

Complexity reduction techniques for quantified diagnosability of stochastic systems

Hugo Bazille *, Eric Fabre *, Blaise Genest **

* Univ Rennes, INRIA, team SUMO, IRISA

** Univ Rennes, CNRS, team SUMO, IRISA

Abstract: In a discrete event stochastic system, the natural notion of diagnosability, called A-diagnosability, requires that each fault event is eventually detected with probability one. Several definitions of diagnosability degree have been derived from this notion. They examine the detection probability after a fault occurs. To check diagnosability and compute diagnosability degrees, one usually attaches to the original stochastic system the information of a so-called diagnoser, which is in general exponentially larger than the original system. In this paper, we show that the full complexity of such diagnosers is not necessary, and that one can rely on simpler systems, with up to an exponential gain in complexity.

Keywords: discrete event system, stochastic system, diagnosability degree, diagnoser

1. INTRODUCTION

The diagnosis problem for discrete event systems was introduced two decades ago [11]. It consists in detecting the occurrence of a specific event, called a “fault,” given the observations collected along this run. If all fault events can be detected in bounded time after their occurrence, the system is said to be diagnosable. Diagnosability was proved to be decidable in polynomial time [9, 13]. Despite the simplicity of the original setting, the diagnosis problem is important as it is the simplest paradigm of more general observability problems, where one tries to guess some property on runs of a partially observable system. As a matter of fact, numerous extensions and refinements have been explored, to more complex models of discrete event systems (Petri nets, distributed systems, partially known systems, stochastic systems), more complex observation settings (distributed observations, on demand observations, active diagnosis, games...), more complex properties to guess (repairable faults, secrets, fault prediction...).

In this paper, we consider the diagnosis problem for stochastic discrete event systems, modeled as partially observable stochastic automata. Several definitions of diagnosability were introduced [12], like A-diagnosability (the fact that, after a fault, its detection will occur with probability one) and AA-diagnosability (the fact that, after a fault, fault likelihood given observations will converge to one). Even the most natural one (A-diagnosability) revealed to hide surprising phenomena: there exist several non-equivalent definitions, one of them being undecidable, the other one being PSPACE-complete [2].

A related problem of practical importance is to embed the diagnosability analysis into a quantitative setting. This consists in determining “how much” a stochastic system is A-diagnosable, which also characterizes how much of a hidden property (e.g. a secret) an external observer could extract from observations, or conversely how much of a secret leaks out through observations. Several approaches have addressed the problem. For example [10] based on the analysis of the stationary distribution over states of the stochastic system. Alternately, [3] examined the exact analysis of the probability of detecting a fault after

it occurs. The usual method to check A-diagnosability or to compute the diagnosability degree is to combine the stochastic automaton with a (non-probabilistic) diagnoser. The problem with this approach is that diagnosers are in general exponential in the size of the system [8]. In this paper, we propose to use pseudo quantitative diagnosers, weaker than diagnosers, but also with much fewer states (by up to an exponential factor). The set of paths pseudo quantitative diagnosers misdiagnose is a 0-probability set, hence they suffice to compute the diagnosability degree and check for A-diagnosability. In the worst case, our pseudo-diagnosers might still be exponential in the size of the systems, which is obvious as otherwise it would show that $P = PSPACE$, which is highly unlikely. But worst-case complexity may not be the norm [5].

This paper is organized as follows. After recalling the basic definitions of diagnosability analysis for stochastic systems (Sec. 2), we recall some key features of the quantitative analysis, which can be expressed in terms of detection probability, and/or in detection speed. We then examine situations where the computation of these quantities can be much simplified. The first approach (Sec. 3) consists in merging faulty states of the original system, which preserves all diagnosability degrees. The second one (Sec. 4) examines the removal of non-faulty states in the system, without altering the detection probability. The two methods can of course be combined, and we exhibit examples where they lead to exponential gains for computing the diagnosability degree of a system. A long version with omitted proofs can be found in [4].

2. NOTIONS OF DIAGNOSABILITY

2.1 Stochastic automata

A weighted automaton $\mathcal{A} = (S, \Sigma, s_0, w)$ over a semi-ring \mathbb{K} consists of a set of states S , an initial state $s_0 \in S$, a finite alphabet of actions Σ and a weight function $w : S \times \Sigma \times S \rightarrow \mathbb{K}$ that associates a weight to any triple $(s, a, s') \in S \times \Sigma \times S$. A stochastic automaton (also called a labeled Markov Chain) is a weighted automaton with $\mathbb{K} = \mathbb{R}^+$ and satisfying $\forall s \in S, \sum_{(a,s') \in \Sigma \times S} w(s, a, s') = 1$. In a stochastic automaton, transition

