

Active diagnosis for probabilistic systems^{*}

Nathalie Bertrand¹, Éric Fabre¹,
Stefan Haar^{1,2}, Serge Haddad², and Loïc Hélouët¹

¹ Inria, France

² LSV, ENS Cachan & CNRS & Inria, France

Abstract. The diagnosis problem amounts to deciding whether some specific “fault” event occurred or not in a system, given the observations collected on a run of this system. This system is then diagnosable if the fault can always be detected, and the active diagnosis problem consists in controlling the system in order to ensure its diagnosability. We consider here a stochastic framework for this problem: once a control is selected, the system becomes a stochastic process. In this setting, the active diagnosis problem consists in deciding whether there exists some observation-based strategy that makes the system diagnosable with probability one. We prove that this problem is EXPTIME-complete, and that the active diagnosis strategies are belief-based. The *safe* active diagnosis problem is similar, but aims at enforcing diagnosability while preserving a positive probability to non faulty runs, i.e. without enforcing the occurrence of a fault. We prove that this problem requires non belief-based strategies, and that it is undecidable. However, it belongs to NEXPTIME when restricted to belief-based strategies. Our work also refines the decidability/undecidability frontier for verification problems on partially observed Markov decision processes.

1 Introduction

Diagnosis for discrete event systems was introduced in [SSL⁺95], and can be described as follows: a labeled transition system performs a run, which may fire some specific events called *faults*. Some of the transition labels are observable, so one gets information about the performed run through its trace, i.e. its sequence of observed labels. The diagnosis problem then amounts to determining whether a fault event occurred or not given the observed trace. The trace is called faulty (resp. correct) if all runs that can have produced it contain (resp. do not contain) a fault. In the remaining cases the trace is called ambiguous. Along with the diagnosis problem comes the diagnosability question: does there exist an infinite ambiguous

^{*} This work was supported by project ImpRo ANR-2010-BLAN-0317 and the European Union Seventh Framework Programme [FP7/2007-2013] under grant agreement 257462 HYCON2 NOE.

trace (thus forbidding diagnosis)? For finite transition systems, checking diagnosability was proved to have a polynomial complexity [YL02].

Diagnosis and diagnosability checking have been extended to numerous models (Petri nets [CGLS09], pushdown systems [MP09], etc.) and settings (centralized, decentralized, distributed), and have had an impact on important application areas, e.g. for telecommunication network failure diagnosis. Several contributions have considered enforcing the diagnosability of a system. Under the generic name of active diagnosis, the problems take quite different shapes. They range from the selection of minimal sets of observable labels that make the system diagnosable [CT08], to the design of controllers that select a diagnosable sublanguage of a system [SLT98], and to online aspects that either turn on and off sensors [TT07, CT08] or modify an action plan [CP09] in order to reduce the amount of ambiguity. Probabilistic systems have also received some attention [TT05, FJ10], with two essential motivations: determining the likelihood of a fault given an observed trace and defining diagnosability for probabilistic systems. Two definitions have been proposed: The A-diagnosability, which requires that the ambiguous traces have a null probability, and the weaker AA-diagnosability, which requires that fault likelihood will converge to one with probability one. Interestingly, the A-diagnosability does not depend on the specific values of transition probabilities, but only on their support: it is thus a structural property of a system, which can be checked in polynomial time on finite state systems.

Here we address the question of active diagnosis for stochastic systems. We elaborate on two recent contributions. The first one [HHMS13] improves the work in [SLT98] and designs an observation-based controller that enables a subset of actions in the system in order to make it diagnosable while preserving its liveness. Optimal constructions are then proposed the most relevant for our work being the characterization of unambiguous traces by a deterministic Büchi automaton with minimal size. The second one [BBG12] considers probabilistic Büchi automata, a subclass of partially observed Markov decision processes (POMDP), and proves that checking the existence of strategies that almost surely achieve a Büchi condition on POMDP is EXPTIME-complete. The result was later extended in [BGG09]. This motivates the use of POMDP as semantics for the models we consider.

The first contribution of this paper is a framework for the active diagnosis problem of probabilistic systems. The models we consider are weighted and labeled transition systems, where some transitions represent a fault. Some of the transition labels are observable, and similarly

some are controllable. From a given state of the system, and given a set of enabled labels, one derives a transition probability by normalization of transition weights. The active diagnosis problem amounts to designing a label activation strategy that enforces the stochastic diagnosability of the system while preserving its liveness. As a second contribution, this problem is proved to be decidable, and EXPTIME complete. The resulting strategies are belief-based, i.e. they only depend on the set of possible states of the system given past observations, regardless of the exact values of transition weights. As a third contribution, we introduce and analyze the *safe* active diagnosis problem. It extends the active diagnosis by enforcing a positive probability to correct runs. In other words, this rules out strategies that would reach diagnosability only by enforcing the occurrence of a fault. We prove that safe active diagnosis may require non belief-based strategies, and that the existence of such strategies is an undecidable problem. This result refines the decidability/undecidability frontier for POMDP: the existence of a strategy simultaneously ensuring a Büchi condition almost-surely and a safety condition with positive probability is undecidable. This may seem surprising since the existence of strategies for each objective taken separately is decidable. As a last contribution, we prove that, restricted to belief-based strategies, the safe active diagnosis problem becomes decidable and belongs to NEXPTIME.

The paper is organized as follows: section 2 introduces the active diagnosis problem for probabilistic systems, and compares it with the state of the art. Section 3 proposes resolution techniques for active diagnosis. Section 4 analyzes the safe active diagnosis problem. Section 5 concludes this work. Due to lack of space, several proofs are provided in appendix.

2 The active diagnosis problem

This section recalls diagnosis problems from the literature, and formalizes the new problems we are interested in.

2.1 Passive (probabilistic) diagnosis

When dealing with stochastic discrete event systems diagnosis, systems are often modeled using labeled transition systems.

Definition 1. A probabilistic labeled transition system (*pLTS*) is a tuple $\mathcal{A} = \langle Q, q_0, \Sigma, T, \mathbf{P} \rangle$ where:

- Q is a set of states with $q_0 \in Q$ the initial state;

- Σ is a finite set of events;
- $T \subseteq Q \times \Sigma \times Q$ is a set of transitions;
- \mathbf{P} is the transition matrix from T to $\mathbb{Q}_{>0}$ fulfilling for all $q \in Q$:

$$\sum_{(q,a,q') \in T} \mathbf{P}[q, a, q'] = 1.$$

Observe that a pLTS is a labeled transition system (LTS) equipped with transition probabilities. The transition relation of the underlying LTS is defined by: $q \xrightarrow{a} q'$ for $(q, a, q') \in T$; this transition is then said to be *enabled* in q . A *run* over the word $\sigma = a_1 a_2 \dots \in \Sigma^\omega$ is a sequence of states $(q_i)_{i \geq 0}$ such that $q_i \xrightarrow{a_{i+1}} q_{i+1}$ for all $i \geq 0$, and we write $q_0 \xrightarrow{\sigma}$ if such a run exists. A finite run over $w \in \Sigma^*$ is defined analogously, and we write $q \xrightarrow{w} q'$ if such a run ends at state q' . A state q is *reachable* if there exists a run $q_0 \xrightarrow{w} q$ for some $w \in \Sigma^*$. On the other hand, forgetting the labels and merging the transitions with same source and target, one obtains a discrete time Markov chain (DTMC).

