

Diagnosis in Infinite-State Probabilistic Systems (long version)

Nathalie Bertrand, Serge Haddad, Engel Lefaucheux

► **To cite this version:**

Nathalie Bertrand, Serge Haddad, Engel Lefaucheux. Diagnosis in Infinite-State Probabilistic Systems (long version). [Research Report] Inria Rennes; LSV, ENS Cachan. 2016. <hal-01334218>

HAL Id: hal-01334218

<https://hal.inria.fr/hal-01334218>

Submitted on 20 Jun 2016

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

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

Diagnosis in Infinite-State Probabilistic Systems (long version)

Nathalie Bertrand¹, Serge Haddad², and Engel Lefaucheux^{1,2}

¹ Inria, France nathalie.bertrand@inria.fr

² LSV, ENS Cachan & CNRS & Inria, France
{serge.haddad, engel.lefaucheux}@ens-cachan.fr

Abstract

In a recent work, we introduced four variants of diagnosability (FA, IA, FF, IF) in (finite) probabilistic systems (pLTS) depending whether one considers (1) finite or infinite runs and (2) faulty or all runs. We studied their relationship and established that the corresponding decision problems are PSPACE-complete. A key ingredient of the decision procedures was a characterisation of diagnosability by the fact that a random run almost surely lies in an open set whose specification only depends on the qualitative behaviour of the pLTS. Here we investigate similar issues for infinite pLTS. We first show that this characterisation still holds for FF-diagnosability but with a G_δ set instead of an open set and also for IF- and IA-diagnosability when pLTS are finitely branching. We also prove that surprisingly FA-diagnosability cannot be characterised in this way even in the finitely branching case. Then we apply our characterisations for a partially observable probabilistic extension of visibly pushdown automata (POpVPA), yielding EXPSPACE procedures for solving diagnosability problems. In addition, we establish some computational lower bounds and show that slight extensions of POpVPA lead to undecidability.

1 Introduction

Diagnosis. Monitoring (hardware and/or software) systems prone to faults involves several critical tasks: controlling the system to prevent faults as much as possible, deducing the cause of the faults, etc. Most of these tasks assume that an observer has the capability to assess the *status* of the current run based on the outputs of the system: providing information about the possible occurrence of faults. Such an observer is called a *diagnoser* and its associated task is called *diagnosis*. This framework leads to interesting decision and synthesis problems: “Does there exist a diagnoser?” and in the positive case “How to build such a diagnoser?”, “Which kind of diagnoser is sufficient?”, etc. The decision problem, on which we focus here, is called *diagnosability* [14].

Diagnosis of discrete event systems. In order to formally reason about diagnosability, the systems were first modelled by finite labelled transition systems (LTS). Then the specification of a diagnoser is defined by two requirements: *correctness*, meaning that the information provided by the diagnoser is accurate, and *reactivity*, ensuring that a fault will eventually be detected. Within the framework of finite LTS, the decision problem was shown to be solvable in PTIME [9] and it is in fact NLOGSPACE-complete.

Diagnosis of probabilistic systems. A natural way of modelling partially observable systems consists in introducing probabilities (*e.g.* when the design is not fully known or the effects of the interaction with the environment is not predictable). Thus the notion of diagnosability was later extended to Markov chains with labels on transitions, also called probabilistic labelled transition systems (pLTS) [15]. In this context, the reactivity requirement now asks that faults will be almost surely eventually detected. Regarding correctness, two specifications have been proposed: either one sticks to the original definition



and requires that the provided information is accurate, defining *A-diagnosability*; or one weakens the correctness by admitting errors in the provided information that should, however, have an arbitrary small probability defining *AA-diagnosability*. From a computational viewpoint, we recently proved that A-diagnosability is PSPACE-complete [3] and that AA-diagnosability can be solved in PTIME [4].

In case a system is not diagnosable, one may be able to control it, by forbidding some controllable actions, so that it becomes diagnosable. This property of *active diagnosability* has been studied for discrete-event systems [13, 8], and for probabilistic systems [2]. Interestingly, the diagnosability notion in the latter work slightly differs from the original one in [15]. Building on this variation, in [3] semantical issues have been investigated and four relevant notions of diagnosability (FA, IA, FF, IF) have been defined depending on (1) whether one considers finite or infinite runs and (2) faulty or all runs. In finite pLTS, it was shown that all these notions can be characterized by the fact that a random run almost surely lies in an open set, whose specification only depends on the qualitative behaviour of the pLTS.

Diagnosis of infinite-state systems. Diagnosability in infinite-state systems has been studied, on the one hand for restricted Petri nets [5], for which an accurate diagnoser can be designed, and on the other hand for visibly pushdown automata (VPA) [11], for which diagnosability can be decided via the determinisation procedure of [1]. However to the best of our knowledge diagnosis of probabilistic infinite-state systems has not yet been studied.