$t = (s, a, s')$ exists iff $w(s, a, s') \neq 0$, i.e. it has a positive likelihood. We denote by $s^-(t) = s$ the starting state of transition $t = (s, a, s')$, by $s^+(t) = s'$ its resulting state and by $\sigma(t) = a$ its label (or signature). The *support* of \mathcal{A} is the ordinary automaton denoted $\mathcal{A} = (S, \Sigma, s_0, T)$ where the transition set $T \subseteq S \times \Sigma \times S$ is the support of w . A path in \mathcal{A} (and by extension in \mathcal{A}) is a sequence $\pi = t_1 \dots t_n$ of transitions such that for all i in $\{1, 2, \dots, n-1\}$, $s^+(t_i) = s^-(t_{i+1})$. The length of π is denoted $|\pi|$ and is equal to the number of transitions in π . A path π' is a prefix of π iff there exists π'' such that $\pi = \pi'\pi''$ (and π'' is a suffix). Operators s^-, s^+, σ and w naturally extend to paths by $s^-(\pi) = s^-(t_1)$, $s^+(\pi) = s^+(t_n)$, $\sigma(\pi) = \sigma(t_1) \dots \sigma(t_n)$ and $w(\pi) = \prod_{1 \leq i \leq n} w(t_i)$. A run of \mathcal{A} is a path π such that $s^-(\pi) = s_0$. We denote respectively $\mathcal{P}(\mathcal{A})$, $\mathcal{R}(\mathcal{A})$ and $\mathcal{L}(\mathcal{A}) = \{\sigma(\pi) : \pi \in \mathcal{R}(\mathcal{A})\}$ the set of paths of \mathcal{A} , the set of runs of \mathcal{A} and the language of \mathcal{A} . These notions naturally extend to infinite sequences and the according sets are denoted $\mathcal{P}^\infty(\mathcal{A})$, $\mathcal{R}^\infty(\mathcal{A})$ and $\mathcal{L}^\infty(\mathcal{A})$.

In a stochastic automaton \mathcal{A} , let $\pi \in \mathcal{R}(\mathcal{A})$ be a run of \mathcal{A} such that $|\pi| = n$. We denote by $\text{Cyl}(\pi) \subseteq \mathcal{R}^\infty(\mathcal{A})$ the set of all infinite runs of \mathcal{A} that admit π as a prefix. In the set of infinite runs of \mathcal{A} , let \mathcal{C}_n be the sigma-field generated by $\{\text{Cyl}(\pi) : \pi \in \mathcal{R}(\mathcal{A}), |\pi| = n\}$, the set of cylinders generated by runs of length n , and let \mathbb{P}_n be the probability distribution over \mathcal{C}_n generated by the $\mathbb{P}_n(\pi) = w(\pi)$. Then $(\mathcal{C}_n, \mathbb{P}_n)_{n \geq 0}$ forms a projective family, i.e. each \mathbb{P}_{n+m} restricted to \mathcal{C}_n coincides with \mathbb{P}_n . By Kolmogorov's extension theorem, this results in a unique probability space $(\mathcal{C}, \mathbb{P})$ over $\mathcal{R}^\infty(\mathcal{A})$, and \mathbb{P} coincides with \mathbb{P}_n on cylinders of \mathcal{C}_n . This is the probability distribution we consider in the sequel, and we write $\mathbb{P}(\pi)$ instead of $\mathbb{P}(\text{Cyl}(\pi))$ for short. Notice that \mathbb{P} is additive on finite runs (=cylinders), i.e. $\mathbb{P}(\{\pi, \pi'\}) = \mathbb{P}(\pi) + \mathbb{P}(\pi')$, for disjoint cylinders, that is for π, π' not prefixes of one another.

2.2 Diagnosability

Given automaton \mathcal{A} , we partition the set of states into *faulty* and *non-faulty* states: $S = S_F \uplus S_N$. In the examples, the faulty states are the double circles. We also assume that the set of faulty states is absorbing: there are no transitions in \mathcal{A} from S_F to S_N , i.e. faults cannot be repaired. This setting is equivalent to the more usual one where faults are attached to specific transitions. Traditionally, the action alphabet Σ is partitioned into observable and unobservable letters (and faults are generally unobservable transitions). Without loss of generality, we assume that all actions are observable, as any property stated about runs of \mathcal{A} could be equivalently expressed on runs of the non-deterministic epsilon-reduction of \mathcal{A} , which is totally observable [6]: the observation from a path π is its label $\sigma(\pi)$.

We are interested in the standard notion of *diagnosability* for (non-stochastic) automata [11], in its natural extension to stochastic systems, named the *A-diagnosability* [2, 12], and in the definition of more precise *degrees of A-diagnosability* [3, 10], which is 1 iff \mathcal{A} is A-diagnosable. We briefly define these notions hereafter.

Let $\pi \in \mathcal{R}(\mathcal{A})$ be a run of \mathcal{A} . π is *faulty* iff $s^+(\pi) \in S_F$, otherwise π is said to be *non-faulty*. π is said to be *faulty ambiguous* iff it is faulty and there exists π' such that $\sigma(\pi) = \sigma(\pi')$ and π' is non-faulty. Let us denote by $\mathcal{R}_N(\mathcal{A})$ and $\mathcal{R}_F(\mathcal{A})$ the sets of non-faulty and faulty runs of \mathcal{A} . They form the partition $\mathcal{R}(\mathcal{A}) = \mathcal{R}_N(\mathcal{A}) \uplus \mathcal{R}_F(\mathcal{A})$. By extension, let us define by $\mathcal{L}_N(\mathcal{A}) = \sigma(\mathcal{R}_N(\mathcal{A}))$ and $\mathcal{L}_F(\mathcal{A}) = \sigma(\mathcal{R}_F(\mathcal{A}))$

the non-faulty and faulty languages of \mathcal{A} . One has $\mathcal{L}(\mathcal{A}) = \mathcal{L}_N(\mathcal{A}) \cup \mathcal{L}_F(\mathcal{A})$, but this is generally not a partition.