Definition 2 (Languages of a pLTS). Let $\mathcal{A} = \langle Q, q_0, \Sigma, T, \mathbf{P} \rangle$ be a pLTS. The finite language $\mathcal{L}^*(\mathcal{A}) \subseteq \Sigma^*$ of \mathcal{A} and the infinite language $\mathcal{L}^\omega(\mathcal{A}) \subseteq \Sigma^\omega$ of \mathcal{A} are defined by:

$$\mathcal{L}^*(\mathcal{A}) = \{ w \in \Sigma^* \mid \exists q : q_0 \xrightarrow{w} q \} \quad \mathcal{L}^\omega(\mathcal{A}) = \{ \sigma \in \Sigma^\omega \mid q_0 \xrightarrow{\sigma} \}$$

Observations. In order to formalize problems related to diagnosis, we partition Σ into two disjoint sets Σ_o and Σ_u , the sets of *observable* and of *unobservable events*, respectively. Moreover, we distinguish a special *fault* event $f \in \Sigma_u$. Let σ be a finite word; its length is denoted $|\sigma|$. For $\Sigma' \subseteq \Sigma$, define $\mathcal{P}_{\Sigma'}(\sigma)$, the projection of σ on Σ' , inductively by: $\mathcal{P}_{\Sigma'}(\varepsilon) = \varepsilon$; for $a \in \Sigma'$, $\mathcal{P}_{\Sigma'}(\sigma a) = \mathcal{P}_{\Sigma'}(\sigma) a$; and $\mathcal{P}_{\Sigma'}(\sigma a) = \mathcal{P}_{\Sigma'}(\sigma)$ for $a \notin \Sigma'$. Write $|\sigma|_{\Sigma'}$ for $|\mathcal{P}_{\Sigma'}(\sigma)|$, and for $a \in \Sigma$, write $|\sigma|_a$ for $|\sigma|_{\{a\}}$. When σ is an infinite word, its projection is the limit of the projections of its finite prefixes. This projection can be either finite or infinite. As usual the projection is extended to languages. In the rest of the paper, we will only use \mathcal{P}_{Σ_o} , the projection onto observable events, and hence we will drop the subscript and simply write \mathcal{P} instead of \mathcal{P}_{Σ_o} .

With respect to the partition of $\Sigma = \Sigma_o \uplus \Sigma_u$, a pLTS \mathcal{A} is *convergent* if $\mathcal{L}^\omega(\mathcal{A}) \cap \Sigma^* \Sigma_u^\omega = \emptyset$ (i.e. there is no infinite sequence of unobservable events from any reachable state). When \mathcal{A} is convergent, then for all $\sigma \in \mathcal{L}^\omega(\mathcal{A})$, one has $\mathcal{P}(\sigma) \in \Sigma_o^\omega$. In the rest of the paper we assume that pLTS are convergent and we will call a *sequence* a finite or infinite word over Σ , and an *observed sequence* a finite or infinite sequence over Σ_o . Clearly, the projection of a sequence on Σ_o yields an observed sequence. Intuitively, a

sequence describes the behavior of a system during an execution, and an observed sequence represents how such a run is perceived. Now, the role of diagnosis is to decide, for any observed sequence, whether a fault has occurred or not.

Ambiguity. A finite (resp. infinite) sequence σ is *correct* if it belongs to $(\Sigma \setminus \{f\})^*$ (resp. $(\Sigma \setminus \{f\})^\omega$). Otherwise σ is called *faulty*. A correct sequence and a faulty sequence may have the same observed projection, yielding ambiguity.

Definition 3 (Classification of observed sequences). Let \mathcal{A} be a pLTS. An observed sequence $\sigma \in \Sigma_o^\omega$ is called *ambiguous* if there exist two sequences $\sigma_1, \sigma_2 \in \mathcal{L}^\omega(\mathcal{A})$ such that $\mathcal{P}(\sigma_1) = \mathcal{P}(\sigma_2) = \sigma$, σ_1 is correct and σ_2 is faulty. An observed sequence $\sigma' \in \mathcal{P}(\mathcal{L}^\omega(\mathcal{A}))$ is *surely faulty* if $\mathcal{P}^{-1}(\sigma') \cap \mathcal{L}^\omega(\mathcal{A}) \subseteq \Sigma^* f \Sigma^\omega$. An observed sequence $\sigma' \in \mathcal{P}(\mathcal{L}^\omega(\mathcal{A}))$ is *surely correct* if $\mathcal{P}^{-1}(\sigma') \cap \mathcal{L}^\omega(\mathcal{A}) \subseteq (\Sigma \setminus \{f\})^\omega$. These notions are defined analogously for finite observed sequences.




Fig. 1. Two examples of pLTS (cLTS), with $\Sigma_u = \{f, u\}$ and $\Sigma_o = \{a, b, c, d, e\}$.

Example. Consider the (convergent) pLTS to the left in Fig. 1, where $\Sigma_u = \{f, u\}$. We assume uniform distributions so we do not represent the probability matrix \mathbf{P} . This pLTS contains infinite ambiguous sequences: immediately after a is observed, an ambiguity appears, and this ambiguity remains in all infinite observed sequences without occurrence of d and finishing with ab^ω . Removing the loop at q_2 and/or q_4 makes all infinite ambiguous sequences disappear.

In the sequel, we will use the characterization of unambiguous sequences using deterministic Büchi automata [HHMS13].

Definition 4 (Büchi automaton). A Büchi automaton over Σ is a tuple $\mathcal{B} = \langle Q, q_0, \Sigma, T, F \rangle$ with $\langle Q, q_0, \Sigma, T \rangle$ its underlying LTS and $F \subseteq Q$ an acceptance condition. A run $(q_i)_{i \geq 0}$ is accepting if $q_i \in F$ for infinitely many values of i . The language $\mathcal{L}(\mathcal{B})$ consists of all words in Σ^ω for which there exists an accepting run. A Büchi automaton is deterministic if for all q, a , $\{q' \mid q \xrightarrow{a} q'\}$ is either a singleton or the empty set.

Theorem 1 ([HHMS13]). Given a pLTS \mathcal{A} with n states, one can build in exponential time a deterministic Büchi automaton \mathcal{B} with $2^{O(n)}$ states whose language is the set of unambiguous sequences of \mathcal{A} .

We briefly sketch the structure of \mathcal{B} . Its states are triples $\langle U, V, W \rangle$, where $U, V, W \subseteq Q$, $U \cup V \cup W \neq \emptyset$ and $V \cap W = \emptyset$, and its transitions are labeled by events from Σ_o , that is \mathcal{B} recognizes observed sequences. The initial state of \mathcal{B} is $\langle \{q_0\}, \emptyset, \emptyset \rangle$. Given an observed sequence σ reaching state $\langle U, V, W \rangle$, U is the set of states of \mathcal{A} reached by a correct sequence with projection σ , and $V \cup W$ is the set of states of \mathcal{A} reached by a faulty sequence with projection σ . When $U = \emptyset$, σ is the projection of faulty sequences of \mathcal{A} . The decomposition between V and W reflects the fact that \mathcal{B} tries to “solve the ambiguity” between U and W (when both are non empty), while V corresponds to a waiting room of states reached by faulty sequences that will be examined when the current ambiguity is resolved. Given some new observation a , a transition from $\langle U, V, W \rangle$ to the new state $\langle U', V', W' \rangle$ is defined as follows. U' is the set of states reached from U by a correct sequence with projection a . Let Y be the set of states reached from U by a faulty sequence with projection a , or reached from V by a sequence with projection a . When W is non empty then W' is the set of states reached from W by a sequence with projection a and $V' = Y$. Otherwise, the faulty sequences ending in states memorized by W cannot be extended by a sequences with projection a , and we set $V' = \emptyset$ and $W' = Y$. The ambiguity between U and W has been resolved, but new ambiguity may arise between U' and W' . Accepting states in F are triples $\langle U, V, W \rangle$ with $U = \emptyset$ or $W = \emptyset$. Hence, all infinite observed sequence of \mathcal{A} passing infinitely often through F are not ambiguous (they resolve ambiguities one after another) and are accepted by \mathcal{B} .

We are now in position to define diagnosability. It is well-known that given a pLTS \mathcal{A} and a Büchi automaton \mathcal{B} , the set of sequences of \mathcal{A} accepted by \mathcal{B} is measurable [Var85]. So the following definition is sound.

Definition 5 (Diagnosability). A pLTS \mathcal{A} is diagnosable if the set of sequences yielding ambiguous observed sequences has null measure. It is safely diagnosable if it is diagnosable and the set of correct sequences has positive measure.