Contributions. The characterisations of diagnosability established in [3] strongly relied on the finiteness of the models. Our first aim is thus to establish characterisations in the infinite-state case. FF-diagnosability (the original notion of diagnosability) states that almost surely a faulty run will be detected in finite time. We establish that FF-diagnosability can be characterised by the fact that a random run almost surely lies in a G_δ set, only depending on the qualitative behaviour of the system. This characterisation also applies to IF-diagnosability for finitely-branching systems, since then the two notions coincide. An *ambiguous* infinite correct (resp. faulty) run is a run indistinguishable from a faulty (resp. correct) run. IA-diagnosability states that almost surely a run is unambiguous. The set of ambiguous runs is an analytic set (so a priori not known to be a Borel set). However in the finitely-branching case, we establish that the set of unambiguous runs is a G_δ set, yielding a characterisation of IA-diagnosability. FA-diagnosability states that the probability that a finite run is unambiguous goes to 1 when its length goes to infinity. Surprisingly, despite the fact that IA-diagnosability and FA-diagnosability are very close, we prove that FA-diagnosability cannot be characterised by the fact that a random run almost surely lies in a G_δ set. Furthermore we strengthen this result by another inexpressiveness result also related to FA-diagnosability.

We then introduce partially observable probabilistic visibly pushdown automata (POpVPA), a model generating infinite-state probabilistic systems. We show how to exploit the above characterisations to design a decision procedure for diagnosability in POpVPA. More precisely we show that we can “encode” our characterisations in an enlarged probabilistic VPA and then exploit the decision procedures of [7] leading to an EXPSPACE algorithm. Since our characterisations are not regular, this requires some tricky machinery. Finally we complete this work by exhibiting an EXPTIME lower-bound and showing that slight extensions of POpVPA lead to undecidability of the diagnosability problem.

Organisation. In Section 2, we successively introduce probabilistic infinite-state systems, equip them with partial observation and faults, and define diagnosability notions. In Section 3, we establish characterisations of the diagnosability notions and inexpressiveness results. We exploit the characterisations to design decision procedures for POpVPA in Section 4, also

proving hardness and undecidability results. We conclude and give some perspectives in Section 5. All the proofs are given in Appendix.

2 Diagnosis specifications of infinite-state probabilistic systems

2.1 Probabilistic labelled transition systems

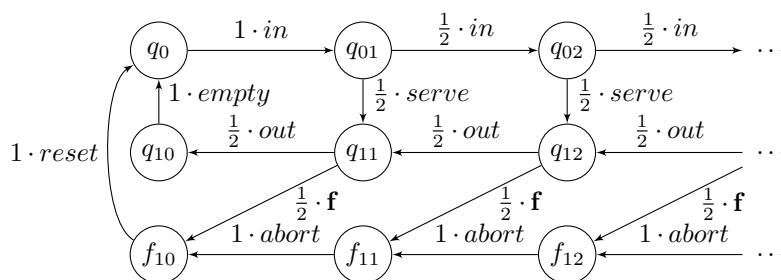
Probabilistic labelled transition systems (pLTS) are labelled transition systems equipped with probability distributions on transitions outgoing from a state.

► **Definition 1.** A pLTS is a tuple $\mathcal{M} = \langle Q, q_0, \Sigma, T, \mathbf{P} \rangle$ where:

- Q is a finite or countable 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;
- $\mathbf{P} : T \rightarrow \mathbb{Q}_{>0}$ is the transition probability fulfilling: $\forall q \in Q, \sum_{(q,a,q') \in T} \mathbf{P}[q, a, q'] = 1$.

Given a pLTS \mathcal{M} , the transition relation of the *underlying LTS* \mathcal{L} is defined by $q \xrightarrow{a} q'$ for $(q, a, q') \in T$; this transition is then said to be *enabled* in q . In order to emphasise the relation between the pLTS and the LTS, we sometimes write $\mathcal{M} = (\mathcal{L}, \mathbf{P})$. Note that since we assume the state space to be at most countable, a pLTS is by definition at most countably branching: from every state q , there are at most countably many transitions enabled in q .

► **Example 2.** The pLTS of Figure 1 represents a server that accepts jobs (event *in*) until it randomly decides to serve the jobs (event *serve*). When a job is done the result is delivered (event *out*). When all jobs are done, the server waits for a new batch of jobs. However randomly, the server may trigger a fault (event *f*) and then abort all remaining jobs (event *abort*). Afterwards, the server is reset (event *reset*). In the figure, the label of a transition (q, a, q') is depicted as $\mathbf{P}[q, a, q'] \cdot a$.



■ **Figure 1** An infinite-state pLTS.

