

Diagnosis and Degradation Control for Probabilistic Systems

Nathalie Bertrand · Serge Haddad · Engel
Lefauchaux

Received: date / Accepted: date

Abstract Systems prone to faults are often equipped with a controller whose aim consists in restricting the behaviour of the system in order to perform a diagnosis. Such a task is called *active diagnosis*. However to avoid that the controller degrades the system in view of diagnosis, a second objective in terms of quality of service is usually assigned to the controller. In the framework of stochastic systems, a possible specification, called *safe active diagnosis* requires that the probability of correctness of the infinite (random) run is non null. We introduce and study here two alternative specifications that are in many contexts more realistic. The notion of (γ, v) -fault freeness associates with each run a value depending on the discounted length of its correct prefix where the discounting factor is γ . The controller has to ensure that the average of this value is above the threshold v . The notion of α -resiliency requires that asymptotically, at every time step, a proportion greater than α of correct runs remain correct. From a semantic point of view, we determine the equivalences and (non) implications between the three notions of degradations both for finite and infinite systems. From an algorithmic point of view, we establish the border between decidability and undecidability of the diagnosability problems. Furthermore in the positive case, we exhibit their precise complexity and propose a synthesis of the controller which may require an infinite memory.

Keywords Stochastic systems · Partial observation · Fault tolerance · Diagnosis

The work of serge Haddad was supported by the project ERC EQualIS (FP7-308087).

N. Bertrand
Univ Rennes, Inria, CNRS, IRISA
E-mail: nathalie.bertrand@inria.fr

S. Haddad
LSV, ENS Paris-Saclay, CNRS, Inria, Université Paris-Saclay
E-mail: haddad@lsv.fr

E. Lefauchaux
Univ Rennes, Inria, CNRS, IRISA
LSV, ENS Paris-Saclay, CNRS, Inria, Université Paris-Saclay
E-mail: engel.lefauchaux@inria.fr

1 Introduction

Diagnosis. The designer of a system aims at eliminating faults that could trigger unwanted behaviours. However, for embedded systems interacting with an unpredictable environment, the absence of faults is not a reasonable hypothesis. Thus diagnosis, whose goal consists to detect faults from the observations of the run of the system, is a crucial task. One of the approach frequently used to analyse *diagnosability* (i.e. the existence of a diagnoser) consists in modelling the system by a transition system whose states (depending on the internal part of the system) are unobservable and events may, depending on their nature, be observable or not. A diagnoser must fulfill two requirements: *correctness* and *reactivity*. A diagnoser is correct if it never erroneously claims a fault. It is reactive if every fault is announced after a finite delay. For finite systems, the diagnosability problem is decidable in polynomial time¹ while the synthesis of a diagnoser may require an exponential time [8].

Active diagnosis. Embedded systems are often equipped with one (or more) controller(s) in order to maintain some functionalities of the system in case of a pathological behaviour of the environment. It is thus tempting to add to the controller a diagnosis task. Formally some of the observable events are controllable and considering its current observation, the controller chooses which subset of actions should be allowed to make the system diagnosable. A system is said *actively diagnosable* if there exists a controller ensuring the role of diagnoser. In [11], the authors showed that the active diagnosability problem is decidable in doubly exponential time. Then in [7], the authors designed a single exponential time algorithm and proved the optimality of this complexity.

Probabilistic diagnosis. In transition systems, the unpredictable behaviours of the environment are modelled by a nondeterministic choice between the possible events from the current state. However, in order to quantify the risks induced by the faults of the systems, the designer often substitutes to the non deterministic choice by a random choice or equivalently by a weighted one. Then the model becomes a discrete time Markov chain in the *passive* case (i.e. without controller) and a weighted transition system in the *active* case (i.e. with a controller). The reactivity requirement is then adapted by requiring that *almost surely* (i.e. with probability 1) a fault is announced [12]. The passive probabilistic diagnosability is a PSPACE-complete problem [4] while the active probabilistic diagnosability is an EXPTIME-complete problem [2].

Active diagnosis and degradation. However the choices performed by the controller ensuring active diagnosis may have a pernicious effect: to detect faults, the controller sometimes could favour the occurrence of these faults! Aiming to manage the degradation of a system, a controller ensuring *safe active diagnosis* ensures the diagnosis task and a positive probability that an infinite run is correct. A quantitative version of

¹ In this paper, we assume some familiarity with basic complexity notions, and refer the interested reader to [9].

this requirement fixes a probabilistic threshold ε to achieve. Safe active probabilistic diagnosability is undecidable; however, when limited to finite memory controllers, the problem becomes decidable in NEXPTIME [2].

Contributions. Ensuring a positive probability for correct runs is only one possible way to express a requirement on the degradation control of a system and it is not necessarily appropriate for all contexts. For instance, some systems are designed to correctly behave for a long period of time at the end of which they will be replaced by a new system. In order to address such requirements, we introduce two new specifications of degradation control:

- A system is (γ, v) -fault free if, when applying a temporal discount $\gamma \leq 1$, the mean value of the discounted length of the maximal correct prefix of a run is greater or equal to v . The qualitative version of this specification, called *lasting fault freeness* is obtained for $\gamma = 1$ and $v = \infty$. This means that the average length of the maximal correct prefix of a run is infinite.
- A system is α -resilient for $\alpha < 1$ if the proportion of correct runs decreases asymptotically slower than a factor α at every time step. There are two qualitative versions of this specification: a system is strongly resilient (resp. weakly resilient) if for all $\alpha < 1$ (resp. there exists $\alpha < 1$ such that) it is α -resilient.

First we study these specifications in a passive framework. More precisely we focus on the qualitative notions. We establish that the safeness of a system implies its lasting fault freeness and its strong resiliency and that no other implication exists between the three notions. However they coincide for finite systems.

Then we analyse the active framework. We show that diagnosability combined with (γ, v) -fault freeness or with α -resiliency is undecidable. Afterward we improve the complexity result related to safe active diagnosis with finite memory showing that the problem is EXPTIME-complete. Contrary to safe active diagnosis, diagnosability combined with (1) lasting fault freeness, (2) strong resiliency or (3) weak resiliency, remains decidable and more precisely is EXPTIME-complete. Moreover, we establish that the additional constraints of lasting fault freeness and strong resiliency coincide in the active framework. Those decidability results are all the more surprising since the corresponding diagnosers may require infinite memory.

Organisation. In Section 2, we define the probabilistic transition systems and introduce diagnosis and the different specifications of the degradation of these systems. We also present the links between the qualitative versions. In Section 3, we establish the decidability status of the active diagnosability problems and, when decidable, their complexity. We then conclude and give perspectives of this work in Section 4.

2 Diagnosis and degradation of a probabilistic system

2.1 Probabilistic labelled transition system

We introduce a standard probabilistic model of discrete time event system based on discrete time Markov chains (see [1]).

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

- Q is a countable set of states with $q_0 \in Q$ being 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 a function from T to $\mathbb{Q}_{>0}$ verifying:

$$\forall q \in Q, \quad \sum_{(a,q') \in \Sigma \times Q \mid (q,a,q') \in T} \mathbf{P}[q, a, q'] = 1 .$$

A pLTS is a labelled transition system (LTS) enhanced by probabilities on the transitions. The transition relation of the induced LTS is defined by: $q \xrightarrow{a} q'$ for $(q, a, q') \in T$; such a transition is called *enabled* in state q . By definition, in every state q of the pLTS, at least one transition is enabled, *i.e.* a pLTS is *live*.

Notations. Given a countable set E , we denote $\text{Dist}(E)$ the set of probability distributions on E . Let $q \in Q$, the function associating $\mathbf{P}[q, a, q']$ with a pair (a, q') if $(q, a, q') \in T$ and 0 otherwise, is an element of $\text{Dist}(\Sigma \times Q)$. The support of a distribution $p \in \text{Dist}(E)$, written $\text{Supp}(p)$, is defined by $\text{Supp}(p) = \{e \in E \mid p(e) > 0\}$. Thus The support of the above distribution is $\{(a, q') \mid (q, a, q') \in T\}$. When the support of a distribution is a singleton $\{e\}$, we denote this Dirac distribution $\mathbf{1}_e$.




Fig. 1 An example of (finite) pLTS.

Example 1 A pLTS is represented by a labelled oriented graph whose vertices are the states and edges are the transitions labelled by the associated event and its probability. In Figure 1, the edge from q_0 to f_1 is triggered by the event f with probability $\frac{2}{3}$. We will often omit the probabilities when they are equal to 1, and, more generally, when the distribution on the transitions exiting a state is uniform.

We now introduce some important notions and notations used in the sequel. A *run* ρ of a pLTS \mathcal{A} is a (finite or infinite) sequence $\rho = q_0 a_0 q_1 \dots$ such that for all $i \geq 0$, $q_i \in Q$, $a_i \in \Sigma$ and when q_{i+1} is defined, $q_i \xrightarrow{a_i} q_{i+1}$. The notion of a run may be generalised by allowing to start in an arbitrary state q . We write Ω for the set of infinite runs starting in q_0 , assuming the pLTS \mathcal{A} is clear from context. A finite run ρ ends in a state denoted $\text{last}(\rho)$ and its *length*, denoted $|\rho|$, is the number of events in ρ . Let $\rho = q_0 a_0 q_1 \dots q_n$ be a finite run and $\rho' = q_n a_n q_{n+1} \dots$ a (finite or infinite) run starting in the last state of ρ , we call *concatenation* of ρ and ρ' the run

$\rho\rho' = q_0a_0q_1 \dots q_n a_n q_{n+1} \dots$. The run ρ is called a *prefix* of ρ' , which we will write $\rho \preceq \rho'$, if there exists another run ρ'' such that $\rho' = \rho\rho''$. The *cylinder* generated by a finite run ρ is the set of infinite runs extending ρ : $\text{Cyl}(\rho) = \{\rho' \in \Omega \mid \rho \preceq \rho'\}$. The sequence associated with $\rho = qa_0q_1 \dots$ is the word $w_\rho = a_0a_1 \dots$, and we write $q \xrightarrow{\rho} q'$ or $q \xrightarrow{w_\rho} q'$ (resp. $q \xrightarrow{\rho} q'$ or $q \xrightarrow{w_\rho} q'$) for an infinite (resp. finite) run ρ . A state q is *reachable* (from the initial state q_0) if there exists a run ρ such that $q_0 \xrightarrow{\rho} q$, also written $q_0 \Rightarrow q$. The language of a pLTS \mathcal{A} is the set of infinite words labelling runs of \mathcal{A} and is formally defined by $\mathcal{L}^\omega(\mathcal{A}) = \{w \in \Sigma^\omega \mid \exists q_0 \xrightarrow{w}\}$.

Forgetting the labels and merging (and adding up the probabilities) the transitions with same source and destination, a pLTS becomes a discrete time Markov chain (DTMC). In a DTMC, the set of infinite runs of \mathcal{A} is the support of a probability measure extended from the probabilities of the cylinders by the Caratheodory's extension theorem:

$$\mathbb{P}_{\mathcal{A}}(\text{Cyl}(q_0a_0q_1 \dots q_n)) = \mathbf{P}[q_0, a_0, q_1] \dots \mathbf{P}[q_{n-1}, a_{n-1}, q_n] .$$

When \mathcal{A} is fixed, we will often omit the subscript, and write \mathbb{P} for $\mathbb{P}_{\mathcal{A}}$. Let ρ be a finite run, with a small abuse of notation we write $\mathbb{P}(\rho)$ for $\mathbb{P}(\text{Cyl}(\rho))$. If R is a countable set of finite runs such that no run is prefix of another, we write $\mathbb{P}(R)$ for $\sum_{\rho \in R} \mathbb{P}(\rho)$ which is consistent as the intersections of the associated cylinders are empty.

2.2 Partial observation and ambiguity

In order to formalise the problems related to fault diagnosis, we partition the set of events Σ in two subsets Σ_o and Σ_u , the *observable* events and *unobservable* ones, respectively. Moreover, we distinguish a special event, the *fault* $\mathbf{f} \in \Sigma_u$.

Example 2 The set of events of the pLTS of Figure 1 is defined by $\Sigma_o = \{a, b\}$ and $\Sigma_u = \{\mathbf{f}, u\}$. Transitions labelled by unobservable events are represented by dashed edges.