The notion of a safely diagnosable pLTS is introduced to ensure that fault occurrence is not almost sure. This property is important: a diagnosable system which is not safely diagnosable contains only faulty infinite runs. In the rest of the paper, we will consider active diagnosis, that is, ways to force a system to become diagnosable using a controller. If a controlled system is not safely diagnosable, then the diagnosis solution enforced by the controller is not acceptable.

Example. Consider again the pLTS to the left in Fig. 1. The only ambiguous observed (infinite) sequences necessarily terminate with ab^ω . But the probability to produce such a sequence is null, as the system will reach q_5 with probability one. In other words, ambiguity vanishes at the first occurrence of d or cb . Since cb occurs with probability one, this pLTS is diagnosable. This pLTS is also safely diagnosable, as it can produce correct sequences with a positive probability: there is a positive probability to reach q_5 by sequence uac . If one removes state q_5 and its connected transitions, the system remains diagnosable, but is not safely diagnosable anymore: as the graph of the pLTS is strongly connected, every transition will be visited (infinitely often) with probability 1 implying that f occurs.

2.2 Active probabilistic diagnosis

In order to allow control over the actions of a system while preserving the possibility of a probabilistic semantic, we introduce controllable weighted labelled transition system where probabilities are replaced by weights.

Definition 6. A controllable weighted labelled transition system (cLTS) is a tuple $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ where:

- Q is a finite set of states with $q_0 \in Q$ the initial state;
- the event alphabet Σ is partitionned into observable Σ_o and unobservable Σ_u events, and also partitionned into controllable Σ_c and uncontrollable Σ_e (e for environment) events;
- $\Sigma_u = \{f, u\}$ contains a faulty event, and a non-faulty one;
- $T : S \times \Sigma \times S \rightarrow \mathbb{N}$ is the transition function, labelling transitions with integer weights.

A cLTS has an underlying LTS where the transition relation is defined by $q \xrightarrow{a} q'$ if $T(q, a, q') > 0$. All previous definitions that do not depend on probabilities equally apply to cLTS. We denote by $\text{Ena}(q)$ the set of events that are enabled in q : $\text{Ena}(q) = \{a \in \Sigma \mid \exists q', T(q, a, q') > 0\}$. We assume that the cLTS is convergent and *live*: for all q , $\text{Ena}(q) \neq \emptyset$.

Let $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ be a cLTS. For $q \in Q$ and $\Sigma^\bullet \subseteq \Sigma$, we define

$$G^{\Sigma^\bullet}(q) = \sum_{a \in \Sigma^\bullet, q' \in Q} T(q, a, q')$$

as the (possibly null) global outgoing weight from q restricted to Σ^\bullet -events. Similarly, we define a normalization of the transition relation restricted to Σ^\bullet by

$$T^{\Sigma^\bullet}(q, a, q') = \begin{cases} \frac{T(q, a, q')}{G^{\Sigma^\bullet}(q)} & \text{if } a \in \Sigma^\bullet \text{ and } T(q, a, q') > 0 \\ 0 & \text{otherwise} \end{cases}$$

For a given finite set X , we define by $\text{Dist}(X)$ the set of probabilistic distributions over X . Let $x \in X$, we denote by $\mathbf{1}_x$ the Dirac distribution on x . For a distribution $\delta \in \text{Dist}(X)$, the support of δ is the set $\text{Supp}(\delta) = \{x \in X \mid \delta(x) > 0\}$.

A *strategy* for a cLTS \mathcal{C} is a mapping $\pi : \Sigma_o^* \rightarrow \text{Dist}(2^\Sigma)$ such that for every $\sigma \in \Sigma_o^*$, for every $\Sigma' \in \text{Supp}(\pi(\sigma))$, $\Sigma' \supseteq \Sigma_e$. A strategy consists in, given some observation, randomly choosing a subset of allowed events that includes the uncontrollable events. Given a cLTS \mathcal{C} and a strategy π , we consider configurations of the form $(\sigma, q, \Sigma^\bullet) \in \Sigma_o^* \times Q \times 2^\Sigma$ where σ is the observed sequence, q is the current state and Σ^\bullet is the set of allowed events. We define inductively the set $\text{Reach}_\pi(\mathcal{C})$ of reachable configurations under π :

- for all $\Sigma^\bullet \in \text{Supp}(\pi(\varepsilon))$, $(\varepsilon, q_0, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$;
- for all $(\sigma, q, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, for all $a \in \Sigma_u \cap \Sigma^\bullet$, such that $q \xrightarrow{a} q'$ $(\sigma, q', \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, denoted $(\sigma, q, \Sigma^\bullet) \xrightarrow{a}_\pi (\sigma, q', \Sigma^\bullet)$;
- for all $(\sigma, q, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, for all $a \in \Sigma_o \cap \Sigma^\bullet$ such that $q \xrightarrow{a} q'$ and $\Sigma^{\bullet'} \in \text{Supp}(\pi(\sigma a))$, $(\sigma a, q', \Sigma^{\bullet'}) \in \text{Reach}_\pi(\mathcal{C})$, denoted $(\sigma, q, \Sigma^\bullet) \xrightarrow{a}_\pi (\sigma a, q', \Sigma^{\bullet'})$.

A strategy π is said to be *live* if for every configuration $(\sigma, q, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, $G^{\Sigma^\bullet}(q) \neq 0$. Live strategies are the only relevant strategies as the other strategies introduce deadlocks. We are now in position to introduce the semantics of a cLTS. It is defined w.r.t. to some live strategy π as a pLTS. Its set of states is $\text{Reach}_\pi(\mathcal{C})$ with an initial state whose goal is to randomly select w.r.t. π the initial control. The transition probabilities are defined by T^{Σ^\bullet} accordingly to the current control Σ^\bullet except that when an observable action occurs it must be combined with the random choice (w.r.t. π) of the next control.

Definition 7. Let \mathcal{C} be a CLTS and π be a live strategy, the pLTS \mathcal{C}_π induced by strategy π on \mathcal{C} is defined as $\mathcal{C}_\pi = \langle Q_\pi, \Sigma, q_{0\pi}, T_\pi, \mathbf{P}_\pi \rangle$ where:

- $Q_\pi = \{q_{0\pi}\} \cup \text{Reach}_\pi(\mathcal{C})$;
- for all $(\varepsilon, q_0, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, $(q_{0\pi}, u, (\varepsilon, q_0, \Sigma^\bullet)) \in T_\pi$;
- for all $(\sigma, q, \Sigma^\bullet), (\sigma', q', \Sigma^{\bullet'}) \in \text{Reach}_\pi(\mathcal{C})$,
 $((\sigma, q, \Sigma^\bullet), a, (\sigma', q', \Sigma^{\bullet'})) \in T_\pi$ iff $(\sigma, q, \Sigma^\bullet) \xrightarrow{a}_\pi (\sigma', q', \Sigma^{\bullet'})$;
- for all $(\varepsilon, q_0, \Sigma^\bullet) \in \text{Reach}_\pi(\mathcal{C})$, $\mathbf{P}_\pi(q_{0\pi}, u, (\varepsilon, q_0, \Sigma^\bullet)) = \pi(\varepsilon)(\Sigma^\bullet)$;
- for all $((\sigma, q, \Sigma^\bullet), a, (\sigma, q', \Sigma^\bullet)) \in T_\pi$, for all $a \in \Sigma_u \cap \Sigma^\bullet$,
 $\mathbf{P}_\pi((\sigma, q, \Sigma^\bullet), a, (\sigma, q', \Sigma^\bullet)) = T^{\Sigma^\bullet}(q, a, q')$;
- for all $((\sigma, q, \Sigma^\bullet), a, (\sigma a, q', \Sigma^{\bullet'})) \in T_\pi$, for all $a \in \Sigma_o \cap \Sigma^\bullet$,
 $\mathbf{P}_\pi((\sigma, q, \Sigma^\bullet), a, (\sigma a, q', \Sigma^{\bullet'})) = T^{\Sigma^\bullet}(q, a, q') \cdot \pi(\sigma.a)(\Sigma^{\bullet'})$.