Let $o \in \mathcal{L}(\mathcal{A})$. *Diagnosing* observation sequence o means determining if $\sigma^{-1}(o) \subseteq \mathcal{R}_N(\mathcal{A})$, or if $\sigma^{-1}(o) \subseteq \mathcal{R}_F(\mathcal{A})$, or if none of the above holds. In the first case, one has $o \in \mathcal{L}_N(\mathcal{A}) \setminus \mathcal{L}_F(\mathcal{A})$, and o is declared *non-faulty*, denoted $D(o) = N$: the fault did not occur in π . In the second case, one has $o \in \mathcal{L}_F(\mathcal{A}) \setminus \mathcal{L}_N(\mathcal{A})$, and o is declared *faulty*, denoted $D(o) = F$: the fault did occur for sure in π . In the last case, $o \in \mathcal{L}_N(\mathcal{A}) \cap \mathcal{L}_F(\mathcal{A})$, and o is declared *ambiguous*, denoted $D(o) = A$.

A faulty run $\pi \in \mathcal{R}_F(\mathcal{A})$ is *k-diagnosable* iff $\forall \pi\pi' \in \mathcal{R}(\mathcal{A}), |\pi'| \geq k \Rightarrow D(\sigma(\pi\pi')) = F$. It means that after at most k observations after π one knows for sure that a fault has occurred. A run π is *diagnosable* if there exists k such that π is k -diagnosable. \mathcal{A} is (k-)diagnosable if all faulty runs of \mathcal{A} are (k-)diagnosable. It is well known [6] that for a finite system \mathcal{A} , \mathcal{A} is diagnosable iff there exists k such that \mathcal{A} is k -diagnosable. The diagnosability of an automaton can be characterized in polynomial time [6, 9, 13] using the *twin-machine* $\mathcal{T}_{\mathcal{A}} = \mathcal{A} \times \mathcal{A}'_N$ derived from \mathcal{A} , where \mathcal{A}'_N is the restriction of \mathcal{A} to the non-faulty states of \mathcal{A} and \times is the synchronous product of automata, see Section 4.

2.3 A-Diagnosability

These definitions being purely structural and involving no probability, they can also be applied to stochastic automata. However, this is not totally satisfactory as it is possible that a stochastic system is non-diagnosable while every faulty run will eventually be diagnosed with probability 1. The notion of *A-diagnosability* has been proposed in [12]. Formally, a faulty run π of \mathcal{A} is A-diagnosable iff the probability of producing an infinite ambiguous extension of π is 0. Equivalently, let π' be an extension of length k of π : $\pi\pi' \in \mathcal{R}(\mathcal{A}), |\pi'| = k$, and consider the random variable $D_k = D(\sigma(\pi\pi'))$. D_k is a binary variable, that can only assume value A or F. When it takes value F, it corresponds to the detection of the fault. Then a faulty run π of \mathcal{A} is A-diagnosable iff the series of random variables $(D_k)_{k \geq 0}$ converges to F with probability 1.

Several notions of A-diagnosability were defined: the uniform one asking that $D_k(\pi)$ converges towards 0 uniformly over all faulty path π , and the non-uniform one not requiring it. It was recently proved that the uniform diagnosability is undecidable, while the non-uniform one is decidable, and more precisely PSPACE-complete [2]. In this paper, we focus on the non-uniform A-diagnosability i.e. given an automaton \mathcal{A} , \mathcal{A} is A-diagnosable iff $\mathbb{P}(\pi|\pi \in \mathcal{L}_F^\infty(\mathcal{A}), \exists \pi' \in \mathcal{L}(\mathcal{A}), \sigma(\pi) = \sigma(\pi')) = 0$. We then recall one structural characterization of it relying on the notion of diagnoser [8, 11].

A *diagnoser* for \mathcal{A} is a pair (\mathcal{D}, ϕ) formed by a deterministic automaton $\mathcal{D} = (Q, \Sigma, q_0, T_{\mathcal{D}})$ over the same alphabet as \mathcal{A} , such that $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{D})$, and a labeling function $\phi : Q \rightarrow \{N, F, A\}$ satisfying: for all observed sequence $o \in \sigma(\mathcal{R}(\mathcal{A}))$, denoting by $q^+(o)$ the unique state reached in \mathcal{D} by reading word o from the initial state q_0 , one has $D(o) = \phi(q^+(o))$. Let $\mathcal{D} = \text{Det}(\mathcal{A})$ be the determinized version of \mathcal{A} (the support of \mathcal{A}) obtained by the classical subset construction. So $Q = 2^S$, $q_0 = \{s_0\}$, and for $X \subseteq S$ let $\phi(X) = F$ (resp. N) when $X \subseteq S_F$ (resp. $X \subseteq S_N$), and let $\phi(X) = A$ otherwise. The pair (\mathcal{D}, ϕ) is a diagnoser for \mathcal{A} . In the worst case, the smallest diagnoser of an automaton \mathcal{A} is exponential in $|\mathcal{A}|$ [8].

If the faulty run $\pi \in \mathcal{R}_F(\mathcal{A})$ is k -diagnosable in \mathcal{A} , then denoting $o = \sigma(\pi)$ its observed sequence, every path in \mathcal{D} from $q^+(o)$ of length k must lead to a state q' labeled $\phi(q') = F$. As \mathcal{D} is finite, the above diagnoser construction proves the existence of a uniform maximal delay for detecting a fault when \mathcal{A} is diagnosable. When \mathcal{A} is non-diagnosable, a diagnoser (\mathcal{D}, ϕ) of \mathcal{A} necessarily exhibits an ambiguous cycle accessible after some faulty run of \mathcal{A} .