Let w be a finite word on the alphabet Σ , its length is denoted by $|w|$ and $\mathbf{1}$ represents the empty string. The projection of words of Σ^* on the observable events is inductively defined by: $\pi(\mathbf{1}) = \mathbf{1}$, for $a \in \Sigma_o$, $\pi(wa) = \pi(w)a$ and for $a \in \Sigma_u$, $\pi(wa) = \pi(w)$. We write $|w|_o$ for the *observable length* of w , i.e. $|\pi(w)|$. When w is an infinite word on Σ , its projection is the limit of the projections of its finite prefixes, and by convention $|w| = \infty$. A pLTS \mathcal{A} is called *convergent* with respect to a partition $\Sigma = \Sigma_o \uplus \Sigma_u$ if, from every reachable state, there is no infinite sequence of unobservable events: $\mathcal{L}^\omega(\mathcal{A}) \cap \Sigma^* \Sigma_u^\omega = \emptyset$. When \mathcal{A} is convergent, for all $w \in \mathcal{L}^\omega(\mathcal{A})$, $\pi(w) \in \Sigma_o^\omega$. In the sequel, we assume that the pLTS are convergent. We use the terminology *sequence* for a word $w \in \Sigma^* \cup \Sigma^\omega$, and *observed sequence* for a word $w \in \Sigma_o^* \cup \Sigma_o^\omega$. The projection of a sequence is therefore an observed sequence.

The observable length of a run ρ denoted $|\rho|_o \in \mathbb{N} \cup \{\infty\}$, is the number of occurrences of observable events: $|\rho|_o = |w_\rho|_o$. A *signalling run* is a finite run

$q_0 a_0 q_1 \cdots a_{n-1} q_n$ such that a_{n-1} is an observable event. The signalling runs are precisely the relevant runs from the point of view of partial observation as every observable event gives an additional information on the run to an external observer. In the following, SR denote the set of signalling run and SR_n the set of signalling runs of observable length n . Since the pLTS are convergent, for all $n > 0$, SR_n is equipped with a probability distribution defined by assigning the measure $\mathbb{P}(\rho)$ to every $\rho \in \text{SR}_n$. By convention the empty run q_0 is defined as the single run of length 0. Let $w \in \Sigma_o^*$ be an observed sequence, we define its cylinder $\text{Cyl}(w) = w \Sigma_o^\omega$ and the associated probability $\mathbb{P}(\text{Cyl}(w)) = \mathbb{P}(\{\rho \in \text{SR}_{|w|} \mid \pi(\rho) = w\})$, often abbreviated by $\mathbb{P}(w)$.

We now classify the runs depending on the occurrence of faults. A run ρ is *faulty* if the associated sequence w_ρ contains \mathbf{f} , otherwise it is *correct*. Let $n \in \mathbb{N}$, we write F_n (resp. C_n) for the set of infinite runs such that the signalling prefix of observable length n is faulty (resp. correct). We define the sets of finite (resp. infinite) faulty and correct signalling runs F (resp. F_∞) and C (resp. C_∞). Without loss of generality, by considering two copies of every state of the pLTS, we suppose that the state space Q of \mathcal{A} is partitioned between correct and faulty states: $Q = Q_f \uplus Q_c$ such that the faulty (resp. correct) states are only reachable by faulty (resp. correct) runs. An infinite (resp. finite) observed sequence $w \in \Sigma_o^\omega$ (resp. Σ_o^*) is *ambiguous* if there exists an infinite correct run (a correct signalling run) ρ and an infinite faulty run (a faulty signalling run) ρ' such that $\pi(\rho) = \pi(\rho') = w$. Otherwise it is either *surely faulty* or *surely correct* depending whether $\pi^{-1}(w) \cap \text{SR} \subseteq F$ or $\pi^{-1}(w) \cap \text{SR} \subseteq C$. A run is ambiguous, surely correct or surely faulty if its observed sequence is ambiguous, surely correct or surely faulty respectively.

Example 3 Consider the pLTS of Figure 1. The correct states are q_0 and q_1 while the faulty states are f_1 and f_2 . The run $\rho_f = q_0 \mathbf{f} (f_1 a)^\omega$ is faulty and ambiguous as the single correct run $\rho_c = q_0 u (q_1 a)^\omega$ has the same observed sequence a^ω . For every n , the finite sequence a^n is ambiguous while the sequence $a^n b$ is surely faulty as b does not occur in ρ_c .

2.3 Diagnosability

Diagnosability of a pLTS is defined in terms of probability of sets of runs. Toward this goal, we define FAmb_∞ the set of infinite faulty ambiguous runs.

Definition 2 (Diagnosability) A pLTS \mathcal{A} is *diagnosable* if $\mathbb{P}(\text{FAmb}_\infty) = 0$.

Let n be an integer, FAmb_n is the set of infinite runs of \mathcal{A} which signalling prefix of observable length n is faulty and ambiguous. We recall the following result which allows us to use an alternative definition of diagnosability.

Lemma 1 ([4]) *Let \mathcal{A} be a pLTS. Then $\lim_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_\infty \setminus \text{FAmb}_n) = 0$. Moreover, if \mathcal{A} is finitely branching, then $\lim_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n \setminus \text{FAmb}_\infty) = 0$ and consequently \mathcal{A} is diagnosable iff $\limsup_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n) = 0$.*

The alternative definition of diagnosability given by Lemma 1 is used implicitly multiple times throughout this paper. Among other things, it allows to synthesise a diagnoser (with infinite memory) in a simple way when the (finite) pLTS is diagnosable. After an observed sequence w , the diagnoser claims a fault if w is surely faulty. By construction, the diagnoser is correct and since $\limsup_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n) = 0$, it is reactive. In fact, we can build a diagnoser using finite memory by only remembering the current possible states and claim a fault when all these states are faulty [12].

Example 4 Consider the pLTS of Figure 1. FAmb_∞ is a singleton reduced to the run $\rho_f = q_0 \mathbf{f}(f_1 a)^\omega$ with a null probability. Thus this pLTS is diagnosable. $\text{FAmb}_n = \text{Cyl}(q_0 \mathbf{f}(f_1 a)^n f_1) \cup \text{Cyl}(q_0 \mathbf{f}(f_1 a)^n f_2) = \text{Cyl}(q_0 \mathbf{f}(f_1 a)^{n-1} f_1)$. The probability of FAmb_n is thus equal to $\frac{2^{-n+2}}{3}$ and converges to 0 as announced by the previous lemma. In this particular case, the diagnoser does not require any memory and claims a fault at the first occurrence of event b .

2.4 Degradation

We describe and study here three notions of degradation of a system: safeness, fault freeness and resiliency. A pLTS is *safe* if it guarantees a positive probability of infinite correct runs. We can refine this notion by quantifying it: a pLTS is ε -safe if this probability is greater or equal than ε .

Definition 3 (Safe pLTS) Let \mathcal{A} be a pLTS.

- For $\varepsilon > 0$, \mathcal{A} is ε -safe if $\mathbb{P}(\text{C}_\infty) \geq \varepsilon$;
- \mathcal{A} is *safe* if $\mathbb{P}(\text{C}_\infty) > 0$.

As pointed out in the introduction, in some cases, safeness is a too strong requirement. We formalise now two alternatives: fault freeness and resiliency. Fault freeness aims at quantifying the period of time during which the pLTS is correct. In order to (possibly) take into account the importance of the immediate future, we introduce a discount factor $\gamma \leq 1$ for counting this duration. The expectation of this discounted value is then compared to a threshold v .

Definition 4 (Fault free pLTS) Let \mathcal{A} be a pLTS.

- For $0 < \gamma \leq 1$ and $v \in [0, \infty]$, \mathcal{A} is (γ, v) -*fault free* if $\sum_{n \geq 1} \mathbb{P}(\text{C}_n) \gamma^n \geq v$.
- \mathcal{A} is *lasting fault free* if it is $(1, \infty)$ -fault free.

Observe that when γ equals 1, $\sum_{n \geq 1} \mathbb{P}(\text{C}_n) \gamma^n$ is the mean observable length of the maximal correct signalling prefix of a random run. This justifies the name *lasting fault free* for an infinite expectation.

The notion of resiliency is an alternative measure of degradation based on a factor degradation ratio per time unit $\alpha < 1$. A pLTS is α -*resilient* if the proportion of finite correct runs which stays correct on the next occurrence of an observable event is asymptotically greater than α . This requirement has two qualitative variants: strong resiliency (resp. weak resiliency) requires α -resiliency for every (resp. for at least one) $\alpha < 1$.

Definition 5 (Resilient pLTS) Let \mathcal{A} be a pLTS.

- For $0 < \alpha < 1$, \mathcal{A} is α -resilient if $\limsup_{n \rightarrow \infty} \frac{\alpha^n}{\mathbb{P}(C_n)} = 0$;
- \mathcal{A} is *strongly resilient* if for all $0 < \alpha < 1$, \mathcal{A} is α -resilient;
- \mathcal{A} is *weakly resilient* if there exists $0 < \alpha < 1$ such that \mathcal{A} is α -resilient.




Fig. 2 An example of infinite pLTS.

Example 5 The pLTS \mathcal{A} of Figure 2 has a single correct run $\rho = q_0 a q_1 a q_2 \dots$ while every faulty run contains an infinite number of b . \mathcal{A} is thus diagnosable. Moreover, the probability of ρ is equal to $\prod_{n \geq 1} p_n$ and the probability of its prefix of length n is equal to $r_n = \prod_{i \leq n} p_i$. Consequently, \mathcal{A} is safe iff $\lim_{n \rightarrow \infty} r_n > 0$. By direct application of the definition, \mathcal{A} is lasting fault free iff $\sum_{n \geq 1} r_n = \infty$. Let us consider different values of $(p_i)_{i \in \mathbb{N}}$.

- Let $p_i = \frac{i}{i+1}$. Then $r_n = \frac{1}{n+1}$. Thus \mathcal{A} is not safe but is lasting fault free. For every $\alpha < 1$, $\lim_{n \rightarrow \infty} (n+1)\alpha^n = 0$. Thus \mathcal{A} is also strongly resilient.
- Let $p_i = \frac{i^2}{(i+1)^2}$. Then $r_n = \frac{1}{(n+1)^2}$. Thus \mathcal{A} is neither safe nor lasting fault free. For every $\alpha < 1$, $\lim_{n \rightarrow \infty} (n+1)^2 \alpha^n = 0$. Thus \mathcal{A} is strongly resilient.
- We inductively define two sequences m_k and n_k by:

$$n_k = 2^{\sum_{j < k} m_j} \text{ (hence } n_0 = 1) \text{ and } m_k = n_k + \sum_{j < k} m_j + n_j.$$

Define:

- $I_k = [n_k + \sum_{j < k} m_j + n_j, \sum_{j \leq k} m_j + n_j[$;
- $J_k = [\sum_{j \leq k} m_j + n_j, n_{k+1} + \sum_{i \leq k} m_i + n_i[$.

When $i \in I_k$, $p_i = \frac{1}{2}$. When $i \in J_k$, $p_i = 1$.

Observe that for all $n \in J_k$, $r_n = 2^{-\sum_{j \leq k} m_j}$.

Consequently $\sum_{n \geq 1} r_n \geq \sum_{k \geq 0} \sum_{n \in J_k} r_n = \sum_{k \geq 0} 2^{\sum_{j \leq k} m_j} 2^{-\sum_{j \leq k} m_j} = \infty$.

Thus \mathcal{A} is lasting fault free.

Let $n = \sum_{j \leq k} m_j + n_j$. Consequently, $r_n = 2^{-\sum_{j \leq k} m_j}$. Fix $\alpha = \frac{1}{\sqrt{2}}$.

$$\frac{\alpha^n}{r_n} = 2^{\sum_{j \leq k} m_j} (\sqrt{2})^{-\sum_{j \leq k} m_j + n_j} \geq 2^{m_k} (\sqrt{2})^{-2m_k} = 1.$$

Therefore \mathcal{A} is not α -resilient.

The next theorem establishes the relationships between the qualitative versions of the three degradation notions. Note that the pLTS from Example 5 serves as witness for the last two statements.