We can now formalize the decision problems we are interested in.

Definition 8 ((Safe) Active probabilistic diagnosis). Given a cLTS $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$, the active probabilistic diagnosis problem asks, whether there exists a live strategy π in \mathcal{C} such that the pLTS \mathcal{C}_π is diagnosable. The safe active probabilistic diagnosis problem asks whether there exists a live strategy π in \mathcal{C} such that the pLTS \mathcal{C}_π is safely diagnosable. The synthesis problems consists in building a live strategy π in \mathcal{C} such that the pLTS \mathcal{C}_π is (safely) diagnosable.

Example. Consider the cLTS to the right in Fig. 1 with all weights equal to 1 and $\Sigma_o = \Sigma_c$. Without control, the system is not diagnosable as all observed sequences starting by aa are ambiguous. So the strategy should disable action a for each correct observed sequence ending by ab^* . In addition, if this strategy always forbids c , the system becomes diagnosable, but the occurrence of a fault is enforced: so it is not safely diagnosable. Alternatively, if the strategy always forbids e , the system becomes safely diagnosable, as we obtain a pLTS “weakly probabilistically bisimilar” to the one on the left in Fig. 1.

3 Analysis of the active probabilistic diagnosis problem

To solve the active probabilistic diagnosis problem, we reduce it to a decidable problem on POMDP: namely, the existence of a strategy ensuring a Büchi objective with probability one [BBG12, BGG09].

Definition 9 (POMDP). A partially observable Markov decision process (POMDP) is a tuple $\mathcal{M}_\mathcal{C} = \langle Q, q_0, \text{Obs}, \text{Act}, T \rangle$ where

- Q is a finite set of states with q_0 the initial state;
- $\text{Obs} : Q \rightarrow \mathcal{O}$ assigns an observation $O \in \mathcal{O}$ to each state.
- Act is a finite set of actions;
- $T : Q \times \text{Act} \rightarrow \text{Dist}(Q)$ is a partial transition function. Letting $\text{Ena}(q) = \{a \in \text{Act} \mid T(q, a) \text{ is defined}\}$, we assume that:
 - for all $q \in Q$, $\text{Ena}(q) \neq \emptyset$, and
 - whenever $\text{Obs}(q) = \text{Obs}(q') = O$, then $\text{Ena}(q) = \text{Ena}(q') = \text{Ena}(O)$.

A *decision rule* is an item of $\text{Dist}(\text{Act})$ that resolves non-determinism by randomization. A *strategy* maps histories of observations to decision rules. Formally, a strategy is a function $\pi : \mathcal{O}^+ \rightarrow \text{Dist}(\text{Act})$ such that for all $O_1 \cdots O_i$, $\text{Supp}(\pi(O_1 \cdots O_i)) \subseteq \text{Ena}(O_i)$. Given a strategy π and an initial distribution δ over states, a POMDP M becomes a stochastic process that can be represented by a possibly infinite pLTS denoted $M(\pi)$. One denotes $\mathbb{P}_\pi^\delta(\text{Ev})$ the probability that event Ev is realized in this process.

A *belief* is a subset of some observation O that corresponds to the possible reachable states w.r.t. some sequence of observations. The initial belief is $\{q_0\}$ and given a current belief B , a decision rule δ and a observation O , the belief $\Delta(B, (\delta, O))$ obtained after δ has been applied and O has been observed is defined by: $\bigcup_{q \in B, a \in \text{Supp}(\delta)} \text{Supp}(T(q, a)) \cap \text{Obs}^{-1}(O)$. A strategy which only depends on the current belief is called a *belief-based strategy*.

In order to provide a POMDP $M_{\mathcal{C}}$ for the diagnosis problems of a cLTS \mathcal{C} , we face several difficulties. First, in a cLTS the observations are related to actions while in a POMDP they are related to states. Fortunately all the information related to ambiguity is included in the deterministic Büchi automaton described in section 2. Thus (with one exception) the states are pairs of a state of the Büchi automaton and a state of the cLTS. In \mathcal{C} , the control is performed by allowing a subset of events. Thus actions of $M_{\mathcal{C}}$ are subset of events that includes the uncontrollable events. Given some control Σ' , for defining the transition probability of $M_{\mathcal{C}}$ from (l, q) to (l', q') , one must consider all paths in \mathcal{C} labelled by events of Σ' from q to q' that end with an observable event b . The probability of any such path is obtained by the product of the individual step probabilities. The latter are then defined by the normalization of weights w.r.t. Σ' . They cannot be infinite paths of unobservable events due to the convergence of \mathcal{C} . However some path can reach a state where no event of Σ' is possible. In other words, the control Σ' applied in (l, q) has a non null probability to reach a deadlock (i.e. the chosen decision rule leads to a non live strategy for the cLTS). In order to capture this behaviour and to obtain a non defective probability distribution, we add an additional state **lost**,

that corresponds to such deadlocks. The next definition formalizes our approach.

Definition 10. *The POMDP $M_C = \langle Q^M, q_0^M, \text{Obs}, \text{Act}, T^M \rangle$ derived from a cLTS $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ and its associated deterministic Büchi automaton $\mathcal{B} = \langle L, l_0, \Sigma_o, T^{\mathcal{B}}, F \rangle$ is defined by:*

- $Q^M = L \times Q \uplus \{\mathbf{lost}\}$ with $q_0^M = ((l_0, q_0))$;
- the set of observations is $\mathcal{O} = L \cup \{\mathbf{lost}\}$, with
 $\text{Obs}((l, q)) = l$ and $\text{Obs}(\mathbf{lost}) = \mathbf{lost}$;
- $\text{Act} = \{\Sigma' \mid \Sigma' \supseteq \Sigma_e\}$;
- for all $(l, q) \in Q^M$ and $\Sigma' \in \text{Act}$ $T^M((l, q), \Sigma') = \mu$ where:
 - $\mu((l', q'))$ is defined by:

$$\sum_{\substack{l \xrightarrow{b} l' \\ b \in \Sigma' \cap \Sigma_o}} \sum_{\substack{q \xrightarrow{a_1} q_1 \dots \xrightarrow{a_n} q_n \xrightarrow{b} q' \\ a_1 \dots a_n \in \Sigma' \cap \Sigma_u}} T^{\Sigma'}(q, a_1, q_1) \cdot \left(\prod_{i=1}^{n-1} T^{\Sigma'}(q_i, a_{i+1}, q_{i+1}) \right) \cdot T^{\Sigma'}(q_n, b, q')$$

- $\mu(\mathbf{lost})$ is defined by:

$$\sum_{\substack{q \xrightarrow{a_1} q_1 \dots \xrightarrow{a_n} q_n \\ a_1 \dots a_n \in \Sigma' \cap \Sigma_u \\ G^{\Sigma'}(q_n) = 0}} T^{\Sigma'}(q, a_1, q_1) \cdot \prod_{i=1}^{n-1} T^{\Sigma'}(q_i, a_{i+1}, q_{i+1})$$

- $T^M(\mathbf{lost}, \Sigma') = \mathbf{1}_{\mathbf{lost}}$ for all $\Sigma' \in \text{Act}$.

Given \mathcal{C} , the construction of the Büchi automaton \mathcal{B} is performed in exponential time. The construction of M_C is also done in exponential time. Indeed, there is an exponential blowup for Act but again w.r.t. \mathcal{C} . Finally, while the distributions μ of action effects are presented in the definition as sums over paths of \mathcal{C} , each one can be computed by a matrix inversion in polynomial time (as done in discrete time Markov chains).

The next lemma is a straightforward consequence of the properties of \mathcal{B} and the above definition of M_C . Here we use LTL notations to denote sets of paths in a POMDP, such as \diamond , \square and $\square\diamond$ for eventually, always and infinitely often respectively.