The (non-uniform) A-diagnosability of a stochastic system \mathcal{A} can in turn be characterized in a similar manner [2]. Consider the synchronous product $\mathcal{M} = \mathcal{A} \times \mathcal{D}$ of \mathcal{A} with one of its diagnosers \mathcal{D} , called the *quantified diagnoser*. \mathcal{M} is a stochastic automaton, because \mathcal{D} is deterministic with a larger language than $\mathcal{L}(\mathcal{A})$. Moreover, there is a one to one correspondence between runs of \mathcal{A} and runs of \mathcal{M} , that preserves likelihoods. Through the labeling ϕ , this construction attaches the diagnosis signal N, A or F to states of \mathcal{M} and thus to runs of \mathcal{A} . Notice that along a run of \mathcal{M} (and thus of \mathcal{A}) the labeling of the successive states can only change from N to A or F , and from A to F . Therefore, state labeling ϕ is constant in each strongly connected component of \mathcal{M} . We can then associate each strongly connected component in \mathcal{M} with either non-faulty, ambiguous, or faulty. System \mathcal{A} is A-diagnosable iff there is no ambiguous bottom strongly connected component of \mathcal{M} that is accessible from a state with a faulty first component [2].

2.4 Diagnosability degrees

More than answering to a yes-no question, we can quantify *how diagnosable* a stochastic system is. Several *diagnosability degrees* can be defined [3, 10] and can especially focus on the diagnosability in bounded time or on the probability to eventually detect a fault. In this article, we will focus on the latter, *i.e.* the volume of faulty runs that can be diagnosed in (unbounded) finite time. This degree is defined by $\Delta_\infty(\mathcal{A}) = 1 - \mathbb{P}(\pi \in \mathcal{R}_A^\infty(\mathcal{A}) | \pi \in \mathcal{R}_F^\infty(\mathcal{A}))$. Notice that \mathcal{A} is A-diagnosable iff Δ_∞ is equal to 0.




Fig. 1. An automaton \mathcal{A}

In figure 1, \mathcal{A} is not diagnosable. For the sake of simplicity, we assume the equiprobability of all transitions going out of the same state. In \mathcal{A} , a faulty run will be diagnosed when it produces a b , thus when it reaches state 3. Some faulty runs will reach state 5 and will not be diagnosed. Hence, the probability that a faulty run will be diagnosed (*i.e.* it reaches state 3) is $7/8$.

3. DIAGNOSER COMPRESSION BY MERGING FAULTY STATES

As recalled in the previous section, the construction of a diagnoser \mathcal{D} is the bottleneck to handle quantified diagnosability on stochastic automata, either checking A-diagnosability or computing diagnosability degree. A traditional choice for \mathcal{D} is

$\mathcal{D} = \text{Det}(\mathcal{A})$, the determinized version of the support of \mathcal{A} , which incurs an exponential state explosion in the worst case. We show in the following that smaller structure works as well.

Definition 1. Let $\mathcal{A} = (S, \Sigma, s_0, T)$ be the support of \mathcal{A} . Assume without loss of generality that $s_0 \notin S_F$. The merged system $\bar{\mathcal{A}}$ is defined as $\bar{\mathcal{A}} = (\bar{S}, \Sigma, s_0, \bar{T})$, where $\bar{S} = S_N \cup \{\bar{s}_F\}$ preserves all non-faulty states of \mathcal{A} but merges all faulty states into a single one, \bar{s}_F . Transitions follow accordingly:

- $\forall (s, \alpha, s') \in S_N \times \Sigma \times S_N, [(s, \alpha, s') \in T \Leftrightarrow (s, \alpha, s') \in \bar{T}]$
- $\forall (s, \alpha) \in S_N \times \Sigma, [(s, \alpha, \bar{s}_F) \in \bar{T} \Leftrightarrow \exists s' \in S_F, (s, \alpha, s') \in T]$
- $\forall \alpha \in \Sigma, (s_F, \alpha, \bar{s}_F) \in \bar{T}$




Fig. 2. Automaton \mathcal{A} (top) and its merged version $\bar{\mathcal{A}}$ (below). While \mathcal{A} is 2-diagnosable, $\bar{\mathcal{A}}$ is not diagnosable.

When \mathcal{A} is diagnosable, the merged system $\bar{\mathcal{A}}$ may lose diagnosability, as shown by the example in Fig. 2. It suffices to consider the observation sequence $a(ba)^\infty$ that can both be produced by an infinite faulty path and by an infinite non-faulty path. Observe that $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\bar{\mathcal{A}})$ as well, and specifically $\mathcal{L}_N(\mathcal{A}) = \mathcal{L}_N(\bar{\mathcal{A}})$ while $\mathcal{L}_F(\mathcal{A}) \subseteq \mathcal{L}_F(\bar{\mathcal{A}})$. In other words, merging faulty states can only introduce extra faulty words, but preserves the non-faulty language. Further, from the first item defining \bar{T} in Def. 1, one can easily check that there exists a one to one correspondence between non-faulty runs of \mathcal{A} and non-faulty runs of $\bar{\mathcal{A}}$.

Consider now the natural diagnoser $(\bar{\mathcal{D}}, \bar{\phi})$ of $\bar{\mathcal{A}}$, where $\bar{\mathcal{D}} = \text{Det}(\bar{\mathcal{A}})$. We claim that although $\bar{\mathcal{D}}$ does not detect all faults in $\bar{\mathcal{A}}$, $(\bar{\mathcal{D}}, \bar{\phi})$ is as good as (\mathcal{D}, ϕ) for fault detection on observations produced by \mathcal{A} .