Theorem 1 *Let \mathcal{A} be a pLTS.*

- (a) *If \mathcal{A} is safe then \mathcal{A} is lasting fault free and strongly resilient;*
- (b) *Assuming \mathcal{A} is finite, \mathcal{A} is safe iff \mathcal{A} is lasting fault free iff \mathcal{A} is strongly resilient;*
- (c) *There exists a lasting fault free pLTS that is not strongly resilient;*
- (d) *There exists a strongly resilient pLTS that is not lasting fault free.*

Proof We first prove item (a). Assume \mathcal{A} is a safe pLTS. There exists $\varepsilon > 0$ such that for all n , $\mathbb{P}(C_n) \geq \varepsilon$. Thus, $\sum_{n \geq 1} \mathbb{P}(C_n) \geq \sum_{n \geq 1} \varepsilon = \infty$. On the other hand, for all $\alpha < 1$, $\lim_{n \rightarrow \infty} \frac{\alpha^n}{\mathbb{P}(C_n)} \leq \lim_{n \rightarrow \infty} \frac{\alpha^n}{\varepsilon} = 0$. We conclude that \mathcal{A} is lasting fault free and strongly resilient.

To prove item (b), we pick \mathcal{A} a finite pLTS. Observe that every bottom strongly connected component (BSCC) of \mathcal{A} (here seen as a graph) either contains only correct states or contains only faulty states. Accordingly, we can refer to them as faulty BSCC or correct BSCC. As \mathcal{A} is a finite Markov chain (with events labelling the transitions), almost surely an infinite run reaches a BSCC and the mean time to reach a BSCC is finite. Due to the first result, \mathcal{A} is safe iff there exists a correct BSCC that is reachable from the initial state.

Suppose that \mathcal{A} is not safe.

- Every reachable BSCC is faulty, and this implies that the mean time to reach a faulty BSCC is finite. This mean time is an upper bound on the mean observable length of the maximal signalling prefix of a run. Thus \mathcal{A} is not lasting fault free.
- We write $m = |Q|$. For all $q \in Q_c$, there exists ρ_q a run starting in q composed of an elementary run from q to a faulty BSCC followed by an elementary run (or circuit) in the BSCC of which only the last event is observable (by convergence). This run has an observable length smaller or equal to m . We note μ_q , the probability of that run and $\mu = \min_{q \in Q_c} \mu_q$. Consider a signalling run ρ of observable length n for an arbitrary n and ending in $q \in Q_c$. From the existence of ρ_q , $\mathbb{P}(\{\rho' \in \text{SR}_{n+m} \cap C \mid \rho \leq \rho'\}) \leq (1 - \mu)\mathbb{P}(\rho)$. Thus $\mathbb{P}(C_{n+m}) \leq (1 - \mu)\mathbb{P}(C_n)$. So, $\mathbb{P}(C_n) \in O((1 - \mu)^{\frac{n}{m}})$. Choosing $\alpha = (1 - \mu)^{\frac{1}{m}}$, \mathcal{A} is not α -resilient and thus not strongly resilient.

This concludes the proof. \square

3 Control and diagnosis

3.1 Active diagnosis

The extension of the pLTS formalism allowing to express control requires to fix at least two features of this formalism: the nature of the control and the distribution of probabilities of the controlled system. *Controllable Labelled Transition System* (CLTS) were introduced in [2]. In order to specify the control, a subset of observable events is considered as controllable. The control strategy forbids a subset of controllable events depending on the sequence of observations it has received so far. In particular it cannot change its control in between two observations. The transitions of the system are no longer labelled by (rational) probabilities but rather by (integer)

weights which represent their relative probabilities. Given a state and a set of forbidden controllable actions, in order to obtain a probability distribution on the allowed transitions, the weights of the outgoing transitions labelled by uncontrollable or allowed controllable actions are normalised. Provided that the control strategy does not create any deadlock, the so-obtained controlled obtained is a pLTS.

Definition 6 (CLTS) A *Controllable Labelled Transition System (CLTS)* is a tuple $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ where:

- Q is a set of states with an initial state $q_0 \in Q$;
- $\Sigma = \Sigma_o \uplus \Sigma_u$ is a finite state of events partitioned into the set of observable events Σ_o containing controllable events $\Sigma_c \subseteq \Sigma_o$ and the set of unobservable events Σ_u containing the fault \mathbf{f} ;
- $T : Q \times \Sigma \times Q \rightarrow \mathbb{N}$ is the transition function that associates an integer weight with each transition.

For CLTS, the transition function T simultaneously plays the role of the probability function and the transition function in pLTS. We use weights instead of probabilities in cLTS, since due to the control normalizing the weights is anyway necessary.

A CLTS has an induced transition system which transition relation is defined by $q \xrightarrow{a} q'$ if $T(q, a, q') > 0$. the extended relation \Rightarrow is defined as for pLTS. As for pLTS, we assume that the CLTS is convergent and live (i.e. $\forall q \exists q' \xrightarrow{a} q'$).

Example 6 A CLTS \mathcal{C} is represented in Figure 3. The weights of the transitions, all equal to 1, are omitted. The only controllable event of \mathcal{C} is b . The observable yet uncontrollable transitions (here on event a) are bold.




Fig. 3 An example of CLTS.

We now formalise the ingredients necessary to define the control of the CLTS. Let $\Sigma^\bullet \subset \Sigma$ and $q \in Q$, let us write $G^{\Sigma^\bullet}(q)$ for the sum of the weights of the transitions

exiting q and labelled by an event of Σ^\bullet . Using this sum, we define a normalisation of the transition relation restricted to the events of Σ^\bullet by:

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

A *strategy* of a CLTS \mathcal{C} is a function $\sigma : \Sigma_o^* \rightarrow \text{Dist}(2^\Sigma)$ such that for all $w \in \Sigma_o^*$ and all $\Sigma^\bullet \in \text{Supp}(\sigma(w))$, $\Sigma \setminus \Sigma_c \subseteq \Sigma^\bullet$. Given an observation, a strategy consists in (randomly) choosing a subset of allowed events containing the uncontrollable events. Let \mathcal{C} be a CLTS and σ be a strategy, we consider the *configurations* of the form $(w, q, \Sigma^\bullet) \in \Sigma_o^* \times Q \times 2^\Sigma$ with w the observed sequence, q the current state and Σ^\bullet the set of allowed events by σ after observation of w . We inductively define the set $\text{Reach}_\sigma(\mathcal{C})$ of the reachable configurations under σ by:

- for all $\Sigma^\bullet \in \text{Supp}(\sigma(\mathbf{1}))$, we have $(\mathbf{1}, q_0, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$;
- for all $(w, q, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$ and all $a \in \Sigma_u \cap \Sigma^\bullet$ such that $q \xrightarrow{a} q'$, we have $(w, q', \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$, and the corresponding transition is denoted by $(w, q, \Sigma^\bullet) \xrightarrow{a}_\sigma (w, q', \Sigma^\bullet)$;
- for all $(w, q, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$, all $a \in \Sigma_o \cap \Sigma^\bullet$ such that $q \xrightarrow{a} q'$ and all $\Sigma^{\bullet'} \in \text{Supp}(\sigma(wa))$, we have $(wa, q', \Sigma^{\bullet'}) \in \text{Reach}_\sigma(\mathcal{C})$, and the corresponding transition is denoted by $(w, q, \Sigma^\bullet) \xrightarrow{a}_\sigma (wa, q', \Sigma^{\bullet'})$.

A strategy σ is called *live* if for every configuration $(w, q, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$, we have $G^{\Sigma^\bullet}(q) \neq 0$. Only the live strategies are relevant as the other strategies create deadlocks. We are now in position to introduce the semantic of a CLTS controlled by a live strategy σ in terms of a pLTS. Its set of states is $\text{Reach}_\sigma(\mathcal{C})$ augmented by an initial state whose goal is to randomly choose in accordance with $\sigma(\mathbf{1})$ the initial control. The probability distributions are based on T^{Σ^\bullet} if the current control is Σ^\bullet combined with the random choice of σ in case of an observable event occurrence.

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

- $Q_\sigma = \{q_{0\sigma}\} \cup \text{Reach}_\sigma(\mathcal{C})$;
- for all $(\mathbf{1}, q_0, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$, $(q_{0\sigma}, u, (\mathbf{1}, q_0, \Sigma^\bullet)) \in T_\sigma$;
- for all $(w, q, \Sigma^\bullet), (w', q', \Sigma^{\bullet'}) \in \text{Reach}_\sigma(\mathcal{C})$,
 $((w, q, \Sigma^\bullet), a, (w', q', \Sigma^{\bullet'})) \in T_\sigma$ iff $(w, q, \Sigma^\bullet) \xrightarrow{a}_\sigma (w', q', \Sigma^{\bullet'})$;
- for all $(\mathbf{1}, q_0, \Sigma^\bullet) \in \text{Reach}_\sigma(\mathcal{C})$, $\mathbf{P}_\sigma(q_{0\sigma}, u, (\mathbf{1}, q_0, \Sigma^\bullet)) = \sigma(\mathbf{1})(\Sigma^\bullet)$;
- for all $((w, q, \Sigma^\bullet), a, (w, q', \Sigma^\bullet)) \in T_\sigma$ and all $a \in \Sigma_u \cap \Sigma^\bullet$,
 $\mathbf{P}_\sigma((w, q, \Sigma^\bullet), a, (w, q', \Sigma^\bullet)) = T^{\Sigma^\bullet}(q, a, q')$;
- for all $((w, q, \Sigma^\bullet), a, (wa, q', \Sigma^{\bullet'})) \in T_\sigma$ and all $a \in \Sigma_o \cap \Sigma^\bullet$,
 $\mathbf{P}_\sigma((w, q, \Sigma^\bullet), a, (wa, q', \Sigma^{\bullet'})) = T^{\Sigma^\bullet}(q, a, q') \cdot \sigma(wa)(\Sigma^{\bullet'})$.

Example 7 Consider the CLTS \mathcal{C} depicted in Figure 3. There are two possible allowed subsets Σ and $\Sigma \setminus \{b\}$ that we denote Σ^- . Let us define the strategy σ by $\sigma(a^n) = p_n \cdot \Sigma^- + r_n \cdot \Sigma$ with $p_n + r_n = 1$ for all $n \in \mathbb{N}$ and $\sigma(w) = \mathbf{1}_\Sigma$ for every other w . A part of the pLTS \mathcal{C}_σ is represented in Figure 4. Let us develop the

distribution of probabilities exiting the configuration $(\mathbf{1}, q_1, \Sigma)$. The two transitions exiting q_1 are enabled with equal relative probabilities, thus normalised to 0.5. Since a and b are observable, the new control is chosen, in the case where a a is observed, by a probabilistic choice $p_1 \cdot \Sigma^- + r_1 \cdot \Sigma$ while if a b is observed, there is a deterministic choice $\mathbf{1}_\Sigma$. This result in three transitions with probability $0.5p_1$, $0.5r_1$ and 0.5 respectively.

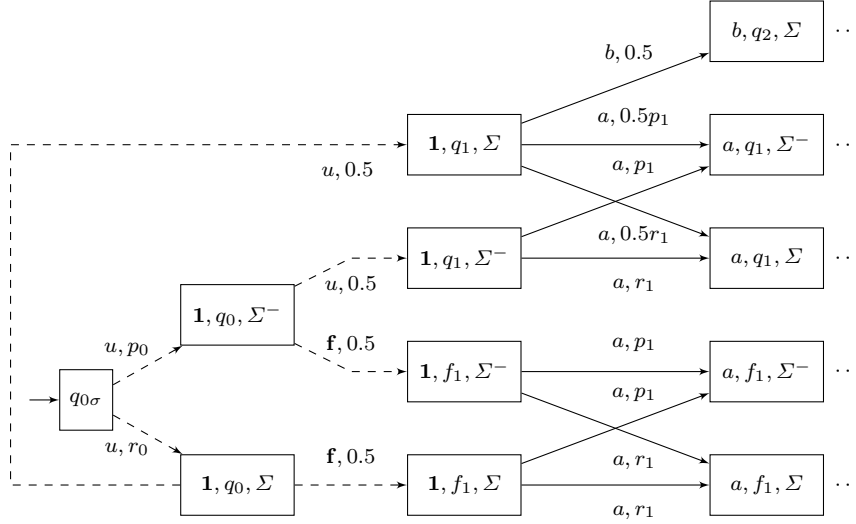


Fig. 4 An example of controlled CLTS.