Lemma 1. *\mathcal{C} is actively diagnosable if and only if there exists a strategy π in M_C such that $\mathbb{P}_\pi^{q_0}(M_C \models \square\diamond(W = \emptyset \vee U = \emptyset)) = 1$.*

Moreover, \mathcal{C} is safely actively diagnosable if and only if there exists a strategy π in M_C such that $\mathbb{P}_\pi^{q_0}(M_C \models \square\diamond(W = \emptyset \vee U = \emptyset)) = 1$ and $\mathbb{P}_\pi^{q_0}(M_C \models \square(U \neq \emptyset)) > 0$.

In the statement of Lemma 1, $W = \emptyset \vee U = \emptyset$ is a shorthand to denote the set of states $(\langle U, V, W \rangle, q)$ in $\mathbf{M}_{\mathcal{C}}$ such that either $W = \emptyset$ or $U = \emptyset$; similarly, $U \neq \emptyset$ represents the set of states $(\langle U, V, W \rangle, q)$ such that $U \neq \emptyset$. As a consequence of Lemma 1, the active diagnosis problem for controllable LTS reduces to the existence of an almost-sure winning strategy for a Büchi objective on some exponential size POMDP.

Theorem 2. *The active probabilistic diagnosis decision and synthesis problems are EXPTIME-complete. There exists a family $(\mathcal{C}_n)_{n \in \mathbb{N}}$ of actively diagnosable cLTS with the size of \mathcal{C}_n in $O(n)$, and such that any winning strategy for $\mathbf{M}_{\mathcal{C}_n}$ diagnosable requires at least $2^{O(n)}$ memory-states.*

The EXPTIME upper bound may seem surprising, since $\mathbf{M}_{\mathcal{C}}$ is exponential in the size of \mathcal{C} , and the procedure to decide whether there exists a strategy in a POMDP to ensure a Büchi objective with probability 1 is in EXPTIME, due to the use of beliefs. However, in the POMDP $\mathbf{M}_{\mathcal{C}}$ we consider, the information on the belief is already contained in the state $(\langle U, V, W \rangle, q)$, as $U \cup V \cup W$. Therefore, a second exponential blowup, due to the beliefs, is avoided and the active probabilistic diagnosis problem remains in EXPTIME.

4 Analysis of the safe active probabilistic diagnosis problem

As will be shown below, the status of the active diagnosis problem changes when the safety requirement is added. The next proposition highlights this difference and it is the basis for the undecidability result.

Proposition 1. *There exists a cLTS which is safely actively diagnosable and such that all belief-based strategies are losing.*

Proof. Let us consider the cLTS of Figure 2 with $\Sigma_u = \{u, f\}$ and $\Sigma_e = \{u, f, c\}$, and where all weights are equal to 1.

Pick any sequence of positive integers $\{\alpha_i\}_{i \geq 1}$ such that $\prod_{i \geq 1} 1 - 2^{-\alpha_i} > 0$. Define $A = \{a\} \cup \Sigma_e$ and $\bar{A} = \{\bar{a}\} \cup \Sigma_e$. We claim that the strategy π that consists in selecting, at instant n , the n^{th} subset in the following sequence $A^{\alpha_1} \bar{A} A^{\alpha_2} \bar{A} \dots$, is winning. Observe that after an observable sequence of length $i \leq \alpha_1$, the system is either after a faulty sequence in r_1 with probability $\frac{1}{2}$, or after a correct sequence in q_1 with probability 2^{-i-1} , or after a correct sequence in q_2 with probability $\frac{1}{2}(1 - 2^{-i})$. So, after an observable sequence of length $\alpha_1 + 1$, the system is either after a faulty sequence in r_2 with probability $\frac{1}{2}$, or after a faulty sequence in r_1 with probability $2^{-\alpha_1-1}$, or after a correct sequence in q_1 with probability




Fig. 2. A cLTS with only non belief-based strategies for safe diagnosis.

$\frac{1}{2}(1 - 2^{-\alpha_1})$. At the next step, the faulty sequence in r_2 is then detected by the occurrence of c .

Iterating this process we conclude that:

- any fault that may occur after π is applied up to $A^{\alpha_1}\bar{A}A^{\alpha_2}\bar{A}\dots A^{\alpha_i}\bar{A}$, is detected after π is applied up to $A^{\alpha_1}\bar{A}A^{\alpha_2}\bar{A}\dots A^{\alpha_i}\bar{A}A$. So the (full) strategy $\pi = A^{\alpha_1}\bar{A}A^{\alpha_2}\bar{A}\dots$ surely detects faults.
- the probability that there is an infinite correct sequence is equal to $\frac{1}{2}\prod_{i \geq 1} 1 - 2^{-\alpha_i} > 0$, due to our choice of the α_i 's. Therefore, correct sequences have positive probability under π .

Consider a belief-based strategy π . There are three possible subsets of allowed events: A , \bar{A} and Σ . The decision rule associated with belief $\{q_0\}$ must allow a in order to get the possibility of a correct sequence which, in case a occurs, leads to belief $\{q_1, q_2, r_1\}$. Consider the (randomized) decision rule of π associated with belief $\{q_1, q_2, r_1\}$: $p_A \cdot A + p_{\bar{A}} \cdot \bar{A} + p_{\Sigma} \cdot \Sigma$ (denoted \mathbf{p}). If $p_A = 1$, then the possible first fault remains undetected, and π is losing. So \bar{a} may occur leading to belief $\{q_1, r_0, r_2\}$.

Consider the decision rule of π associated with belief $\{q_1, r_0, r_2\}$: $p'_A \cdot A + p'_{\bar{A}} \cdot \bar{A} + p'_{\Sigma} \cdot \Sigma$ (denoted \mathbf{p}'). If $p'_{\bar{A}} = 1$, then at the next instant, there is no possible correct sequence, and π is losing.

So $p'_{\bar{A}} < 1$ and $p_A < 1$. Assume now that the current distribution of states is $\alpha q_1 + \beta r_0 + (1 - \alpha - \beta)r_2$ (with belief $\{q_1, r_0, r_2\}$). The distribution after the next occurrence of \bar{a} is defined by $\alpha_{\mathbf{p}, \mathbf{p}'} \alpha q_1 + (1 - \alpha_{\mathbf{p}, \mathbf{p}'}) \alpha r_0 + (1 - \alpha)r_2$, where $\alpha_{\mathbf{p}, \mathbf{p}'} < 1$ only depends on \mathbf{p} and \mathbf{p}' . A correct sequence implies an infinite number of \bar{a} ; after n occurrences of \bar{a} the probability of a correct sequence is bounded by $\alpha_{\mathbf{p}, \mathbf{p}'}^n$. So the probability of an infinite correct sequence is null, and π is losing. \square

Elaborating on the previous example we get the following result.

Theorem 3. *The safe active diagnosis problem for cLTS is undecidable.*

Proof (sketch). We perform a reduction from the following undecidable problem: given a POMDP and a set F of states, does there exist a strategy that ensures the Büchi objective $\Box\Diamond F$ with positive probability. The structure of the cLTS we construct is similar to the one of the example from Fig. 2, except that the states q_1 and q_2 are replaced with two copies of the POMDP. From F states in the first copy, one can move to the second one, and with bared actions, one moves back from the second copy to the first one, or moves from the first copy to r_0 .

The following immediate corollary is interesting since both the existence of a strategy achieving a Büchi objective almost surely, and the existence of strategy achieving a safety objective with positive probability are decidable for POMDP [BGG09, CDGH10].

Corollary 1. *The problem whether, given a POMDP M with subsets of states F and I , there exists a strategy π with $\mathbb{P}_\pi(M \models \Box\Diamond F) = 1$ and $\mathbb{P}_\pi(M \models \Box I) > 0$, is undecidable.*

Given that the general safe active diagnosis problem is undecidable, and that belief-based strategies are not sufficient to achieve safe diagnosability, we consider now the restriction of the safe active diagnosis problem to belief-based strategies. Similarly to the case of active diagnosis, we reduce the safe active probabilistic diagnosis for belief-based-strategies to some verification question on POMDP.