Theorem 1. Let $\mathcal{M} = \mathcal{A} \times \mathcal{D}$ and $\mathcal{M}' = \bar{\mathcal{A}} \times \bar{\mathcal{D}}$. Then the diagnosability, A-diagnosability and diagnosability degree are the same in \mathcal{M} and in \mathcal{M}' .

Proof:[Sketch of.] Let π be a faulty run of \mathcal{A} , and o its observation. Now $\mathcal{D}(o) = A$ iff there exists a non-faulty run π' of \mathcal{A} with the same observation o . As $\bar{\mathcal{A}}$ and \mathcal{A} have the same set of non-faulty runs, this means iff $\bar{\mathcal{D}}(o) = A$. It also implies that $\mathcal{D}(o) = F$ iff $\bar{\mathcal{D}}(o) = F$. \square

It is quite surprising that the diagnoser of a possibly non-diagnosable system $\bar{\mathcal{A}}$ could perform as well as the diagnoser of \mathcal{A} for detecting faults. This readily suggests potential complexity gains for some systems, as $\bar{\mathcal{D}}$ could be exponentially smaller than \mathcal{D} . This is illustrated by Fig. 3. Let \mathcal{A} be the automaton at the top of the figure. Any deterministic automaton accepting the same language as \mathcal{A} (and in particular the traditional \mathcal{D} obtained by subset construction) is known to be exponential in n , as it needs to remember the last n letters seen.




Fig. 3. Automaton \mathcal{A} (top), its merged version $\tilde{\mathcal{A}}$ (center), and the diagnoser \mathcal{D} of \mathcal{A} (bottom).

This result in a machine \mathcal{M} of size 2^n , much larger than \mathcal{A} . By contrast, the merged version $\tilde{\mathcal{A}}$ of \mathcal{A} (center), although non-diagnosable, results in a compressed diagnoser $\tilde{\mathcal{D}}$ (bottom) of size 3. $\tilde{\mathcal{D}}$ performs as well as \mathcal{D} for diagnosing faulty runs of \mathcal{A} . It is much faster to check A-diagnosability and to compute diagnosability degrees using machine \mathcal{M}' instead of \mathcal{M} .

The next section shows that some reduction is also possible on the non-faulty part of \mathcal{A} .

4. REDUCING THE NUMBER OF NON-FAULTY STATES.

This section presents another complexity reduction technique for quantified diagnosability. Specifically, we show that some non-faulty states of \mathcal{A} can be ignored when building a diagnoser for \mathcal{A} , without changing the diagnosability status of \mathcal{A} or altering the diagnosability degree. This method is compatible with the previous one, that aggregates the faulty states of \mathcal{A} .

4.1 Twin machine

When considering an ordinary automaton $\mathcal{A} = (S, \Sigma, s_0, T)$, diagnosability can be decided in polynomial time rather than PSPACE for A-diagnosability. To obtain polynomial time complexity, instead of using the exponential size diagnoser, one can use the twin machine $Twin(\mathcal{A}) = (S \times S_N, \Sigma, (s_0, s_0), T')$ [6, 9, 13], with a quadratic number of states $|S \times S_N|$ in $|S|$, and where T' is defined as follows:

$$((s, t), a, (s', t')) \in T' \text{ iff } (s, a, s') \in T \wedge (t, a, t') \in T$$

Notice that the twin machine restricts the second component to be non-faulty, avoiding some redundancy (the state space would be symmetric otherwise). Using the twin machine, one can check diagnosability in polynomial time, using the following well-known result [6, 9, 13]:

Proposition 2. Let \mathcal{A} be a (non-stochastic) automaton. Then \mathcal{A} is not diagnosable iff there exists a reachable loop ρ in the twin machine $Twin(\mathcal{A})$ such that every pair of states $(f, s) \in \rho$ satisfies $(f, s) \in S_F \times S_N$.

Such a loop ρ , or more generally a path π in $Twin(\mathcal{A})$ going only through states of $S_F \times S_N$, is called an ambiguous loop/path. Its restriction to \mathcal{A} , the first component of the twin machine, is thus a faulty ambiguous loop/path of \mathcal{A} .

We now show that Prop.2 can be used to reduce the complexity of computing the A-diagnosability degree of a stochastic automaton \mathcal{A} . However, recall that deciding the A-diagnosability of \mathcal{A} is PSPACE-complete, so one readily knows that such a complexity reduction cannot be systematic. Nevertheless, the proposed approach is general enough to suggest that complexity gains can often be expected, sometimes with an exponential reduction, as shown in the example of Fig. 4.

4.2 General idea

Let \mathcal{A} be a stochastic automaton, with support \mathcal{A} . Our idea is the following: the purpose of building a quantified diagnoser $\mathcal{M} = \mathcal{A} \times \mathcal{D}(\mathcal{A})$ is to attach to each state s of \mathcal{A} the signal indicating whether the current state estimate q in $\mathcal{D}(\mathcal{A})$ given past observations is non-faulty, faulty or ambiguous. We are mostly interested in pairs $(s, q) \in S \times Q$ where $s \in S_F$ and q is ambiguous, and further in checking whether this ambiguity will last forever with a positive probability. As $q \subseteq S$, the ambiguity of q comes from the existence of a non-faulty state $t \in q$. Using the twin machine, one can easily check whether the ambiguity due to pair $(s, t) \in S_F \times S_N$ can persist forever (and with positive probability), or will vanish in the future and not prevent fault detection. If the ambiguity due to t will for sure vanish, one needs not take it into account in state (s, q) of \mathcal{M} , and may replace (s, q) by (s, q') where $q' = q \setminus \{t\}$. In doing so, one anticipates on the disappearing of an irrelevant ambiguity source due to t . In other words, one may anticipate a fault detection that will take place for sure. So the diagnosis probability does not change, but the detection delay may be shortened.