Let us define the problems of active diagnosis in the context of the degradation control. Roughly speaking, given a CLTS, one asks whether there exists a strategy such that the associated pLTS is diagnosable and satisfies the required property related to degradation. We distinguish, as usually done, the quantitative problems (*i.e.* including numerical values) and the qualitative ones (such as safety, lasting fault free and strong/weak resiliency).

Definition 8 (Quantitative problems) Given a CLTS \mathcal{C} , $0 < \varepsilon, \alpha < 1$, $0 < \gamma \leq 1$ and $v \in [0, \infty]$:

- The ε -safe active diagnosis problem consists in deciding whether there exists a strategy σ such that \mathcal{C}_σ is diagnosable and ε -safe;
- The (γ, v) fault free active diagnosis problem consists in deciding whether there exists a strategy σ such that \mathcal{C}_σ is diagnosable and (γ, v) fault free;
- The α -resilient active diagnosis problem consists in deciding whether there exists a strategy σ such that \mathcal{C}_σ is diagnosable and α -resilient.

Definition 9 (Qualitative problems) Given a CLTS \mathcal{C} :

- The *safe active diagnosis* problem consists in deciding if there exists a strategy σ such that \mathcal{C}_σ is diagnosable and safe;
- The *lasting fault free active diagnosis* problem consists in deciding if there exists a strategy σ such that \mathcal{C}_σ is diagnosable and lasting fault free;
- The *strongly resilient active diagnosis* problem consists in deciding if there exists a strategy σ such that \mathcal{C}_σ is diagnosable and strongly resilient;
- The *weakly resilient active diagnosis* problem consists in deciding if there exists a strategy σ such that \mathcal{C}_σ is diagnosable and weakly resilient.

Example 8 In order to illustrate the impact of taking into account infinite memory strategies, let us examine the CLTS \mathcal{C} of Figure 3. The only ambiguous observed sequence is a^ω . A strategy σ thus makes it diagnosable iff the probability of this observed sequence in \mathcal{C}_σ is 0. However, the only correct run is $\rho = q_0u(q_1a)^\omega$ with observation a^ω . Thus, \mathcal{C} is not actively safely diagnosable.

Denoting as previously by p_n the probability to forbid b after the observed sequence a^n given by the strategy σ . Then $\mathbb{P}_{\mathcal{C}_\sigma}(q_0u(aq_1)^n) = \frac{1}{2} \prod_{i \leq n} \frac{1+p_i}{2}$. Thus, by choosing $p_n = 1 - \frac{1}{n+1}$, \mathcal{C}_σ is diagnosable, lasting fault free and strongly resilient. On the other hand, no finite memory strategy could achieve this goal since otherwise by Theorem 1, \mathcal{C} would be actively safely diagnosable.

3.2 Undecidability of the Quantitative Problems

The quantitative problems related to fault freeness and resiliency turn out to be undecidable. The proofs of these results are obtained by reductions from undecidable problems for *probabilistic automata*, a well-studied model that combines probability, control and partial observation, see *e.g.* [10]. A probabilistic automaton is a finite automaton equipped with a probability distribution on the transitions exiting a given state and labelled by a given letter. Given a finite word, we obtain a distribution on the paths labelled by this word and the *acceptance probability* of this word is the probability of the subset of these paths ending in an accepting state. More formally, a probabilistic automaton $\mathcal{M} = \langle S, s_0, F, \Sigma, \mathbf{P} \rangle$ is defined by:

- S , a finite set of states with $s_0 \in S$ the initial state and $F \subseteq S$ the subset of final states;
- Σ , a finite alphabet;
- \mathbf{P} a matrix $S \times \Sigma \times S$ with rational non-negative coefficients such that for all $s \in S$ and all $a \in \Sigma$, $\sum_{s' \in S} \mathbf{P}(s, a, s') = 1$.

The acceptance probability of a word $w = w_1 \dots w_n$, $\text{val}_{\mathcal{M}}(w)$ is defined by:

$$\text{val}_{\mathcal{M}}(w) = \sum_{s_1, \dots, s_n | s_n \in F} \prod_{i=0}^{n-1} \mathbf{P}(s_i, w_{i+1}, s_{i+1}).$$

Let $0 < \theta < 1$ be an arbitrary threshold. Given \mathcal{M} a probabilistic automaton, the problems of deciding the existence of a word w such that (1) $\text{val}_{\mathcal{M}}(w) \geq \theta$ or (2) $\text{val}_{\mathcal{M}}(w) > \theta$ are undecidable [6]. In the following reductions, we choose $\theta = \frac{1}{2}$.

Before developping it, we give a sketch of proof of the next proposition. Given a probabilistic automaton \mathcal{M} with alphabet Σ , one builds a CLTS \mathcal{C} composed of two independent parts each one initially entered with probability $\frac{1}{2}$ by an unobservable transition. The unobservable event leading to the first part is the fault \mathbf{f} which can only be detected almost surely if the observable event $\sharp \notin \Sigma$ occurs with probability 1. The second part is constituted of a CLTS version of \mathcal{M} augmented by exiting transitions. One exits \mathcal{M} with probability $\frac{1}{2}$ at every step toward a faulty sub-part except if the \sharp event is triggered. In this case, if the system was in a final state of \mathcal{M} it goes back to the initial state of the automaton instead of performing a fault. If there exists a word w with an acceptance probability at least $\frac{1}{2}$, the strategy which consists in forcing the observed sequence $w\sharp$ as long as the run stays in \mathcal{M} ensures an average observable length (without discount) of the maximal correct signalling prefix greater or equal to 1. In the opposite case, we show that no strategy can achieve this threshold.

Proposition 1 (γ, v) -fault free active diagnosability is undecidable.

Proof We proceed here by reduction from the problem of the existence of a word w such that $\text{val}_{\mathcal{M}}(w) \geq \frac{1}{2}$. We consider the probabilistic automaton $\mathcal{M} = \langle S, s_0, F, \Sigma, \mathbf{P} \rangle$ for which w.l.o.g. we assume that: (1) $\Sigma \cap \{u, \mathbf{f}, \sharp, \natural\} = \emptyset$ and (2) the probabilities are fractions $\frac{n}{d}$ with fixed denominator d . One builds the CLTS $\mathcal{C} = \langle Q, q_0, \Sigma', T \rangle$ described in Figure 5 and defined by:

- $Q = S \cup \{q_0, q_c^1, q_c^2, q_c^3, f_1, f_2\}$;
- $\Sigma' = \Sigma \cup \{\mathbf{f}, u, \sharp, \natural\}$, $\Sigma_u = \{\mathbf{f}, u\}$ and $\Sigma_c = \Sigma \cup \{\sharp\}$;
- the transition function T is defined as follows.
 1. $T(q_0, \mathbf{f}, f_1) = T(q_0, u, s_0) = T(q_c^1, \sharp, q_c^3) = T(q_c^3, \sharp, q_c^3) = T(q_c^3, \mathbf{f}, f_2) = T(q_c^2, \mathbf{f}, f_2) = T(f_2, \natural, f_2) = T(f_1, \sharp, f_2) = 1$;
 2. for every $a \in \Sigma$, $T(f_1, a, f_1) = 1$;
 3. for every $s, s' \in S$ and every $a \in \Sigma$, $T(s, a, s') = d \cdot \mathbf{P}(s, a, s')$ and $T(s, a, q_c^2) = d$;
 4. for every $s \in F$, $T(s, \sharp, q_c^1) = 1$ and for every $s \in S \setminus F$, $T(s, \sharp, q_c^2) = 1$;
 5. for every other triplet, T is equal to 0.

As detailed above, the probabilities in \mathcal{M} are all multiplied by their common d , to obtain integer weights, and we write $d \cdot \mathcal{M}$ in the figure to represent this scaling.

Let us show that \mathcal{A} is $(1, 1)$ -fault free iff there exists a word w accepted in \mathcal{M} with probability at least $\frac{1}{2}$.

Let σ be an arbitrary strategy, \mathcal{C}_σ is diagnosable iff \natural occurs almost surely in a run. Indeed an observed sequence $w \in \Sigma^*$ is ambiguous. On the other hand every run ρ leaving $S \cup \{f_1\}$ almost surely reaches f_2 where \natural occurs and, whatever ρ , a fault has occurred.

• Assume that there exists $w = w_1 \dots w_k \in \Sigma^*$ such that $\text{val}_{\mathcal{M}}(w) \geq \frac{1}{2}$. We define the deterministic strategy σ by:

- $\sigma(w) = \{\mathbf{f}, u, \sharp, \natural\}$;
- for all $0 \leq i < k$, $\sigma(w_1 \dots w_i) = \{\mathbf{f}, u, w_{i+1}, \natural\}$;
- $\sigma(w') = \Sigma'$ for any other word w' .

◦ A correct run ρ of observed sequence $w\sharp$ has a conditional probability $\text{val}_{\mathcal{M}}(w)$ that $\text{last}(\rho) = q_c^1$ and $1 - \text{val}_{\mathcal{M}}(w)$ that $\text{last}(\rho) = q_c^2$. Thus:

$$\sum_{w \in \Sigma^{n-1}} \mathbb{P}_\sigma(w\sharp^2 \wedge C) = \sum_{w \in \Sigma^{n-1}} \text{val}_{\mathcal{M}}(w) \mathbb{P}_\sigma(w\sharp \wedge C) \leq \frac{1}{2} \sum_{w \in \Sigma^{n-1}} \mathbb{P}_\sigma(w\sharp \wedge C)$$

with a strict inequality if there exists a word w with $\mathbb{P}_\sigma(w\sharp) > 0$.

By assumption, \mathcal{C}_σ is diagnosable. Thus, according to our characterisation of a strategy ensuring the diagnosis, there exists a word w such that $\mathbb{P}_\sigma(w\sharp) > 0$. As a consequence, $\sum_{n=1}^{\infty} \mathbb{P}(C_n) < \sum_{n=1}^{\infty} (\frac{1}{2})^n = 1$, thus \mathcal{A} is not $(1, 1)$ fault free. \square

Observation. A straightforward adaptation of the proof establishes that for $0 < \gamma < 1$, \mathcal{A} is $(\gamma, \frac{\gamma}{2-\gamma})$ fault free iff there exists a word w such that $\text{val}_{\mathcal{M}}(w) \geq \frac{1}{2}$.

Proposition 2 α -resilient active diagnosability is undecidable.

Proof We proceed here by reduction the problem of the existence of a word w such that $\text{val}_{\mathcal{M}}(w) > \frac{1}{2}$. We consider the probabilistic automaton $\mathcal{M} = \langle S, s_0, F, \Sigma, \mathbf{P} \rangle$ for which we assume w.l.o.g. that: (1) $\Sigma \cap \{u, \mathbf{f}, \sharp, \natural\} = \emptyset$ and (2) the probabilities are fractions $\frac{n}{d}$ with d fixed. One builds the CLTS $\mathcal{C} = \langle Q, s_0, \Sigma', T \rangle$ represented in Figure 6 (with some shortcuts to ease readability) and defined by:

- $Q = S \cup \{q_1, f_1\}$;
- $\Sigma' = \Sigma \cup \{\mathbf{f}, \natural, \natural\}$, $\Sigma_u = \{\mathbf{f}\}$ et $\Sigma_c = \Sigma \cup \{\sharp\}$;
- the transition function T is defined by:
 1. $T(q_1, \mathbf{f}, f_1) = T(f_1, \natural, f_1) = 1$;
 2. for all $s, s' \in S, a \in \Sigma, T(s, a, s') = d \cdot \mathbf{P}(s, a, s')$ and $T(s, a, q_1) = d$;
 3. for all $s \in F, T(s, \sharp, s_0) = 1$ and for all $s \in S \setminus F, T(s, \sharp, q_1) = 1$;
 4. for every other triplet, T is equal to 0.

Here again, the probabilities in \mathcal{M} are multiplied by the constant d , which we abbreviate in the figure by $d \cdot \mathcal{M}$.




Fig. 6 From a probabilistic automaton to (another) CLTS.

As a faulty run is followed by a \natural , whatever the strategy σ , \mathcal{C}_σ is diagnosable.

• Assume there exists $w = w_1 \dots w_k \in \Sigma^*$ such that $\text{val}_{\mathcal{M}}(w) > \frac{1}{2}$. We denote $v = \text{val}_{\mathcal{M}}(w)$. We define the deterministic strategy σ by:

- $\sigma((w\sharp)^*w) = \{\mathbf{f}, \natural, \sharp\}$;
- for all $0 \leq i < k$, $\sigma((w\sharp)^*w_1 \dots w_i) = \{\mathbf{f}, \natural, w_{i+1}\}$;
- $\sigma(w') = \Sigma'$ for any other word w' .