Theorem 4. *The safe active probabilistic diagnosis problem restricted to belief-based strategies is in NEXPTIME and EXPTIME-hard.*

5 Conclusion

We studied the active diagnosis and safe active diagnosis problems for probabilistic discrete event systems, within a unifying POMDP framework. While the active diagnosis problem is EXPTIME-complete, the safe active diagnosis problem is undecidable in general, and belongs to NEXPTIME when restricted to belief-based strategies. Since the lower and upper bounds do not coincide for the latter problem, we strive to close the gap between these bounds in future work. Another problem, closely related to diagnosability, is the predictability problem: given any observation, can we detect that the occurrence of a fault *before* it happens? Last, given the tight relation probabilistic diagnosis has with verification problems for POMDP, we plan to investigate further POMDP problems with multiple objectives.

References

- [BBG08] Christel Baier, Nathalie Bertrand, and Marcus Größer. On decision problems for probabilistic büchi automata. In *FoSSaCS'08*, volume 4962 of *Lecture Notes in Computer Science*, pages 287–301. Springer, 2008.
- [BBG12] Christel Baier, Nathalie Bertrand, and Marcus Größer. Probabilistic ω -automata. *Journal of the ACM*, 59(1):1–52, 2012.
- [BD08] D. Berwanger and L. Doyen. On the power of imperfect information. In *Proc. FSTTCS*, volume 2 of *LIPICs*, Bangalore, India, 2008.
- [BGG09] Nathalie Bertrand, Blaise Genest, and Hugo Gimbert. Qualitative determinacy and decidability of stochastic games with signals. In *Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009*, pages 319–328. IEEE Computer Society, 2009.
- [CDGH10] Krishnendu Chatterjee, Laurent Doyen, Hugo Gimbert, and Thomas A. Henzinger. Randomness for free. In *MFCS*, volume 6281 of *Lecture Notes in Computer Science*, pages 246–257. Springer, 2010.
- [CGLS09] M.P. Cabasino, A. Giua, S. Lafortune, and C. Seatzu. Diagnosability analysis of unbounded Petri nets. In *CDC09: 48th IEEE Conf. on Decision and Control*, pages 1267–1272, 2009.
- [CP09] E. Chantbery and Y. Pencolé. Monitoring and active diagnosis for discrete-event systems. In *Proc. SafeProcess'09*, 2009.
- [CT08] F. Cassez and S. Tripakis. Fault diagnosis with static and dynamic observers. *Fundamenta Informaticae*, 88:497–540, 2008.
- [FJ10] E. Fabre and L. Jezequel. On the construction of probabilistic diagnosers. In *WODES: IFAC Workshop on Discrete Event Systems*, 2010.
- [HHMS13] Stefan Haar, Serge Haddad, Tarek Melliti, and Stefan Schwoon. Optimal constructions for active diagnosis. In *Proceedings of the 33rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'13)*, LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013. To appear.
- [MP09] C. Morvan and S. Pinchinat. Diagnosability of pushdown systems. In *Proceedings of the Haifa Verification Conference*, LNCS 6405, 2009.
- [SLT98] M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, July 1998.
- [SSL⁺95] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE Trans. Aut. Cont.*, 40(9):1555–1575, 1995.
- [TT05] D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Transactions on Automatic Control*, 50(4):476–492, 2005.
- [TT07] D. Thorsley and D. Teneketzis. Active acquisition of information for diagnosis and supervisory control of discrete-event systems. *J. of Discrete Event Dynamic Systems*, 17:531–583, 2007.
- [Var85] Moshe Y. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science (FOCS'85)*, pages 327–338. IEEE Computer Society Press, 1985.
- [YL02] T-S. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *IEEE Trans. Automat. Contr.*, 47(9):1491–1495, 2002.

A Additional proofs

A.1 Proof of Theorem 2

EXPTIME upper-bound

Proposition 2. *The active probabilistic diagnosis problem is in EXPTIME.*

Proof. For the sake of completeness, and to justify the EXPTIME upper-bound, we recall the decision algorithm for POMDP with Büchi condition F . The correctness proof can be found in [BGG09], in the more general framework of 2-player stochastic games with signals. Given a POMDP $M = \langle Q, q_0, \text{Obs}, \text{Act}, T \rangle$, we define its *belief automaton*:

Recall that a *belief* B for M is a non empty subset of states included in some observation $O \in \text{Obs}(Q)$. We write \mathcal{Bl} for the set of all beliefs, and we define the deterministic belief automaton $\mathcal{A}_{\mathcal{Bl}}(M) = \langle \mathcal{Bl}, \{q_0\}, \text{Act} \times \text{Obs}, \Delta \rangle$ such that: for $B \in \mathcal{Bl}$, $\alpha \in \text{Act}$ and $O \in \text{Obs}(Q)$, $\Delta(B, (\alpha, O)) = \bigcup_{q \in B} \text{Supp}(T(q, \alpha)) \cap O$. In words, $\Delta(B, (\alpha, O))$ updates the possible set of states the system is in, given the action that has been triggered and the observation that was made.

For a sequence of actions and observations $(\alpha_1, O_1) \cdots (\alpha_n, O_n)$, we write $\Delta(B, (\alpha_1, O_1) \cdots (\alpha_n, O_n))$ for $\Delta(\cdots \Delta(B, (\alpha_1, O_1)), \cdots), (\alpha_n, O_n)$.

For almost-sure Büchi objectives in POMDP, it was proven that belief-based strategies are sufficient, that is, there exists a strategy to achieve a given Büchi objective iff there exists a belief-based strategy for it. Building on the belief automaton, the set Win of beliefs from which there exists a winning strategy (to ensure the Büchi condition almost-surely), can be computed as a greatest fixpoint. Let $\mathcal{Bl}_F = \{B \in \mathcal{Bl} \mid B \subseteq F\}$. Then, Win is the limit of the non-increasing sequence that starts with $\text{Win}_0 = \mathcal{Bl}$ and is defined inductively by:

$$\text{Win}_{n+1} = \{B \in \text{Win}_n \mid \exists (\alpha_1, O_{i_1}) \cdots (\alpha_n, O_{i_n}), \Delta(B, (\alpha_1, O_{i_1}) \cdots (\alpha_n, O_{i_n})) \in \mathcal{Bl}_F \\ \wedge \forall k, \forall O_{j_k}, \Delta(B, (\alpha_1, O_1) \cdots (\alpha_k, O_{j_k})) \neq \emptyset \Rightarrow \Delta(B, (\alpha_1, O_1) \cdots (\alpha_k, O_{j_k})) \in \text{Win}_n\}.$$

Then, there exists a strategy to ensure the Büchi objective $\square \diamond F$ with probability 1, if and only if $\{q_0\} \in \text{Win}$.

Note that this procedure is exponential w.r.t. the size of the input POMDP, due to the construction of the beliefs. However in the case of the POMDP $M_{\mathcal{C}}$ we consider, the belief is already contained in the state (U, V, W, q) , as $U \cup V \cup W$. Therefore, there is no exponential blowup due to

the resolution on the POMDP; the only exponential blowup comes from the Büchi automaton component, hence the active probabilistic diagnosis problem is in EXPTIME. \square

EXPTIME-hardness The proof relies on a reduction from safety games with imperfect information [BD08] and it is adapted from an original proof in [HHMS13] in a non probabilistic context.

Proposition 3 (hardness). *The following problems are EXPTIME-hard.*

- *The existence of a winning strategy for the active diagnosis of a cLTS.*
- *The existence of a winning belief-based strategy for the safe active diagnosis of a cLTS.*

Proof. A safety game $\mathcal{G} = (L, l_0, \Sigma, \Delta, O, F, obs)$ with imperfect information is defined by:

- L a finite set of locations with $l_0 \in L$ the initial location;
- Σ a finite alphabet;
- $\Delta \subseteq L \times \Sigma \times L$ the transition relation such that for all $l \in L$ and $a \in \Sigma$ there exists at least one l' with $(l, a, l') \in \Delta$;
- O a finite set of observations with $F \subseteq O$ the final observations;
- $obs : L \mapsto O$ the observation mapping.