Let us now focus on the characterization of pairs of states $(s, t) \in S \times S_N$ that can be safely discarded without changing the A-diagnosability degree. Let $\mathcal{P}^\infty(\mathcal{A}, s)$ denote infinite paths of \mathcal{A} starting from state s , and similarly $\mathcal{P}_F^\infty(\mathcal{A}, s)$, $\mathcal{P}_N^\infty(\mathcal{A}, s)$ for faulty and non-faulty paths.

Definition 3. Given an automaton \mathcal{A} , the state pair $(s, t) \in S \times S_N$ is a negligible pair of type 1 iff there is no pair of infinite runs $\rho \in \mathcal{P}_F^\infty(\mathcal{A}, s)$, $\rho' \in \mathcal{P}_N^\infty(\mathcal{A}, t)$ with $\sigma(\rho) = \sigma(\rho')$. We denote by NE_1 the set of negligible pairs of type 1.

From such pairs, the ambiguity that may hold at state (s, t) or that may appear after state (s, t) will vanish for sure in the future. Notice that we do not require s to be faulty.

One can also ignore pairs of states $(s, t) \in S \times S_N$ for which any ambiguity that may appear in the future will later vanish with probability 1 in \mathcal{A} .

Definition 4. Given a stochastic automaton \mathcal{A} , the state pair $(s, t) \in S \times S_N$ is a negligible pair of type 2 iff $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A}, s) : \exists \rho' \in \mathcal{P}_N^\infty(\mathcal{A}, t), \sigma(\rho) = \sigma(\rho')] = 0$. We denote by NE_2 the set of negligible pairs of type 2.

The above probability is computed over trajectories of stochastic automaton \mathcal{A} , and the involved set of runs can be shown to be measurable. We have trivially $NE_1 \subseteq NE_2$. Characterizing pairs of states in NE_2 algorithmically is clearly more difficult than for NE_1 (which only require to consider the twin machine), as this is where the complexity of checking A-diagnosability comes into the picture. Sufficient conditions can be derived that capture most of such pairs, as we show in the next subsection.

Let us define $NE(s) = \{t \mid (s,t) \in NE_2\}$. Consider now the classical quantified diagnoser $\mathcal{M} = \mathcal{A} \times Det(\mathcal{A})$, and assume this machine is in state $(s,q) \in S \times Q$ after some observed sequence $o \in \Sigma^*$, with $t \in q \cap S_N$. Assume pair $(s,t) \in NE_2$. Then t could be safely removed from q without changing the diagnosability degree: the part of ambiguity due to pair (s,t) in (s,q) will almost surely vanish in the future (with probability 1). Thus, it cannot lead to an ambiguous cycle of positive likelihood.

“Removing” such negligible pairs (s,t) from machine \mathcal{M} can be done in several ways. Either abruptly, by replacing each state (s,q) of \mathcal{M} by pairs $(s,q \setminus NE(s))$. Or better, by recursively synchronizing \mathcal{A} with a constrained state estimator, which gives a smaller stochastic automaton: let (s,q) be a state of \mathcal{M}' , such that $q \cap NE(s) = \emptyset$, if (s,a,s') exists in \mathcal{A} , then add transition $((s,q), a, (s',q'))$ to \mathcal{M}' where $q' = \{t : \exists (s,a,t) \text{ in } \mathcal{A}\} \setminus NE(s')$. This recursive construction starts with initial state (s_0, s_0) . The machine \mathcal{M}' obtained in that way is called the *pseudo quantified diagnoser* of \mathcal{A} . Notice that \mathcal{M}' is a well defined stochastic automaton, just like \mathcal{M} , and that there is a one to one correspondence between runs of \mathcal{A} and runs of \mathcal{M}' that preserves likelihood.

We want to compute $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A}) : \nexists \rho' \in \mathcal{P}_N^\infty(\mathcal{A}), \sigma(\rho) = \sigma(\rho')]$ to get the diagnosability degree (by dividing by $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A})]$ which is easy to compute), and also check whether \mathcal{A} is A-diagnosable (iff the degree is 1) [3]. Using the usual quantified diagnoser \mathcal{M} , we have that $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A}) : \nexists \rho' \in \mathcal{P}_N^\infty(\mathcal{A}), \sigma(\rho) = \sigma(\rho')]$ is the probability to reach states of \mathcal{M} labeled $F(\text{aulty})$ [3]. We denote by B the set of Faulty states (s,q) for \mathcal{M} , that is the set of states (s,q) s.t. $q \in S_F$. Computing the probability to reach a set of states can be made in polynomial time in the size of the machine, here $|\mathcal{M}|$. We now show that the probability to reach B in \mathcal{M} can also be computed as the probability to reach a set of state B' in \mathcal{M}' . This gives us a faster algorithm to check A-diagnosability or compute the degree of diagnosability as \mathcal{M}' is smaller than \mathcal{M} (up to an exponential factor as shown in the example of the next section). We set B' to be the set of states $(s,q) \in S_F \times 2^{S_F}$.

Lemma 5. The probability to reach states B' in \mathcal{M}' is $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A}) : \nexists \rho' \in \mathcal{P}_N^\infty(\mathcal{A}), \sigma(\rho) = \sigma(\rho')]$.