Under strategy σ , the observed sequence of a correct run ρ is some $(w\sharp)^m w_1 \dots w_i$ with $0 \leq i \leq k$.

◦ If $\pi(\rho) = (w\sharp)^m w_1 \dots w_i$ with $0 < i$ then with conditional equiprobability, $\text{last}(\rho) \in S$ or $\text{last}(\rho) = q_1$. Thus with probability $\frac{1}{2}$, the run will be correct after the next observation.

◦ If $\pi(\rho) = (w\sharp)^m$ then with conditional probability v , $\text{last}(\rho) = s_0$ and with probability $1 - v$, $\text{last}(\rho) = q_1$. Thus with probability v , the run will be correct after the next observation.

Consider an arbitrary n and write its Euclidian division by $k + 1$ as $n = m(k + 1) + i$ with $i \leq k$. One has $2^{-(n-1)} \mathbb{P}_\sigma(\mathcal{C}_n) = \left(\frac{v}{2}\right)^m$.

Hence $\frac{2^{-(n-1)}}{\mathbb{P}_\sigma(\mathcal{C}_n)} = \left(\frac{2}{v}\right)^{\lfloor \frac{n-1}{k+1} \rfloor}$ implying $\lim_{n \rightarrow \infty} \frac{2^{-n}}{\mathbb{P}_\sigma(\mathcal{C}_n)} = 0$. So \mathcal{C}_σ is $\frac{1}{2}$ -resilient.

• Assume now that for all word $w \in \Sigma^*$, $\text{val}_{\mathcal{M}}(w) \leq \frac{1}{2}$. Let σ be an arbitrary strategy. The observed sequence of a correct run ρ is some $u_1\sharp \dots \sharp u_m$ such that for all i , $u_i \in \Sigma^*$.

◦ Si $u_m \neq \mathbf{1}$ with $0 < i$ then with conditional equiprobability, $\text{last}(\rho) \in S$ or $\text{last}(\rho) = q_1$. Thus with probability $\frac{1}{2}$, the run will be correct after the next observation.

◦ If $u_m = \mathbf{1}$ then with conditional probability $\text{val}_{\mathcal{M}}(u_{m-1})$, $\text{last}(\rho) = s_0$ and with probability $1 - \text{val}_{\mathcal{M}}(u_{m-1})$, $\text{last}(\rho) = q_1$. Thus with probability $\text{val}_{\mathcal{M}}(u_{m-1})$, the run will be correct after the next observation.

Summarising one has: $\mathbb{P}_\sigma(\mathcal{C}_n) \leq 2^{-(n-1)}$ implying $\limsup_{n \rightarrow \infty} \frac{2^{-n}}{\mathbb{P}_\sigma(\mathcal{C}_n)} \geq \frac{1}{2}$.

So \mathcal{C}_σ is not $\frac{1}{2}$ -resilient. \square

3.3 Decidability of the Qualitative Problems

In contrast to the quantitative notions, and to the notable exception of the safe active diagnosis, all the qualitative problems of diagnosability under degradation constraints we introduced are decidable and EXPTIME-complete. Moreover, to remedy the undecidability of the safe active diagnosis problem [2], in a second step, we also establish its EXPTIME-completeness when restricted to finite-memory strategies.

We start with the three qualitative problems that turn out to be decidable, even with no restriction on the memory of control strategies. The techniques used in the proofs are similar in spirit to the ones used for solving decision problems in Partially Observable Markov Decision Processes [5]. The proof idea is common to all cases: we establish a necessary and sufficient condition for the existence of a control strategy that ensures the given notion of diagnosability under a degradation constraint. To do so, we revisit the construction given in [2] (for active diagnosis), which builds an enriched model including finite information of the history. On this enriched model, the necessary and sufficient condition consists of graph-based properties.

Let us start by recalling the construction from [2]. To decide the diagnosability of a CLTS, its states are enriched with two subsets of states: U and V that correspond,

respectively, to the subset of correct, or faulty, states, that are reachable by a signalling run corresponding to the current observed sequence. A pair (U, V) is called a *belief*. Formally, from a CLTS $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$, we define its belief version on the same event alphabet $\mathcal{C}^B = \langle Q^B, q_0^B, \Sigma, T^B \rangle$ by:

- $Q^B = Q \times 2^Q \times 2^Q$ and $q_0^B = (q_0, \{q_0\}, \emptyset)$;
- for every $(q, U, V) \in Q \times 2^Q \times 2^Q$, for every $a \in \Sigma$, and every $q' \in Q$
 - if $a \notin \Sigma_o$, $T^B((q, U, V), a, (q', U, V)) = T(q, a, q')$;
 - if $a \in \Sigma_o$, letting
 1. $U' = \{q'_c \in Q_c \mid \exists q_c \in U, \exists \rho \in \text{SR}_1, q_c \xrightarrow{\rho} q'_c \wedge \pi(\rho) = a\}$
 2. $V' = \{q'_f \in Q_f \mid \exists q_x \in U \cup V, \exists \rho \in \text{SR}_1, q_x \xrightarrow{\rho} q'_f \wedge \pi(\rho) = a\}$.
 then $T^B((q, U, V), a, (q', U', V')) = T(q, a, q')$.
- for every other triplet $((q, U, V), a, (q', U', V'))$, T is equal to 0.

The size of the belief CLTS \mathcal{C}^B is exponential in the size of \mathcal{C} . For the properties we are interested in, they have the same behaviour. We introduce Δ , a discrete version of T^B , extended to observed sequences. For $w \in \Sigma_o^*$, $(q', U', V') \in \Delta((q, U, V), w)$ as soon as there exists a run ρ such that $\pi(\rho) = w$ and $(q, U, V) \xrightarrow{\rho} (q', U', V')$.

We will now construct Win in the set of all beliefs (U, V) such that, starting from any (q, U, V) with $q \in U \cup V$, \mathcal{C}^B is actively diagnosable. This set is computed as a greatest fixpoint. We let $\text{Win}_0 = 2^{Q_c} \times 2^{Q_f}$ and for $n \in \mathbb{N}$, Win_{n+1} is the set of the beliefs (U, V) of Win_n such that for all state $q \in U \cup V$, there exists a sequence of sets of allowed events $(\Sigma_i^\bullet)_{1 \leq i \leq k}$ and an observed sequence $w = o_1 \dots o_k$ with $o_i \in \Sigma_i^\bullet$ verifying:

- there exists a run ρ starting in (q, U, V) with $\pi(\rho) = w$ and reaching (q^*, U^*, V^*) with $q^* \in Q_c$ (i.e. the current state is correct) or $U^* = \emptyset$ (the fault is claimed);
- Consider a state q_i reached from $q' \in U \cup V$ by a run with observed sequence $o_1 \dots o_i$ with $0 \leq i < k$, i.e. $(q_i, U_i, V_i) \in \Delta((q', U, V), o_1 \dots o_i)$ for a belief (U_i, V_i) . then:
 1. the control induced by Σ_{i+1}^\bullet does not create any deadlock: $G^{\Sigma_{i+1}^\bullet}(q_i) \neq \emptyset$;
 2. Every new belief obtained by an observable step $o \in \Sigma_{i+1}^\bullet$ starting in q_i belongs to Win_n : $\forall o \in \Sigma_{i+1}^\bullet, \forall (q_o, U_o, V_o) \in \Delta((q_i, U_i, V_i), o), (U_o, V_o) \in \text{Win}_n$.

The computation of Win is in polynomial time in the size of \mathcal{C}^B , given that at every non-terminal iteration at least one belief is removed. The correctness of Win is established in [2], and σ^* a (deterministic finite-memory) strategy ensuring diagnosability consists in, given a belief $(U, V) \in \text{Win}$ choosing the greatest set Σ^\bullet such that every possible belief reached on the next step still belongs to Win.

To decide weakly (resp. strongly) resilient active diagnosability, and lasting fault free active diagnosability, we build on the belief CLTS construction. The simplest case is the weak notion:

Theorem 2 *Weakly resilient active diagnosability is EXPTIME-complete.*

Proof We first establish the membership in EXPTIME. Given a CLTS \mathcal{C} , its belief CLTS \mathcal{C}^B , and the deterministic finite-memory σ^* , we derive a pLTS \mathcal{A} . It is obtained

from \mathcal{C}^B by restricting to the states of with belief in Win and controlled by the strategy σ^* . We claim that \mathcal{C} is actively diagnosable with guarantee of weak resiliency iff there exists in \mathcal{A} a reachable cycle such that the first component of every state along the cycle is a correct state of \mathcal{C} .

- Suppose first that such a cycle exists in \mathcal{A} . We let $\alpha > 0$ be the probability of this cycle, n_1 its length, n_0 the observed length of the shortest run reaching a state of the cycle and μ the probability of this run. For all $n \geq n_0$, $\mathbb{P}_{\mathcal{A}}(\mathcal{C}_n) \geq \mu\alpha^{\lceil \frac{n-n_0}{n_1} \rceil}$. As a consequence, \mathcal{A} is α' -resilient for all $\alpha' < \alpha$. \mathcal{A} is thus weakly resilient. Therefore, \mathcal{C}_{σ^*} , which has the same probabilistic behaviour as \mathcal{A} is weakly resilient too.

- Conversely, suppose that there is no such cycle in \mathcal{A} . Let σ' be a (live) strategy such that $\mathcal{C}_{\sigma'}$ is diagnosable. This strategy can be mimicked in \mathcal{C}^B , ignoring the belief information. The reachable states of $\mathcal{C}_{\sigma'}^B$ are associated with beliefs of Win (due to the characterisation recalled above). As σ^* is the most permissive strategy ensuring to stay in Win, there does not exist any such cycle in $\mathcal{C}_{\sigma'}^B$ either. Consequently, there exists $n_f \in \mathbb{N}$ such that every run ρ in $\mathcal{C}_{\sigma'}^B$ with $|\rho| \geq n_f$ ends in a state which first component is faulty. Thus $\mathbb{P}_{\mathcal{C}_{\sigma'}}(\mathcal{C}_{n_f}) = \mathbb{P}_{\mathcal{C}_{\sigma'}^B}(\mathcal{C}_{n_f}) = 0$, which means that $\mathcal{C}_{\sigma'}$ is not weakly resilient.

The complexity lower-bound is obtained by reduction from the active diagnosability, which is known to be EXPTIME-hard [2]. For $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ a CLTS, we define the CLTS $\mathcal{C}' = \langle Q \cup \{q'_0, q_s\}, q'_0, \Sigma \cup \{\#\}, T' \rangle$ with $\#$ a fresh observable event, and such that $T'(q'_0, \#, q_0) = T'(q'_0, \#, q_s) = T'(q_s, \#, q_s) = 1$, for all $q, q' \in Q, a \in \Sigma, T'(q, a, q') = T(q, a, q')$ and for every other triplet $T'(q, a, q') = 0$. Clearly enough, \mathcal{C}' is diagnosable iff \mathcal{C} is diagnosable. Moreover, \mathcal{C}' is safe by construction, and thanks to Theorem 1(a), it is strongly resilient, and thus weakly resilient. \square

The proof of the next theorem also relies on the set of beliefs Win. We build a subset of Win, called WinK. A belief (U, V) of Win belongs to WinK if there exists a strategy σ such that from every distribution with support $U \cup V$, σ guarantees to stay in Win, and to give a positive probability to the set of infinite correct runs. The CLTS is actively diagnosable with guarantee of strong resiliency iff from the initial belief one can reach a belief of WinK while staying in Win. The winning strategy consists in combining cleverly the strategy used to make the system diagnosable and the one allowing to stay in WinK.

