [3] for a discussion on this subject.

The cost of a word w is given by:

$$\text{Cost}(w) = \frac{\sum_{i=0}^{i=n} |L(\delta(s_0, w(i)))|}{n+1}$$

with $n = |w|$.

We now show how to define and compute the cost of an observer Obs that observes a DES A .

Given a run $\rho \in \text{Runs}(A)$, the observer only processes $\pi_{/\Sigma}(tr(\rho))$ (ε and f -transitions are not processed). To have a consistent notion of costs that takes into account the logical time elapsed from the beginning, we need to take into account one way or another the number of *steps* of ρ (the length of ρ) even if some of them are non observable. A simple way to do this is to consider that ε and f are now observable events, let's say u , but that the observer never chooses to observe them. Indeed we assume we have already

checked that A is (Obs, k) -diagnosable, and the problem is now to compute the cost of the observer we have used.

Definition 4 (Cost of a Run) Let $\rho = q_0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} q_n$ be in $\text{Runs}(A)$ and $w_i = \text{Obs}(\pi_{/\Sigma}(tr(\rho(i))))$, $0 \leq i \leq n$. The cost of ρ is defined by:

$$\text{Cost}(\rho, A, \text{Obs}) = \frac{1}{n+1} \cdot \sum_{i=0}^n |L(\delta(s_0, w_i))|.$$

■

Let $\text{Runs}^n(A)$ be the set of runs of length n in $\text{Runs}(A)$. The cost of the runs of length n of A is

$$\text{Cost}(n, A, \text{Obs}) = \max_{\rho \in \text{Runs}^n(A)} \{\text{Cost}(\rho, A, \text{Obs})\}.$$

The cost of the pair (Obs, A) is

$$\text{Cost}(A, \text{Obs}) = \limsup_{n \rightarrow \infty} \text{Cost}(n, A, \text{Obs}).$$

Notice that $\text{Cost}(n, A, \text{Obs})$ is defined for each n because we have assumed A generates runs of arbitrary large length.

As emphasised previously, in order to compute $\text{Cost}(n, A, \text{Obs})$ we consider that ε and f are now observable events, say u , but that the observer never chooses to observe them. Let $\text{Obs}^+ = (S, s_0, \Sigma^u, \delta', L)$ where δ' is δ augmented with u -transitions that loop on each state $s \in S$. Let A^+ be A where ε and f transitions are renamed u . Let $A^+ \times \text{Obs}^+$ be the synchronized product of A^+ and Obs^+ . $A^+ \times \text{Obs}^+ = (Z, z_0, \Sigma^u, \Delta)$ is complete w.r.t. Σ^u and we let $w(q, s) = |L(s)|$ so that $(A^+ \times \text{Obs}^+, w)$ is a weighted automaton [10].

Thus we can compute the cost of a given pair (A, Obs) : this can be done using Karp's maximum mean weight cycle algorithm [10] on weighted graphs. This algorithm is polynomial in the size of the weighted graph and thus:

Theorem 5 Computing $\text{Cost}(A, \text{Obs})$ is in P .

Remark 1 Notice that instead of the values $|L(s)|$ we could use any mapping from states of Obs to \mathbb{Z} and consider that the cost of observing $\{a, b\}$ is less than observing a .

B. Optimal Dynamic Diagnosers

We now focus on the problem of computing a best observer in the sense that diagnosing the DES with it has minimal cost. We address the following problem:

Problem 4 (Bounded Cost Observer)

INPUT: A , $k \in \mathbb{N}$ and $c \in \mathbb{N}$.

PROBLEM:

- (A). Is there an observer Obs s.t. A is (Obs, k) -diagnosable and $\text{Cost}(A, \text{Obs}) \leq c$?
 (B). If the answer to (A) is “yes”, compute a witness optimal observer Obs with $\text{Cost}(A, \text{Obs}) \leq c$.

To compute an optimal observer, we use a result by Zwick and Paterson [11] on *weighted graph games*.

To solve Problem 4, we use the most permissive observer \mathcal{F}_A we computed in section III-C. Given A and \mathcal{F}_A , we build a weighted graph game $WG(A, \mathcal{F}_A)$ s.t. the value of the game is the optimal cost for the set of all observers. Moreover an optimal observer can be obtained by taking an optimal memoryless strategy in $WG(A, \mathcal{F}_A)$. By construction of $WG(A, \mathcal{F}_A)$ and the definition of the value of a weighted graph game, the value of the game is the optimal cost for the set of all observers O s.t. A is (O, k) -diagnosable.