4.3 Negligible pairs in the twin machine

We now explain how to compute a set $NE \subseteq NE_2$ of negligible pairs of states. We first compute the strongly connected components C_1, \dots, C_k of the twin machine $Twin(\mathcal{A})$ using Tarjan’s algorithm, in linear time in the number of states of the twin machine. Remember that the number of states of the twin machine is quadratic at most in the number of states of \mathcal{A} .

We label a strongly connected component of $Twin(\mathcal{A})$ as ambiguous if it contains some state in $S_F \times S_N$. Notice that in this case, as faulty state remains faulty and the second component of $Twin(\mathcal{A})$ is in S_N , the states reachable from a state in $S_F \times S_N$ are also in $S_F \times S_N$.

We can then characterize the set NE_1 of negligible states of type 1 as the set of states of the twin machine which cannot reach any ambiguous SCCs. This can be done in time linear in the number of states of the twin machine, by considering first bottom strongly connected components and then inductively considering components C_i which can reach only components C_j already considered.

Lemma 6. NE_1 is the set of states (s,t) of the twin machine which cannot reach a loop around a state (x,y) with $x \in S_F, y \in S_N$.

We are now ready to define a set NE with $NE_1 \subseteq NE \subseteq NE_2$. It will contain only pairs $(s,t) \in S \times S_N$ such that $\mathbb{P}(\rho \mid s^-(\rho) = (s,t) \wedge \rho = \rho_1 \rho_2, \rho_2 \in (S_F \times S_N)^\omega) = 0$.

To define NE , we define inductively a sequence $P_1 \subsetneq \dots \subsetneq P_\ell$ of sets of states of the twin machine $Twin(\mathcal{A})$ that cannot be used to give a positive probability to stay ambiguous forever. Then, NE will be defined as the set of states that cannot reach an ambiguous cycle avoiding P_ℓ . This can be computed in linear time in the size of $Twin(\mathcal{A})$ by using Tarjan. It suffices to remove states of P_ℓ and to look for SCCs with self loops.

We now define P_i inductively. First, let us define P_1 as NE_1 . Then we define inductively the set P_{i+1} with $(s,t) \in P_{i+1}$ if $(s,t) \in P_i$ or if there exists $(s,a,s') \in T$ such that for all $(s,t) \xrightarrow{a} (s',t')$, we have $(s',t') \in P_i$. The condition $(s,a,s') \in T$ ensures that a is possible from s . When $P_\ell = P_{\ell+1}$, which must happen after a time bounded by the number of states in $Twin(\mathcal{A})$ steps, we stop the process. That is, P_ℓ is a smallest fix point that can be obtained in polynomial time. We have:

Lemma 7. From every state $(s,t) \in P_\ell$ with $s \in S_F$, there exists a path ρ , $s^-(\rho) = s$ such that for every ρ' such that $s^-(\rho') = t$ and $\sigma(\rho) = \sigma(\rho')$, one has that ρ' is faulty.

We can now define formally NE as the set of states that cannot reach an ambiguous cycle avoiding P_ℓ , that is $NE = (S \times S_N) \setminus \{(s,t) \mid \exists \rho = \rho_1 \rho_2, s^-(\rho) = (s,t) \wedge \rho_2 \text{ avoids } P_\ell \wedge s^-(\rho_2) = s^+(\rho_2)\}$. Using lemma 7, we obtain:

Lemma 8. $NE_1 \subseteq NE \subseteq NE_2$.

Proof: Let (s,t) be a pair of type 1. It cannot reach an ambiguous loop, thus in particular it cannot reach an ambiguous loop avoiding P_ℓ .

Similarly, let (s,t) be a pair in NE , i.e. such that (s,t) cannot reach an ambiguous cycle avoiding P_ℓ . Thus, thanks to lemma 7, we know that the probability that for infinite paths ρ, ρ' , $s^-(\rho) = s', s^-(\rho') = t'$ and ρ, ρ' are ambiguous is 0 since they will always have an occasion to have a future that disambiguates them (i.e. the probability to avoid in P_ℓ is 0).

Thus, $\mathbb{P}[\rho \in \mathcal{P}_F^\infty(\mathcal{A}, s) : \exists \rho' \in \mathcal{P}_N^\infty(\mathcal{A}, t), \sigma(\rho) = \sigma(\rho')] = 0$ and $NE \subseteq NE_2$. \square

We can use this to reduce (with low complexity) the size of a pseudo-quantitative-diagnoser:

Theorem 2. From an automaton \mathcal{A} , one can build in quadratic time a pseudo-quantitative-diagnoser \mathcal{M}' such that the probability of a faulty ambiguous run in \mathcal{A} is equal to the probability to reach an ambiguous SCC in \mathcal{M}' . Further, there are automata \mathcal{A} such that the size of \mathcal{M}' is exponentially smaller than the diagnoser built in Section 3.

Proof: The set NE is computable in quadratic time w.r.t to the number of transitions of the original automaton \mathcal{A} . We then define the pseudo-diagnoser $\mathcal{M}' = (S \times Q, \Sigma, (s_0, \{s_0\}), T')$ with $T' = \{((s,q), a, (s',q' \setminus \{t \mid (s,t) \in NE\}))\}$ such that $(s,a,s') \in T$ and $q' = \{t' \mid \exists t \in q(t,a,t') \in T\}$. Since $NE \subseteq NE_2$ (Lemma 8), one can apply Lemma 5 to obtain that the probability of a faulty ambiguous run in \mathcal{A} is equal to the probability to reach an ambiguous SCC in \mathcal{M}' .