Theorem 3 *Strongly resilient active diagnosability is EXPTIME-complete.*

Proof Let \mathcal{C} be a CLTS. As in the construction preliminary to Theorem 2, we build \mathcal{C}^B , Win and σ^* . We then define $\text{WinK}_U \subseteq 2^Q \times \text{Win}$ by a greatest fix point computation. For $(U', (U, V)) \in \text{WinK}_U$, (U, V) is a belief for which there exists a strategy allowing to a set of runs starting in U' to stay in the states of \mathcal{C}^B associated with a belief of Win while staying correct. WinK_U is obtained as the limit of a decreasing sequence $(\text{WinK}_n)_{n \in \mathbb{N}}$ defined inductively by: $\text{WinK}_0 = \{(U', (U, V)) \mid (U, V) \in \text{Win} \wedge \emptyset \neq U' \subseteq U\}$ and for $n \in \mathbb{N}$, WinK_{n+1} is the set of elements $(U', (U, V))$ of WinK_n such that there exist a set of allowed events Σ^\bullet verifying:

- Σ^\bullet does not create a deadlock: $\forall q \in U \cup V, G^{\Sigma^\bullet}(q) \neq \emptyset$;

- under the control Σ^\bullet no run starting in a state of U' will make a fault before the next observation: $\forall q_c \in U', \forall \rho \in \text{SR}_1, q_c \xrightarrow{\rho} q \wedge \pi(\rho) \in \Sigma^\bullet \Rightarrow q \in Q_c$;
- every triplet reached by an observable step $o \in \Sigma^\bullet$ belongs to WinK_n : $(\tilde{U}', (\tilde{U}, \tilde{V})) \in \text{WinK}_n$ with:
 1. $\tilde{U}' = \{q'_c \in Q_c \mid \exists q_c \in U'_1, \exists \rho \in \text{SR}_1, q_c \xrightarrow{\rho} q'_c \wedge \pi(\rho) = a\}$;
 2. (\tilde{U}, \tilde{V}) is obtained by the update of the belief (U, V) following the observation o .

From WinK_U , we define the set $\text{WinK} \subseteq \text{Win}$ by keeping only the second component of WinK_U : $\text{WinK} = \{(U, V) \in \text{Win} \mid \exists U', (U', (U, V)) \in \text{WinK}_U\}$. Let us state some of the properties of this construction.

- By induction, if $(U', (U, V)) \notin \text{WinK}_n$ then for every (live) strategy, there exists a faulty run starting in U' of observable length n ;
- If $\emptyset \neq U'' \subseteq U'$ then $(U', (U, V)) \in \text{WinK}_U$ implies $(U'', (U, V)) \in \text{WinK}_U$. Thus, if $(U, V) \notin \text{WinK}$, for all $q \in U$, $(\{q\}, (U, V)) \notin \text{WinK}_U$.

We also define PreWin the set of states of \mathcal{C}^B of the form $Q \times \text{Win}$ from which a state (q, U, V) with $(U, V) \in \text{WinK}$ is reachable. Let us show that \mathcal{C} is diagnosable and strongly resilient iff the initial state of \mathcal{C}^B belongs to PreWin .

• Suppose that the initial state belongs to PreWin . Let $(U', (U, V))$ be an element of WinK_U . We define $\sigma_{(U', (U, V))}$ the finite-memory strategy with memory states of the form $(\tilde{U}', (\tilde{U}, \tilde{V}))$ and which, starting from $(U', (U, V))$, ensures to stay in WinK_U . This strategy immediately derives from the fixpoint definition of WinK_U .

For $(U, V) \in \text{WinK}$, we also define $\sigma_{(U, V)} = \sigma_{(U', (U, V))}$ for an arbitrary U' such that $(U', (U, V)) \in \text{WinK}_U$.

Finally, we let σ_0 be the following strategy working in three successive phases which may not all be triggered.

1. First σ_0 mimicks σ^* until a belief $(U, V) \in \text{WinK}$ is reached;
2. Then, at every observed sequence w , σ_0 chooses to apply $\sigma_{(U, V)}$ with probability $p_w = \frac{|w|}{|w|+1}$, and to switch to the third phase with probability $1 - p_w$;
3. Finally, σ_0 behaves forever as σ^* .

We observe that \mathcal{C}_{σ_0} is diagnosable. Indeed, on the one hand, the events allowed by σ_0 are included in those allowed by the maximally permissive strategy σ^* , and on the other hand almost-surely, σ^* is applied from some moment on. Therefore every fault will almost surely be detected.

Moreover, let us prove that it is strongly resilient. Indeed, by definition of PreWin , there exists a run ρ starting in the initial state and reaching a state (q, U, V) such that (U, V) belongs to WinK . Let $U' \subseteq U$ the one chosen arbitrarily when defining $\sigma_{(U, V)}$. Without loss of generality, we suppose that ρ reaches a state of U' . As a fault can only be created after ρ if σ_0 switches to its third phase, for $n \geq |\rho|_o$ we have

$$\mathbb{P}_{\sigma_0}(\tilde{\rho} \in \text{C}_n \mid \rho \preceq \tilde{\rho}) \geq \mathbb{P}_{\sigma_0}(\rho) \prod_{i=|\rho|}^n \frac{i}{i+1} = \mathbb{P}_{\sigma_0}(\rho) \frac{|\rho|}{n+1} .$$

Thus, for every $0 < \alpha < 1$, similarly to $n\alpha^n$, $\frac{\alpha^n}{\mathbb{P}_{\sigma_0}(\mathcal{C}_n)}$ converges to 0.

• Conversely, suppose that the initial state does not belong to PreWin. Let σ be a strategy ensuring diagnosability. For every state (q, U, V) with $q \in U$ reachable by a run ρ_0 with σ , $(U, V) \notin \text{WinK}$ and due to our one of our observations $(\{q\}, (U, V)) \notin \text{WinK}_U$. Let K be the number of iterations in the fixpoint computation of WinK. Then, for every sequence of K random choices under σ , there exists a faulty run $\rho \in F$, compatible with these choices, starting in (q, U, V) and of observable length smaller than K . Adding up the probabilities of runs corresponding to every sequence of choices of σ we obtain

$$\mathbb{P}_\sigma(\rho \in F_{|\rho_0|_{\sigma}+K} \mid \rho_0 \preceq \rho) \geq \lambda^{K \cdot |Q|} \mathbb{P}_\sigma(\rho_0)$$

where $\lambda = \min_{q \in Q} \frac{1}{G^{\mathcal{S}}(q)}$. Thus, for every $n \in \mathbb{N}$, $\mathbb{P}_\sigma(\mathcal{C}_{n+K}) \leq \mathbb{P}_\sigma(\mathcal{C}_n)(1 - \lambda^{K \cdot |Q|})$. Letting $\alpha = (1 - \lambda^{K \cdot |Q|})^{\frac{1}{K}}$, we obtain $\lim_{n \rightarrow \infty} \frac{\alpha^n}{\mathbb{P}_\sigma(\mathcal{C}_n)} > 0$, so that \mathcal{C}_{σ_0} is not strongly resilient.

To conclude the proof, we observe that the EXPTIME-hardness derives from the same reduction as in the proof of Theorem 2. \square

It turns out that this same combination of strategies can be used to ensure lasting fault freeness and diagnosability. In fact, the following theorem establishes that the characterisation of the strongly resilient active diagnosability also applies to the lasting fault free active diagnosability.

Theorem 4 *Lasting fault free active diagnosability is equivalent to strongly resilient active diagnosability.*

Proof We will show here that the characterisation given in the proof of Theorem 3 for a CLTS to be actively diagnosable with guarantee of strong resiliency also characterises the fact that the CLTS is actively diagnosable with guarantee of lasting fault freeness. This will show the equivalence of the two notions in the active case.

We reuse the definitions from the proof of Theorem 3. Let us show that \mathcal{C} is actively diagnosable with guarantee of lasting fault freeness iff the initial state of \mathcal{C}^B belongs to PreWin.

• Suppose that the initial state belongs to PreWin. Then, as discussed in the proof of Theorem 3, \mathcal{C}_{σ_0} is diagnosable and there exists a finite run ρ such that $\mathbb{P}(\tilde{\rho} \in \mathcal{C}_n \mid \rho \preceq \tilde{\rho}) \geq \mathbb{P}(\rho) \frac{|\rho|}{n+1}$. Thus:

$$\sum_{n=1}^{\infty} \mathbb{P}(\mathcal{C}_n) \geq \sum_{n=|\rho|}^{\infty} \mathbb{P}(\tilde{\rho} \in \mathcal{C}_n \mid \rho \preceq \tilde{\rho}) \geq \mathbb{P}(\rho) |\rho| \sum_{n=|\rho|}^{\infty} \frac{1}{n+1} = \infty.$$

• Conversely, if the initial state does not belong to PreWin. Let σ be a strategy ensuring diagnosability. For every $n \in \mathbb{N}$, $\mathbb{P}(\mathcal{C}_{n+K}) \leq \mathbb{P}(\mathcal{C}_n)(1 - \lambda^{K \cdot |Q|})$. Thus:

$$\sum_{n=1}^{\infty} \mathbb{P}(\mathcal{C}_n) \leq K \sum_{n=1}^{\infty} (1 - \lambda^{K \cdot |Q|})^n \leq K \cdot |Q_B| \cdot \frac{1}{\lambda^{K \cdot |Q|}} < \infty.$$

\square

Given the equivalence of strong resiliency and lasting fault freeness, from Theorem 3 we derive:

Corollary 1 *Lasting fault free active diagnosability is EXPTIME-complete.*

We now turn our attention to safe active diagnosability. The problem is known to be undecidable in general, and in NEXPTIME when restricting to finite-memory strategies [2]. Note that decidability is not immediate even if the strategies are assumed to be finite-memory, since no *a priori* bound on the memory is known. We refine that complexity result by proving that safe active diagnosis can be solved in EXPTIME when restricting to finite-memory strategies.

To do so, we prove a more general result in the context of a well-known model, quite popular in artificial intelligence and more recently in formal methods, that combines partial observation, probabilities and control, namely *Partially Observable Markov Decision Processes* (POMDP). We establish that the existence of finite-memory schedulers that ensure a Büchi objective with probability 1 and a safety objective with positive probability in a POMDP is decidable in EXPTIME. We then reduce the safe active diagnosis of a CLTS \mathcal{C} restricted to finite-memory strategies to the existence of a finite-memory scheduler in a POMDP $M_{\mathcal{C}}$ ensuring at the same time a Büchi objective with probability 1 and a safety objective with positive probability.

Definition 10 (POMDP) A *partially observable Markov decision process* (POMDP) is a tuple $M = \langle Q, q_0, \text{Obs}, \text{Act}, T \rangle$ 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 $\text{Ena}(q) = \{a \in \text{Act} \mid T(q, a) \text{ is defined}\}$ the set of enabled actions in state q , we assume that:
 - for all $q \in Q$, $\text{Ena}(q) \neq \emptyset$, and
 - whenever $\text{Obs}(q) = \text{Obs}(q')$, then $\text{Ena}(q) = \text{Ena}(q')$ and slightly abusing our notation, we will denote by $\text{Ena}(O)$ the set of events enabled in every state with observation O .

A *decision rule* is a distribution from $\text{Dist}(\text{Act})$ that resolves non-determinism by randomization. A *scheduler* for a POMDP maps histories of observations to decision rules. Formally, a scheduler is a function $\tau : \mathcal{O}^+ \rightarrow \text{Dist}(\text{Act})$ such that for every $O_1 \cdots O_i$, $\text{Supp}(\tau(O_1 \cdots O_i)) \subseteq \text{Ena}(O_i)$. Given a scheduler τ , a POMDP M yields a stochastic process. This stochastic process can be represented by an infinite state pLTS, denoted $M(\tau)$ in which states are histories of observations. One denotes by $\mathbb{P}_{\tau}^{q_0}(\text{Ev})$ the probability that event Ev is realized in this process.

In the context of POMDP, a *belief* is a non-empty set of states that represents the current state estimate, *i.e.* the set of states the system may be in, given the actions and observations so far. The initial belief is $\{q_0\}$, and given a current belief B , a decision rule δ and an observation O , the belief obtained after δ has been applied and O has been observed is defined by:

$$\Delta(B, (\delta, O)) = \bigcup_{q \in B, a \in \text{Supp}(\delta)} \text{Supp}(T(q, a)) \cap \text{Obs}^{-1}(O) .$$

Of course, beliefs can similarly be defined for CLTS. Again, the initial belief is $\{q_0\}$, and given a current belief B and an observed event b , the belief obtained after b has been observed is defined by:

$$\Delta(B, b) = \{q \in Q \mid \exists q' \in B, \rho \in SR_1, q' \xrightarrow{\rho} q \wedge \pi(\rho) = b\} .$$

Intuitively, $\Delta(B, b)$ is the set of states a partially observable systems may be in, given that the previous belief was B and observation O occurred. It does not depend on the strategy as every controllable event is observable. The set of beliefs is denoted \mathcal{Bl}_C and we drop the subscript when there is no risk of confusion. Beliefs are of importance since they formalize the discrete information an observer has on the current state of the system.