\mathcal{G} is a turn-based game played by two players A and B . It starts in location l_0 with A to play. In the first round, A chooses a letter a_0 in Σ , and then B chooses a location l_1 such that $(l_0, a_0, l_1) \in \Delta$. A only observes $o_1 = obs(l_1)$. The next rounds are played similarly. Player A wins if for all i , $o_i \notin F$.

The problem of existence of a winning strategy for player A is EXPTIME-complete [BD08]. We now describe the reduction of this problem to diagnosis problems for a cLTS \mathcal{C} defined as follows.

- Q , the set of states, is defined by $Q = L \uplus ((L \setminus obs^{-1}(F)) \times \Sigma) \uplus \{\perp\}$ and $q_0 = l_0$.
- The alphabet $\Sigma' = \Sigma \uplus O \uplus \{u, f, z\}$. The unobservable events are u and f and the uncontrollable events are $O \uplus \{u, f, z\}$.
- T the transition relation is defined as follows.
 1. For all $l \in L \setminus obs^{-1}(F)$ and $a \in \Sigma$, $T(l, a, (l, a)) = 1$.
 2. For all $l \in L \setminus obs^{-1}(F)$, $a \in \Sigma$ and $l' \in L$, $T((l, a), obs(l'), l') = 1$ if $(l, a, l') \in \Delta$.
 3. For all $l \in obs^{-1}(F)$, $T(l, u, \perp) = 1$ and $T(l, f, \perp) = 1$.




Fig. 3. A cLTS \mathcal{C}_n with $\Sigma_o = \{a, b, c, d\}$, $\Sigma_c = \{c, d\}$ and weights are all one.

4. $T(\perp, z, \perp) = 1$.
5. All other weights are null.

From the very definition of \mathcal{C} , a sequence is ambiguous if and only if it contains an occurrence of z . Thus a strategy of the \mathcal{C} is winning for the active diagnosis problem if and only if it avoids states $obs^{-1}(F)$. In addition, such a strategy only “controls” the subset of states $L \setminus obs^{-1}(F)$ and due to the assumptions on \mathcal{G} , it can safely restrict the allowed events to a single one. Furthermore since the information available to the strategy is exactly that of player A , a winning strategy for player A in \mathcal{G} provides a winning strategy for the active diagnosis problem of \mathcal{C} and vice versa. In addition, a winning strategy for the active diagnosis problem of \mathcal{C} only allows correct sequences. So it also solves the safe active diagnosis problem. Finally, it is known that in safety games with imperfect information if there is a winning strategy then there is a belief-based winning strategy. So the second problem of the proposition is also EXPTIME-hard. \square

Optimality of belief-based strategies Again, the proof of the next proposition is adapted from an original proof in [HHMS13] in a non probabilistic context.

Proposition 4 (memory lower bound for strategy). *There exists a family $(\mathcal{C}_n)_{n \geq 1}$ of actively diagnosable cLTS with the size of \mathcal{C}_n in $\mathcal{O}(n)$ such that any winning strategy has at least 2^n different memory states.*

Proof. The family of LTS $(\mathcal{C}_n)_{n \geq 1}$ is depicted in Figure 3, where $\Sigma_o = \{a, b, c, d\}$, $\Sigma_c = \{c, d\}$, and the initial state is q_0 . Intuitively, during the n first steps a fault can occur leading to the upper (resp. lower) “branch” of the LTS when followed by a (resp. b).

Formally, let $\sigma = w_1 w_2 y a^\omega \in \Sigma_o^*$ be an observed sequence, where $w_1 w_2 \in \{a, b\}^*$, $1 \leq |w_1| \leq n$, $|w_2| = n - 1$, $y \in \{c, d\}$. Such a sequence has a positive probability to occur. Let $x_1 \cdots x_{|w_1|}$ be the letters of w_1 . There are two possible execution sequences that have triggered $\sigma' = w_1 w_2 y$: the correct sequence σ' itself and the faulty sequence $x_1 \cdots x_{|w_1|-1} f x_{|w_1|} w_2 y$. If $x_{|w_1|} = a$, before the occurrence of y , the current state is q_n in the correct sequence and ℓ_n in the faulty sequence. So if $y = d$ the two sequences will lead to the same state q_{n+1} while if $y = c$ one sequence will lead to ℓ_{n+1} and the other one to q_{n+1} and they will be discriminated by the next observation. The case $x_{|w_1|} = b$ is symmetrical. So σ is ambiguous iff $x_{|w_1|} = a$ and $y = d$ or $x_{|w_1|} = b$ and $y = c$.

The LTS \mathcal{C}_n , is actively diagnosable. However assume that one observes a word $\sigma = a_1 \dots a_m \in \{a, b\}^*$ such that $n \leq m \leq 2n - 1$. Then when $a_{m-n+1} = a$, \mathcal{C} may be in either q_n or ℓ_n , and when $a_{m-n+1} = b$, \mathcal{C} may be in either q_n or r_n . In the former case the controller must forbid d while in the latter it must forbid c . This implies that a winning strategy π must be in two different states after seeing two different words from $\{a, b\}^n$, therefore it must have at least 2^n states. \square

A.2 Proof of Theorem 3

Proof. We proceed by establishing a reduction from the problem of achieving a positive probability in a blind POMDP with a Büchi objective. This problem was shown undecidable in [CDGH10] by establishing that pure strategies are enough for POMDP. Indeed, in the case of a blind POMDP, a pure strategy is an infinite word, and so the problem boils down to the undecidable problem of the existence of an infinite word achieving a positive probability in a Büchi probabilistic automaton [BBG08].




Fig. 4. A POMDP M and the derived cLTS C for the undecidability proof.

Here the cLTS has rational weights that could be transformed in integers by an appropriate scaling without changing the problem.

Let $M = \langle Q^M, i, \text{Obs}, \text{Act}, T^M \rangle$ be a blind POMDP (that is $\text{Obs}(Q^M)$ is a constant function) over the action alphabet $\text{Act} = \{a, b\}$, and $F \subseteq Q^M$ a set of states. From M one builds a cLTS $C = \langle Q, q_0, \Sigma, T \rangle$ over $\Sigma = \{a, b, \bar{a}, \bar{b}, u, f, c\}$, and with $\Sigma_u = \{u, f\}$ and $\Sigma_e = \{u, f, c\}$. The construction is illustrated in Figure 4, where the POMDP is represented above right, with $v \in F$, and the cLTS below. Let us detail the reduction:

- The set of states is $Q = Q_1^M \cup Q_2^M \cup \{q_0, r_0, r_1, r_2\}$, where Q_1^M and Q_2^M are two copies of the states Q^M of the POMDP. The structure of

the CLTS in Figure 2 mimics that of the example of Figure 1. While in the example of Figure 1, the objective was to design a strategy increasing the probability to be in state q_2 when allowing action \bar{a} to be safely diagnosable, the main objective for the CLTS of Figure 2 is to increase the probability for the system to be in a state of Q_2^M before allowing \bar{a} or \bar{b} .

- The transitions from states in $\{q_0, r_0, r_1, r_2\}$ are fully depicted in the figure, with i_1 being the first copy of the initial state of the POMDP. The weights labelling these transitions are all equal to 1.
- For $s, t \in Q^M$ and $x \in \{a, b\}$, we let $T(s_1, \bar{x}, r_0) = 1$, $T(s_2, x, t_2) = T^M(s, x)(t)$ and $T(s_2, \bar{x}, t_1) = T^M(s, x)(t)$.
- For $s \in S \setminus F$, $t \in S$ and $x \in \{a, b\}$, we let $T(s_1, x, t_1) = T^M(s, x)(t)$.
- For $s \in F$, $t \in S$ and $x \in \{a, b\}$, we let $T(s_1, x, t_2) = T^M(s, x)(t)$.
- All other weights are null.