Fig. 4. An example of automaton \mathcal{A}_3 that has an exponential sized diagnoser, and its twin machine (bottom)

Figure 4 presents an example where the pseudo diagnoser \mathcal{M}' is exponentially smaller than the natural diagnoser based on the determinized of \mathcal{A} , which is the same as the diagnoser presented in Section 3 as there is a unique faulty state.

Indeed, the number of states of the natural diagnoser of the automaton \mathcal{A}_3 is a $O(2^n)$, as safe runs can produce a c only $n - 1$ steps after producing a b . That is, the diagnoser needs to distinguish between 2^{n-1} cases, depending on the last $n - 1$ letters in $\{a, b\}$.

Using the twin machine, the number of states of the pseudo-quantitative-diagnoser diminishes tremendously. First, NE_1 (and hence P_1) is the set $\{(i, j) | i > 0, j > 0\}$. Then, $P_2 = P_1 \cup \{(f, i) | i \leq n\}$. Indeed, for all $i < n$, there is a transition (f, c, f) but there is no transition starting in i labeled by c and then no successor to (f, i) by c . Thus, for all transition $(f, i) \xrightarrow{c} (s', t')$, we have $(s', t') \in P_1$, trivially because there is no such transition $(f, i) \xrightarrow{c} (s', t')$. Thus $(f, i) \in P_2$ for all $i < n$. Similarly, we obtain that $(f, n) \in P_2$ using the transition (f, a, f) .

Now, state $(0, 0)$ only has successors in P_2 . Thus $(0, 0) \in P_3$. That is, P_3 is made up of all the states of the twin machine and since there is no ambiguous cycle outside of P_3 , NE contains all the states of the twin machine. Hence, the pseudo-diagnoser is very simple: for all s , $NE(s) = S$ and then every state in the pseudo-diagnoser is in the form of (s, \emptyset) with $s \in S$. That is, the pseudo-diagnoser is isomorph to the original automaton, that is it is of size linear in $|S|$. Therefore, this transformation avoids the exponential blow-up required by using an exact diagnoser. \square

5. CONCLUSION

In this paper, we have proposed the notion of pseudo-diagnosers of systems, which are smaller than diagnosers by up to an exponential factor, removing both faulty and non-faulty states. Yet,

they provide enough information to allow computing complex diagnosability degrees in terms of reachability probabilities.

Our notion of negligible sets of pairs of states makes A-diagnosability trivial, as a system is A-diagnosable iff all pairs of states (s, t) in the twin machine $Twin(\mathcal{A})$ are negligible (ie belong to NE_2). As A-diagnosability is PSPACE-complete, it means that finding every negligible pair (of type 2) cannot be done in polynomial time in the worst case. Still, we proposed a polynomial time algorithm to find most of these negligible pairs, and show on an example that it can reduce the size of the pseudo-diagnoser by an exponential factor.

ACKNOWLEDGEMENT

The authors who like to thank Alexandre Blanché who contributed to part of the results presented in this paper.

REFERENCES

- [1] N. Bertrand, E. Fabre, S. Haar, S. Haddad, L. Helouet. Active diagnosis for probabilistic systems. In proc. FoS-SaCS'14, LNCS 8412, pp29-42, Springer, 2014.
- [2] N. Bertrand, S. Haddad, E. Lefauchaux. Foundation of Diagnosis and Predictability in Probabilistic Systems. In proc. 34th IARCS Annual Conf. on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14), LIPIcs 29, pp. 417-429, 2014.
- [3] H. Bazille, E. Fabre, B. Genest. Diagnosability Degree of Stochastic Discrete Event Systems. In *CDC'17*, 5726-5731, IEEE, 2017.
- [4] H. Bazille, E. Fabre, B. Genest. Complexity reduction techniques for quantified diagnosability of stochastic systems. Long version available at <http://perso.crans.org/~genest/BFG18b.pdf>.
- [5] A. Boussif, M. Ghazel. An Experimental Comparison of Two Approaches for Diagnosability Analysis of Discrete Event Systems - A Railway Case-Study. In *VECoS'17*, p. 92-107, LNCS 10466, 2017.
- [6] E. Fabre. Diagnosis and automata. In *Control of Discrete-Event Systems - Automata and Petri Net Perspectives*, Lecture Notes in Control and Information Sciences 433:85-106, Springer, 2013.
- [7] E. Fabre, L. Jezequel. On the construction of probabilistic diagnosers. In Proc. WODES'10, pp. 229-234, Elsevier, 2010.
- [8] S. Haar, S. Haddad, T. Melliti, S. Schwoun. Optimal constructions for active diagnosis. *J. Comput. Syst. Sci.* 83(1): 101-120, Elsevier, 2017.
- [9] S. Jiang, Z. Huang, V. Chandra, R. Kumar. A polynomial algorithm for testing diagnosability of discrete-event systems. *IEEE Trans. Aut. Cont.* 46(8):1318-1321, 2001.
- [10] F. Nouioua, P. Dague. A probabilistic analysis of diagnosability in discrete event systems. In Proc. ECAI'08, Frontiers in Artif. Intel. and Appl. 178:224-228, IOS Press, 2008.
- [11] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamo-hideen, D. Teneketzis. Diagnosability of discrete event systems. *IEEE Trans. Aut. Cont.* 40(9):1555-1575, 1995.
- [12] D. Thorsley, D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Trans. Aut. Cont.* 50(4):476-492, 2005.
- [13] T.-S. Yoo, S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *IEEE Trans. Aut. Cont.* 47(9):1491-1495, 2002.