Aiming at providing a POMDP M_C for the safe active diagnosis problems of a CLTS \mathcal{C} , we face several difficulties. First, in a CLTS the observations are related to events while in a POMDP they are related to states. Fortunately, the relevant information pertaining to the observations, namely the information about ambiguity of observed sequences, is available in the belief. Thus (with one exception) the states are pairs of a state q of the CLTS and a belief B . A second adaptation concerns the control mechanism. In \mathcal{C} , the control is performed by choosing (possibly randomly) a subset of allowed controllable events. Thus actions of M_C are subsets of events that include the uncontrollable events. Given some control decision Σ^\bullet , to define the transition probability of M_C from (q, B) to (q', B') , one must consider all paths in \mathcal{C} labelled by events of Σ^\bullet from q to q' such that the last event is the only observable one. 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^\bullet . Finally, there cannot be infinite paths of unobservable events due to the convergence of \mathcal{C} . However some paths can reach, via unobservable events, a state from which no event of Σ^\bullet is enabled. In other words, the control Σ^\bullet applied in (q, B) may have a positive probability to reach a deadlock (*i.e.* the chosen decision rule leads to a strategy for the CLTS which is not live). 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 11 The POMDP $M_C = \langle Q^{M_C}, q_0^{M_C}, \text{Obs}, \text{Act}, T^{M_C} \rangle$ derived from a CLTS $\mathcal{C} = \langle Q, q_0, \Sigma, T \rangle$ is defined by:

- $Q^{M_C} = Q \times \mathcal{Bl}_C \uplus \{\text{lost}\}$ with $q_0^{M_C} = (q_0, \{q_0\})$;
- the set of observations is $\mathcal{O} = \mathcal{Bl}_C \cup \{\text{lost}\}$, with $\text{Obs}(\text{lost}) = \text{lost}$ and for $(q, B) \in Q^{M_C}$, $\text{Obs}((q, B)) = B$;
- $\text{Act} = \{\Sigma^\bullet \subseteq \Sigma \mid \Sigma^\bullet \supseteq \Sigma \setminus \Sigma_c\}$;
- for every $(q_1, B) \in Q^{M_C}$ and $\Sigma^\bullet \in \text{Act}$, $T^{M_C}((q_1, B), \Sigma^\bullet) = \mu \in \text{Dist}(Q^M)$ where:

$$- \mu((q', B')) =$$

$$\sum_{\substack{\Delta(B, b) = B' \\ b \in \Sigma^\bullet \cap \Sigma_o}} \sum_{\substack{q_1 \xrightarrow{a_1} q_2 \cdots \xrightarrow{a_n} q_{n+1} \xrightarrow{b} q' \\ a_1 \cdots a_n \in \Sigma^\bullet \cap \Sigma_u}} \left(\prod_{i=1}^n T^{\Sigma^\bullet}(q_i, a_i, q_{i+1}) \right) \cdot T^{\Sigma^\bullet}(q_{n+1}, b, q');$$

$$\begin{aligned}
- \mu(\text{lost}) &= \sum_{\substack{q_1 \xrightarrow{a_1} q_2 \cdots \xrightarrow{a_n} q_{n+1} \\ a_1 \cdots a_n \in \Sigma^\bullet \cap \Sigma_u \\ G^{\Sigma^\bullet}(q_{n+1})=0}} \prod_{i=1}^n T^{\Sigma^\bullet}(q_i, a_i, q_{i+1}); \\
- \text{for every } \Sigma^\bullet \in \text{Act}, T^{\text{M}_C}(\text{lost}, \Sigma^\bullet) &= \mathbf{1}_{\text{lost}}.
\end{aligned}$$

Given \mathcal{C} , the construction of M_C , which is of size in $2^{O(|Q|+|\Sigma|)}$, can be done in exponential time. Also, the probability distributions over next states (μ in Definition 11) are presented as sums over paths of \mathcal{C} , but they can be computed in polynomial time by matrix operations (as for DTMC).

A CLTS \mathcal{C} and its associated POMDP M_C are closely related. In particular, strategies in \mathcal{C} and schedulers in M_C are in a one-to-one correspondence. On the one hand, let us explain how to naturally derive a strategy σ for \mathcal{C} from a scheduler τ in M_C . For an observed sequence $a_1 \cdots a_n \in \Sigma_o^*$, there is a unique sequence of beliefs B_0, \dots, B_n such that $B_0 = \{q_0\}$ and for all $i < n$, $B_{i+1} = \Delta(B_i, a_{i+1})$. We then set $\sigma(a_1 \cdots a_n) = \tau(B_0 \cdots B_n)$. Notice that the strategy σ obtained that way is not necessarily live: for example, if after B_1, \dots, B_n the choice of σ leads with positive probability to lost , then σ is not live. However, as soon as τ ensures to avoid state lost , then the corresponding strategy σ is live.

On the other hand, to a live strategy σ for \mathcal{C} , we can associate a scheduler τ in M_C that always avoids lost . For a sequence of observations that does not contain lost , thus of the form $B_0 \cdots B_n$, with $B_i \subseteq Q$ for all i , we pick $a_1 \cdots a_n \in \Sigma_o^*$ an observed sequence such that for all $i < n$, $B_{i+1} = \Delta(B_i, a_{i+1})$. We then set $\tau(B_0 \cdots B_n) = \sigma(a_1 \cdots a_n)$. Note that the observed sequence is not uniquely defined from $B_0 \cdots B_n$. However, if $a_1 \cdots a_n$ and $a'_1 \cdots a'_n$ both lead to the belief B_n , the set of possible states of the CLTS after both observed sequences is the same. Therefore, the same subsets Σ^\bullet after both sequences leave the system live, and the same actions Σ^\bullet yield a probability distribution μ such that $\mu(\text{lost}) = 0$.

Moreover, if (σ, τ) is a pair of live strategy and corresponding scheduler (that always avoids lost), the probability measures $\mathbb{P}_{\mathcal{C}_\sigma}$ and $\mathbb{P}_\tau^{\text{M}_C}$ are essentially equivalent. More precisely, the product in M_C with the belief does not change the probability measure defined by \mathcal{C}_σ .

We now show how to decide for POMDP the existence of a finite-memory scheduler that ensures a Büchi objective with probability one and a safety objective with positive probability. 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.

Theorem 5 *The problem whether, given a POMDP M with subsets of states F and I , there exists a finite-memory scheduler τ such that $\mathbb{P}_\tau^{\text{M}}(\square\diamond F) = 1$ and $\mathbb{P}_\tau^{\text{M}}(\square I) > 0$ is EXPTIME-complete.*

Theorem 5 derives from Propositions 3 and 4 below, that state, respectively, the upper bound in the general case, and the lower bound in a particular case, namely for the safe active diagnosability under finite-memory strategies.

Proposition 3 *Given a POMDP M with subsets of states F and I , one can decide in EXPTIME whether there exists a finite-memory scheduler τ such that $\mathbb{P}_\tau^{\text{M}}(\square\diamond F) = 1$ and $\mathbb{P}_\tau^{\text{M}}(\square I) > 0$.*

Proof In this proof, the POMDP $M = \langle Q, q_0, \text{Obs}, \text{Act}, T \rangle$ is fixed, and we use notation $\mathbb{P}_\tau^{\delta_0}(\text{Ev})$ to denote the probability of the event Ev under scheduler τ assuming that instead of q_0 , the initial state in M is given by the distribution $\delta_0 \in \text{Dist}(Q)$.

Let us first explain how to compute the following set of pairs of beliefs:

$$\begin{aligned} \text{Win}_{=1} = \{ & (B', B) \mid B' \subseteq I, B' \subseteq B \text{ s.t. } \exists \tau \text{ s.t.} \\ & \forall \delta_0 \text{ with } \text{Supp}(\delta_0) = B, \mathbb{P}_\tau^{\delta_0}(\Box \Diamond F) = 1, \text{ and} \\ & \forall \delta'_0 \text{ with } \text{Supp}(\delta'_0) = B', \mathbb{P}_\tau^{\delta'_0}(\Box I) = 1 \} . \end{aligned}$$

Intuitively, $\text{Win}_{=1}$ denotes pairs of beliefs such that there exists a scheduler that ensures a Büchi objective almost-surely from the larger belief, and a safety objective almost-surely from the smaller one. Note that, in the definition of $\text{Win}_{=1}$, we do not require the scheduler τ to be finite-memory. Given that we consider pairs of beliefs, we introduce the following notation: $\Delta((B', B), O_1) = (\Delta(B', O_1), \Delta(B, O_1))$, and similarly for sequences of actions and observations. Also, for $X \subseteq Q$ a subset of states, we denote by $\mathcal{Bl}_{\subseteq X} = \{B \in \mathcal{Bl} \mid B \subseteq X\}$ the set of beliefs contained in X .

Lemma 2 *Let Win_∞ be the greatest fixed point starting from $\{(q, B', B) \in Q \times \mathcal{Bl} \times \mathcal{Bl} \mid q \in B, B' \subseteq B, B' \subseteq I\}$ of the following operator:*

$$\begin{aligned} W \mapsto \{ & (q, B'_1, B_1) \mid \exists n \geq 1, \exists q_0 \dots q_n \in Q, \exists \alpha_1, \dots, \alpha_n \exists O_1 \dots O_n, \\ & (B'_2, B_2) = \Delta((B'_1, B_1), (\alpha_1, O_1) \dots (\alpha_n, O_n)), \forall q' \in B_2, (q', B'_2, B_2) \in W, \\ & q_0 = q, q_n \in F, \forall i < n, T(q_i, \alpha_{i+1})(q_{i+1}) > 0, \forall 1 \leq j \leq n, \text{Obs}(q_j) = O_j, \\ & \forall i \leq n, \forall O'_i, \text{ for } (B'_3, B_3) = \Delta((B', B), (\alpha_1, O_1) \dots (\alpha_{i-1}, O_{i-1})(\alpha_i, O'_i)) \\ & \text{ we have } \forall q' \in B_3, (q, B'_3, B_3) \subseteq W \cap Q \times \mathcal{Bl}_{\subseteq I} \times \mathcal{Bl} \} . \end{aligned}$$

We have $\text{Win}_{=1} = \{(B', B) \mid \forall q \in B, (q, B', B) \in \text{Win}_\infty\}$.

Proof (of Lemma 2) To establish that $\text{Win}_{=1}$ corresponds to the projection on the pair of beliefs of Win_∞ , we first assume that for all $q \in B$, (q, B', B) belongs to Win_∞ , and exhibit a scheduler τ that witnesses $(B', B) \in \text{Win}_{=1}$. Let us define τ as follows. The scheduler τ has finite memory $\mathcal{Bl} \times \mathcal{Bl}$. From memory state (B', B) , τ dictates to play uniformly all actions α such that for every observation O and every $q \in \Delta(B, \alpha, O)$, we have $(q, \Delta((B', B), \alpha, O)) \in \text{Win}_\infty$. Note that this set of “safe” actions is necessarily non empty because $(q, B', B) \in \text{Win}_\infty$. If α is played, and O is observed, the memory state of τ is updated to $\Delta((B', B), \alpha, O)$, which is still in Win_∞ , by assumption on α . The scheduler τ then continues similarly with memory state $\Delta((B', B), \alpha, O)$.

So defined, let us show that τ witnesses $(B', B) \in \text{Win}_{=1}$. First, let δ_0 be a distribution with support B . The scheduler τ ensures to stay (surely) in Win_∞ . Moreover, for every $q \in B$, with a positive probability, say $p_{(q, B', B)} > 0$, the sequence $(\alpha_1, O_1) \dots (\alpha_n, O_n)$ of actions and observations leading to F that derives from the fixpoint definition, happens from q . There are finitely many $p_{(q, B', B)}$, all are positive, so they are lower bounded by some positive value p . Playing τ forever thus ensures visiting F almost surely, and iterating this reasoning, even visiting F infinitely often with probability 1. Now, assuming $B' \neq \emptyset$ let δ'_0 be a distribution with support

B' . Any action picked by τ ensures that, whatever the observation, the first belief-component remains in I . Therefore, surely, from distribution δ'_0 the plays stay in the invariant I .