This reduction ensures that there exists a strategy π in M such that $\mathbb{P}_\pi^i(M \models \square \diamond F) > 0$ if and only if the cLTS \mathcal{C} is safely active diagnosable. To prove it, we start with some preliminary remarks. A reasoning similar to the one in Proposition 1 implies that the faulty unambiguous sequences are those ending by $\bar{x}c^\omega$ with $x \in \{a, b\}$, while the correct unambiguous sequences are those belonging $\{a, b, \bar{a}, \bar{b}\}^\omega$ with an infinite number of occurrences of actions in $\{\bar{a}, \bar{b}\}$.

Assume first that there is a winning strategy in the blind POMDP M for the Büchi objective with positive probability. W.l.o.g., this strategy may be assumed to be pure [CDGH10], and thus can be given as an infinite word $\sigma = \sigma_1\sigma_2 \dots \in \{a, b\}^\omega$. We write p for the probability of infinitely meeting F under σ : $p = \mathbb{P}_\sigma^i(M \models \square \diamond F)$.

Pick some infinite sequence $(\beta_j)_{j \in \mathbb{N}}$ with $0 < \beta_j < 1$ and such that $\prod_{j \geq 0} \beta_j > 0$. We iteratively build an infinite increasing sequence $(n_j)_{j \in \mathbb{N}}$ of integers as follows. First $n_0 = 0$, and if n_0, \dots, n_j are defined, $n_{j+1} > n_j$ is set as the smallest integer that satisfies

$$\mathbb{P}_\sigma^i \left(M \models \diamond^{[n_j+1, n_{j+1}]} F \mid M \models \bigwedge_{k=0}^j \diamond^{[n_k+1, n_{k+1}]} F \wedge \square \diamond F \right) \geq \beta_j ,$$

where the notation $\diamond^{[m, M]} F$ stands for the event “ F is visited between the m -th and M -th time instants”. Because σ is a winning strategy, and due to the induction hypothesis, the above conditional probability is well defined, and it tends to 1 as n_{j+1} goes to infinity. So n_{j+1} is well defined. By construction:

$$\mathbb{P}_\sigma^i \left(M \models \bigwedge_{j \geq 0} \diamond^{[n_j+1, n_{j+1}]} F \right) \geq p \prod_{j \geq 0} \beta_j > 0 .$$

We are now in a position to define a winning strategy π in the cLTS \mathcal{C} . Writing $X = \{x\} \cup \Sigma_e$ and $\bar{X} = \{\bar{x}\} \cup \Sigma_e$ for $x \in \{a, b\}$, at time instant k different from any n_j , π selects X with $x = \sigma_k$, and at time instant n_j , π selects \bar{X} with $x = \sigma_{n_j}$. Due to the choice of time instants n_j , any sequence triggered by π is unambiguous. Furthermore, the probability that, for all j , at time instant n_j the current state is in Q_2^M , is at least $\frac{p}{2} \prod_{j \geq 0} \beta_j$. Thus with at least this probability, the random sequence in \mathcal{C} generated by π will never leave $Q_1^M \cup Q_2^M$, and will be a correct sequence. Putting all together, π ensures the safe diagnosis of cLTS \mathcal{C} .

Assume now that there is a strategy π in the cLTS, that renders it safe diagnosable. Observe that every decision rule for \mathcal{C} corresponds to a possibly randomized decision rule in M . For instance, if π chooses $\Sigma' = \{a, \bar{b}, u, o, f\}$ as set of enabled actions, then for a state in $Q_1^M \cup Q_2^M$ it corresponds to choosing with probability $\frac{1}{2}$ either a or \bar{b} (due to the definition of weights).

Using the preliminary observations, the set Ex of executions that contain infinitely often actions in $\{\bar{a}, \bar{b}\}$ or an occurrence of c , is exactly the set of unambiguous sequences and therefore has probability one under π . Let $\text{Ex}' \subseteq \text{Ex}$ be the subset of executions that contain infinitely often actions in $\{\bar{a}, \bar{b}\}$. By assumption on π , Ex' has positive probability and corresponds to executions that only visit states of $Q_1^M \cup Q_2^M$. Observe that such sequences visit infinitely often F_1 .

One builds a winning strategy π' for M as follows. Recall that the POMDP is blind, and therefore a strategy in M can only base its decisions on the events that have happened thus far, not on the current state. Thus, π' will be defined on $\{a, b\}^*$. Define $\text{proj}(x) = \text{proj}(\bar{x}) = x$ for $x \in \{a, b\}$. Let $\sigma = \sigma_1 \dots \sigma_n \in \{a, b\}^*$ be a finite sequence such that $\mathbb{P}_\pi^{q_0}(\mathcal{C} \models \text{proj}^{-1}(\sigma)) > 0$, and $\sigma' \in \text{proj}^{-1}(\sigma)$. Intuitively, with positive probability, σ is the projection of an execution following π . We further define $p_{\sigma'} = \mathbb{P}_\pi^{q_0}(\mathcal{C} \models \rho = \sigma' \mid \text{proj}(\rho) = \sigma)$ where ρ is a random sequence of length n . Then $\pi'(\sigma)$ is defined by $\pi'(\sigma) = \sum_{\sigma' \in \text{proj}^{-1}(\sigma)} p_{\sigma'} \pi(\sigma')$, *i.e.* the appropriate weighted combination on decision rules applied in the cLTS. For σ such that $\mathbb{P}_\pi^{q_0}(\mathcal{C} \models \text{proj}(\rho) = \sigma) = 0$, $\pi'(\sigma)$ is arbitrarily defined. By construction (and induction) one gets: $\mathbb{P}_{\pi'}^i(M \models \rho = \sigma) = \sum_{\sigma' \in \text{proj}^{-1}(\sigma)} \mathbb{P}_\pi^{q_0}(\mathcal{C} \models \rho = \sigma')$. Thus $\mathbb{P}_{\pi'}^i(M \models \text{proj}(\text{Ex}')) = \mathbb{P}_\pi^{q_0}(\mathcal{C} \models \text{Ex}')$ and $\mathbb{P}_{\pi'}^i(M \models \text{proj}(\text{Ex}')) > 0$. On the other hand, any sequence of $\text{proj}(\text{Ex}')$ visits infinitely often F . As a consequence, π' is a winning strategy. \square

A.3 Proof of Theorem 4

Proof. Here we solve the more general problem of the existence of a belief-based strategy in a POMDP that simultaneously achieves a Büchi condition F and a safety condition I . Let π be a belief-based strategy. We define a finite discrete-time Markov chain (DTMC) as follows.

- The set of states is $\{(q, B) \mid q \in Q \wedge q \in B \subseteq \text{Obs}(q)\}$.
- The initial state is $(q_0, \{q_0\})$.
- The transition matrix is defined by:
 $\mathbf{P}[(q, B), (q', B')] = \sum_{a \in \text{Act}} \pi(B)(a)T(q, a)(q')$
 if $B' = \{q'' \in \text{Obs}(q') \mid \exists q^* \in B \sum_{a \in \text{Act}} \pi(B)(a)T(q^*, a)(q'') > 0\}$
 and $\mathbf{P}[(q, B), (q', B')] = 0$ otherwise.

Assume that π is winning from the initial state q_0 for the almost-sure Büchi objective F and positive safe objective I . This property is equivalent to the fact that in the the underlying graph of the DTMC the following holds:

- every bottom strongly connected component (BSCC) reachable from $(q_0, \{q_0\})$ contains a state (q, B) with $B \subseteq F$,
- and there is a BSCC reachable from $(q_0, \{q_0\})$ by some path such that all the beliefs along this path and in the BSCC belong to I .

These properties only depend on the underlying graph of the DTMC which in turns only depends on the supports of the decisions of strategy π . Therefore a NEXPTIME procedure consists in guessing the supports of some belief-based strategy, building the underlying graph of the corresponding DTMC and checking the two previous properties.

The EXPTIME hardness has been proved in proposition 3. \square