Assume A has n states and m transitions. From Theorem 4 we know that \mathcal{F}_A has at most $O(2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}})$ states and $O(2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}} \times n^2 \times k \times m)$ transitions. Hence $G(A, \mathcal{F}_A)$ has at most $O(n \times 2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}})$ vertices and $O(m \times 2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}})$ edges. To make the game complete we may add at most half the number of states and hence $WG(A, \mathcal{F}_A)$ has the same size. We thus obtain the following results:

Theorem 6 *Problem 4 can be solved in time $O(|\Sigma| \times m \times 2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}})$.*

We can even solve the optimal cost computation problem:

Problem 5 (Optimal Cost Observer)

INPUT: A , $k \in \mathbb{N}$.

PROBLEM: Compute the least value m s.t. there exists an observer Obs with $\text{Cost}(A, \text{Obs}) \leq m$.

Theorem 7 *Problem 5 can be solved in time $O(|\Sigma| \times m \times 2^{n^2} \times 2^k \times 2^{2^{|\Sigma|}})$.*

A consequence of Theorem 7 and Zwick and Paterson’s results is that the cost of the optimal observer is a rational number.

V. CONCLUSIONS

In this paper we have reviewed recent results on sensor minimization problems in the context of fault diagnosis, using dynamic observers. We proved that, for an observer given by a finite automaton, diagnosability can be checked in polynomial time (as in the case of static observers). We also solved a synthesis problem of dynamic observers and showed

that a most-permissive dynamic observer can be computed in doubly-exponential time, provided an upper bound on the delay needed to detect a fault is given. Finally we have defined a notion of cost for dynamic observers and shown how to compute the minimal-cost observer that can be used to detect faults within a given delay.

There are several directions we are currently investigating.

Problem 2 has not been solved so far. The major impediment to solve it is that the reduction we propose in section III yields a Büchi game in this case. More generally we plan to extend the framework we have introduced for fault diagnosis to control under dynamic partial observation and this will enable us to solve Problem 2.

Problem 3 is solved in doubly exponential time. Nevertheless to reduce in practice the number of states of the most permissive observer, we point out that only *minimal* sets of events need to be observed. Indeed, if we can diagnose a system by observing only Σ from some point on, we surely can diagnose it using any superset $\Sigma' \supseteq \Sigma$. So far we keep all the sets that can be used to diagnose the system. We could possibly take advantage of the previous property using techniques described in [12].

Acknowledgements: The authors would like to thank Karine Altisen for her contribution in the early stage of this work.

REFERENCES

- [1] F. Cassez, S. Tripakis, and K. Altisen, “Synthesis of optimal dynamic observers for fault diagnosis of discrete-event systems,” in *1st IEEE & IFIP Int. Symp. on Theoretical Aspects of Soft. Engineering (TASE’07)*. IEEE Computer Society, 2007, pp. 316–325.
- [2] —, “Sensor minimization problems with static or dynamic observers for fault diagnosis,” in *7th Int. Conf. on Application of Concurrency to System Design (ACSD’07)*. IEEE Computer Society, 2007, pp. 90–99.
- [3] F. Cassez and S. Tripakis, “Fault diagnosis with static and dynamic observers,” *Fundamenta Informaticae*, 2008, accepted for Publication.
- [4] P. Ramadge and W. Wonham, “Supervisory control of a class of discrete event processes,” *SIAM J. Control Optim.*, vol. 25, no. 1, Jan. 1987.
- [5] J. Tsitsiklis, “On the control of discrete event dynamical systems,” *Mathematics of Control, Signals and Systems*, vol. 2, no. 2, 1989.
- [6] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzi, “Diagnosability of discrete event systems,” *IEEE Transactions on Automatic Control*, vol. 40, no. 9, Sept. 1995.
- [7] T.-S. Yoo and S. Lafortune, “On the computational complexity of some problems arising in partially-observed discrete event systems,” in *American Control Conference (ACC’01)*, 2001, arlington, VA.
- [8] S. Jiang, Z. Huang, V. Chandra, and R. Kumar, “A polynomial algorithm for testing diagnosability of discrete event systems,” *IEEE Transactions on Automatic Control*, vol. 46, no. 8, Aug. 2001.
- [9] W. Thomas, “On the synthesis of strategies in infinite games,” in *Proc. 12th Annual Symposium on Theoretical Aspects of Computer Science (STACS’95)*, vol. 900. Springer, 1995, pp. 1–13, invited talk.
- [10] R. Karp, “A characterization of the minimum mean cycle in a digraph,” *Discrete Mathematics*, vol. 23, pp. 309–311, 1978.
- [11] U. Zwick and M. Paterson, “The complexity of mean payoff games on graphs,” *Theoretical Computer Science*, vol. 158, no. 1–2, pp. 343–359, 1996.
- [12] L. Doyen, K. Chatterjee, T. Henzinger, and J.-F. Raskin, “Algorithms for omega-regular games with imperfect information,” in *Computer Science Logic (CSL)*, ser. Lecture Notes in Computer Science 4207. Springer, 2006, pp. 287–302.