Let us now assume that the triplet (q, B', B) is removed during the iterative computation of the fixed point W_∞ . We prove, by induction on k , that if (q, B', B) is removed at iteration k , then, $(B', B) \notin \text{Win}_{=1}$. If $k = 0$, the pair is removed at initialization, hence $B' \not\subseteq I$ or $B' \not\subseteq B$, and obviously $(B', B) \notin \text{Win}_{=1}$. Otherwise it happens at the k -th iteration, for some $k \geq 1$. Assume, towards a contradiction, that there exists a scheduler τ , witnessing that $(B', B) \in \text{Win}_{=1}$. In particular, there exists a sequence of pairs of actions and observations allowed by the scheduler $(\alpha_1, O_1) \cdots (\alpha_n, O_n)$ so that there exists $q_0 \dots q_n \in Q$ with $q_0 = q$, $q_n \in F$, $\forall i < n, T(q_i, \alpha_{i+1})(q_{i+1}) > 0, \forall 1 \leq j \leq n$ and $\text{Obs}(q_j) = O_j$. Because the triple (q, B', B) was removed at iteration k , it must be that, either (1) for $(B'_2, B_2) = \Delta((B', B), (\alpha_1, O_1) \cdots (\alpha_n, O_n))$, there exists $q_2 \in B_2$ such that $(q_2, B', B) \notin W_{k-1}$, (2) no path corresponding to a sequence $(\alpha_1, O_1) \cdots (\alpha_n, O_n)$ satisfying (1) and starting in q ends in F or (3) there exists an index i and an observation O'_i such that for $(B'_3, B_3) = \Delta((B', B), (\alpha_1, O_1) \cdots (\alpha_{i-1}, O_{i-1})(\alpha_i, O'_i))$ there exists $q \in B_3$, $(q, B'_3, B_3) \notin W_{k-1} \cap Q \times \mathcal{Bl}_{\subseteq I} \times \mathcal{Bl}$. In the first case, it means that either there is a positive probability, under τ to reach a pair of beliefs out of W_{k-1} , and thus out of $\text{Win}_{=1}$ by induction hypothesis. As the sequence of action and observations was chosen so that one can reach F from q , the second case implies that the first case holds with our selected sequence of actions and observations. For the third case, let $(B'_3, B_3) = \Delta((B', B), (\alpha_1, O_1) \cdots (\alpha_{i-1}, O_{i-1})(\alpha_i, O'_i))$. Either there exists $q' \in B_3$ such that $(q, B'_3, B_3) \notin W_{k-1}$, then it is treated similarly to the first case. Else $B'_3 \notin \mathcal{Bl}_{\subseteq I}$. Observe that, in this case, the second requirement on τ is not satisfied since $\mathbb{P}_\tau^{\delta'_0}(\Box I) < 1$. \square

Thanks to Lemma 2, $\text{Win}_{=1}$ can be computed in EXPTIME. Let us now define Lose as the set of beliefs that are clearly losing:

$$\text{Lose} = \{B \in \mathcal{Bl} \mid \neg \exists \tau \forall \delta_0 \text{ with } \text{Supp}(\delta_0) = B, \mathbb{P}_\tau^{\delta_0}(\Box \Diamond F) = 1\} .$$

As established *e.g.* in [3] in the more general framework of 2-player stochastic games with signals, Lose can also be computed in EXPTIME.

Informally, we now consider the set of beliefs from which one can reach, while staying in I , and not risking to fall in Lose, some belief B such that there exists $B' \neq \emptyset$ with $(B', B) \in \text{Win}_{=1}$. Formally, let Win be the following set of beliefs:

$$\begin{aligned} \text{Win} = \{ & B_0 \in \mathcal{Bl} \mid \exists (B', B) \in \text{Win}_{=1} \text{ s.t. } B' \neq \emptyset \text{ and} \\ & \exists \alpha_1 \cdots \alpha_n, \exists O_1 \cdots O_n, \Delta(B_0, (\alpha_1, O_1) \cdots (\alpha_n, O_n)) = B \\ & \forall i \leq n, \forall O'_i, \Delta(B_0, (\alpha_1, O_1) \cdots (\alpha_{i-1}, O_{i-1})(\alpha_i, O'_i)) \notin \text{Lose}\}. \end{aligned}$$

The set Win characterizes winning beliefs, that is, beliefs from which there exists a finite-memory scheduler ensuring at the same time, the Büchi objective $\Box \Diamond F$ almost-surely, and the safety objective $\Box I$ with positive probability. Formally:

Lemma 3 $B_0 \in \text{Win}$ if and only if for every δ_0 with $\text{Supp}(\delta_0) = B_0$, there exists a finite-memory scheduler τ such that $\mathbb{P}_\tau^{\delta_0}(\Box \Diamond F) = 1$ and $\mathbb{P}_\tau^{\delta_0}(\Box I) > 0$.

Proof (of Lemma 3) Assume first that $B_0 \in \text{Win}$. We design a finite memory scheduler τ that is winning from any initial distribution δ_0 with support B_0 . In a first mode, τ aims at reaching a pair of beliefs $(B', B) \in \text{Win}_{=1}$ from B_0 . More precisely, τ plays the path that leads from B_0 to some $B \in \mathcal{Bl}$ such that there exists $B' \neq \emptyset$ with $(B', B) \in \text{Win}_{=1}$. If this succeeds, τ then switches to another mode, where it behaves as the winning scheduler that starts from (B', B) in Lemma 2. If it fails, the play ends in a belief $B_1 \notin \text{Lose}$ (by definition of Win), and from there τ plays to ensure visiting F infinitely often with probability 1. All in all, τ ensures almost surely visiting F infinitely often, and with positive probability (the probability of the prefix leading to B , times the probability that the play is in B' at that time point) to stay in I . Note that the size of the memory τ uses is in $O(|\mathcal{Bl}|^2)$.

Let now δ_0 be an initial distribution with support B_0 , and assume that there exists a finite-memory scheduler τ such that $\mathbb{P}_\tau^{\delta_0}(\Box\Diamond F) = 1$ and $\mathbb{P}_\tau^{\delta_0}(\Box I) > 0$. We consider \mathcal{M}_τ the Markov chain generated by τ , with finite state space $Q \times \text{Mem}$, where Mem is a finite set of memory states. Without loss of generality, we iteratively tag each state of \mathcal{M}_τ with its associated belief. Since τ is winning, there must exist a BSCC \mathcal{C} in \mathcal{M}_τ , reachable from some (q_0, m_0, B_0) via an I -path (a path where all belief tags are included in I), and such that all states $(q, m, B) \in \mathcal{C}$ satisfy $q \in I$, and there exists a state $(q_f, m_f, B_f) \in \mathcal{C}$ such that $q_f \in F$. Pick any state $(q, m, B) \in \mathcal{C}$. From (q, m, B) , under scheduler τ , all plays stay in I . Moreover, for any $q' \in B$, from (q', m, B) , under scheduler τ , almost all plays visit F infinitely often. As a consequence, by the definition of $\text{Win}_{=1}$, $(\{q\}, B) \in \text{Win}_{=1}$. Then, we conclude that $B_0 \in \text{Win}$, exploiting the I -path from (q_0, m_0, B_0) to \mathcal{C} (and thus to any of its states), and the fact that τ ensures $\Box\Diamond F$ almost-surely, and thus always avoids Lose . \square

Win characterizes the winning beliefs, and can be computed in EXPTIME. We thus showed the computability in EXPTIME of the set of supports B from which there exists a finite-memory scheduler τ such that $\mathbb{P}_\tau^B(\Box\Diamond F) = 1$ and $\mathbb{P}_\tau^B(\Box I) > 0$. \square

Now the safe active diagnosis restricted to finite-memory strategies can be reduced to the existence for POMDP of a finite-memory scheduler that ensures a Büchi objective almost surely, and a safety objective with positive probability. As M_C is exponential in the size of \mathcal{C} and the algorithm on the POMDP is in EXPTIME, we obtain a 2EXPTIME complexity upper-bound. Fortunately, in order to avoid a doubly exponential blowup and to establish the EXPTIME complexity, we observe that the exponential comes in both cases from the computation of beliefs depending *only* on the original CLTS. This implies that the safe active probabilistic diagnosis problem is in EXPTIME when restricted to finite-memory strategies.

Corollary 2 *The safe active diagnosis problem restricted to finite-memory strategies is decidable in EXPTIME.*

Proof Given a CLTS \mathcal{C} , we build M_C and decide if $\{q_0\}$ is a support from which there exists a scheduler τ ensuring $\mathbb{P}_\tau^{\{q_0\}}(\Box\Diamond F) = 1$ and $\mathbb{P}_\tau^{\{q_0\}}(\Box I) > 0$ with $I = \{(q, B) \mid q \in Q_c\}$ and $F = \{(q, B) \mid B \subseteq Q_f \vee q \in Q_c\}$. Due to the link between M_C and \mathcal{C} , this choice of F corresponds to runs that are either correct or surely faulty in \mathcal{C} and this choice of I corresponds to runs that are correct. Thus there exists a

finite-memory scheduler τ as defined above iff the corresponding strategy σ in \mathcal{C} ensures safe active diagnosis. Moreover, as explained above the corollary, deciding the existence of this scheduler can be done in EXPTIME. \square

A matching lower-bound is already known from the literature:

Proposition 4 ([2]) *The safe active diagnosis problem restricted to finite-memory strategies is EXPTIME-hard.*

Obviously, this lower bound also holds for the more general problem: on POMDP, whether there exists a finite-memory strategy ensuring a Büchi objective almost-surely and a safety objective with positive probability.

4 Conclusion

We have studied the active diagnosis of partially observable probabilistic transition systems combined with some degradation control. More precisely we have introduced two new notions of degradation both in a qualitative and a quantitative ways. We have established their links with the notion of safety in the finite, infinite and finite controllable cases. Afterwards we have proved that the quantitative versions of the corresponding decision problems were undecidable. Contrary to the safe active diagnosis, the qualitative versions of these problems are EXPTIME-complete even though the associated diagnosers may require infinite memory.

We now have a set of algorithmic results both in the passive and active framework which could justify the development of a tool. At first, this will require to choose and study a more appropriate formalism than probabilistic transition systems from a modelling point of view. Another direction would consist in studying a different notion of faulty runs. Here a run is faulty once a fault has occurred. A fault could only represent a degradation of the system which can still be partially available. In this alternative framework, the degradation to be evaluated would be the evolution of the number of faults in a run w.r.t. its length.

References

1. C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008.
2. N. Bertrand, E. Fabre, S. Haar, S. Haddad, and L. Hélouët. Active diagnosis for probabilistic systems. In *Proceedings of FoSSaCS'14*, volume 8412 of *Lecture Notes in Computer Science*, pages 29–42. Springer, 2014.
3. N. Bertrand, B. Genest, and H. Gimbert. Qualitative determinacy and decidability of stochastic games with signals. In *Proceedings of LICS'09*, pages 319–328. IEEE Computer Society, 2009.
4. N. Bertrand, S. Haddad, and E. Lefaucheux. Foundation of diagnosis and predictability in probabilistic systems. In *Proceedings of FSTTCS'14*, volume 29 of *Leibniz International Proceedings in Informatics*, pages 417–429. Leibniz-Zentrum für Informatik, 2014.
5. K. Chatterjee, L. Doyen, and T. A. Henzinger. A survey of partial-observation stochastic parity games. *Formal Methods in System Design*, 43(2):268–284, Oct 2013.
6. H. Gimbert and Y. Oualhadj. Probabilistic automata on finite words: Decidable and undecidable problems. In *ICALP 2010*, volume 6199 of *Lecture Notes in Computer Science*, pages 527–538. Springer, 2010.

7. S. Haar, S. Haddad, T. Melliti, and S. Schwoon. Optimal constructions for active diagnosis. *Journal of Computer and System Sciences*, 83(1):101–120, 2017.
8. S. Jiang, Z. Huang, V. Chandra, and R. Kumar. A polynomial algorithm for testing diagnosability of discrete-event systems. *IEEE Transactions on Automatic Control*, 46(8):1318–1321, 2001.
9. Christos H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
10. A. Paz. *Introduction to Probabilistic Automata*. Academic Press, 1971.
11. M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, 1998.
12. D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Transactions on Automatic Control*, 50(4):476–492, 2005.