Let us now introduce some important notions and notations that will be used throughout the paper. A *run* ρ of a pLTS \mathcal{M} is a (finite or infinite) sequence $\rho = q_0 a_0 q_1 \dots$ such that for all i , $q_i \in Q$, $a_i \in \Sigma$ and when q_{i+1} is defined, $q_i \xrightarrow{a_i} q_{i+1}$. The notion of run can be generalised, starting from an arbitrary state q . We write Ω for the set of all infinite runs of \mathcal{M} starting from q_0 , assuming the pLTS is clear from context. When it is finite, ρ ends in a state q and its *length*, denoted $|\rho|$, is the number of events occurring in it. Given a finite run $\rho = q_0 a_0 q_1 \dots q_n$ and a (finite or infinite) run $\rho' = q_n a_n q_{n+1} \dots$, we call concatenation of ρ and ρ' and we write $\rho\rho'$ the run $q_0 a_0 q_1 \dots q_n a_n q_{n+1} \dots$; the run ρ is then a *prefix* of $\rho\rho'$, which we denote $\rho \leq \rho\rho'$. The *cylinder* defined by a finite run ρ is the set of all infinite runs that extend ρ : $C(\rho) = \{\rho' \in \Omega \mid \rho \leq \rho'\}$. Cylinders are a basis of open sets for the standard topology on the set of runs (which can be viewed as an infinite tree). One equips a pLTS

with a probability measure on Ω with σ -algebra being \mathcal{B} , the set of Borel sets, and which is uniquely defined by Caratheodory's extension theorem from the probabilities of the cylinders:

$$\mathbb{P}(C(q_0 a_0 q_1 \dots q_n)) = \mathbf{P}[q_0, a_1, q_1] \cdots \mathbf{P}[q_{n-1}, a_{n-1}, q_n] .$$

We will sometimes omit the C and write $\mathbb{P}(\rho)$ for $\mathbb{P}(C(\rho))$. It is well-known that once the measure is fixed, one can enlarge the set of measurable sets by considering the smallest σ -algebra containing \mathcal{B} and the "null" sets: $\{A \mid \exists B \in \mathcal{B} A \subseteq B \wedge \mathbb{P}(B) = 0\}$ and then extend the original measure to a (complete) measure on this enlarged σ -algebra. We consider this measure in the sequel.

The sequence associated with $\rho = qa_0q_1\dots$ is the word $\sigma_\rho = a_0a_1\dots$, and we write indifferently $q \xrightarrow{\rho}^*$ or $q \xrightarrow{\sigma_\rho}^*$ (resp. $q \xrightarrow{\rho}^* q'$ or $q \xrightarrow{\sigma_\rho}^* q'$) for an infinite (resp. finite) run ρ . A state q is *reachable* (from q_0) if there exists a run such that $q_0 \xrightarrow{\rho}^* q$, which we alternatively write $q_0 \rightarrow^* q$. The (infinite) language of pLTS \mathcal{M} consists of all infinite words that label runs of \mathcal{M} and is formally defined as $L^\omega(\mathcal{M}) = \{\sigma \in \Sigma^\omega \mid q_0 \xrightarrow{\sigma}^*\}$.

2.2 Partial observation and faults

The observation of a pLTS is given by a mask function. This function projects every event to its observation. This observation is partial as an event can have no observation or shares its observation with another event, but it is deterministic.

► **Definition 3.** A *partially observable pLTS* (POpLTS) is a tuple $\mathcal{N} = \langle \mathcal{M}, \Sigma_o, \mathcal{P} \rangle$ consisting of a pLTS \mathcal{M} equipped with a mapping $\mathcal{P} : \Sigma \rightarrow \Sigma_o \cup \{\varepsilon\}$ where Σ_o is the set of observations.

Note that our setting generalises most existing frameworks of fault diagnosis by considering a mask function \mathcal{P} onto a possibly different alphabet rather than a partition of the event alphabet into observable and unobservable events. An event $a \in \Sigma$ is said *unobservable* if $\mathcal{P}(a) = \varepsilon$, *fully observable* if $\mathcal{P}(a) \neq \varepsilon$ and $\mathcal{P}^{-1}(\{\mathcal{P}(a)\}) = \{a\}$ and *partially observable* if $\mathcal{P}(a) \neq \varepsilon$ and $|\mathcal{P}^{-1}(\{\mathcal{P}(a)\})| > 1$. The set of unobservable events is denoted Σ_u .

Let $\sigma \in \Sigma^*$ be a finite word; its length is denoted $|\sigma|$. The mapping \mathcal{P} is extended to finite words inductively: $\mathcal{P}(\varepsilon) = \varepsilon$ and $\mathcal{P}(\sigma a) = \mathcal{P}(\sigma)\mathcal{P}(a)$. We say that $\mathcal{P}(\sigma)$ is the *mask* of σ . Write $|\sigma|_o$ for $|\mathcal{P}(\sigma)|$. When σ is an infinite word, its mask is the limit of the masks of its finite prefixes. This mask function is applicable to runs via their associated sequence; it can be either finite or infinite. As usual the mask function is extended to languages. With respect to \mathcal{P} , a POpLTS \mathcal{N} is *convergent* if there is no infinite sequence of unobservable events from any reachable state: $L^\omega(\mathcal{M}) \cap \Sigma^* \Sigma_u^\omega = \emptyset$. When \mathcal{N} is convergent, for every $\sigma \in L^\omega(\mathcal{M})$, $\mathcal{P}(\sigma) \in \Sigma_o^\omega$. In the rest of the paper we assume that POpLTS are convergent. \mathcal{P} can also be viewed as a mapping from runs to Σ_o^ω by defining $\mathcal{P}(q_0 a_0 q_1 a_1 \dots) = \mathcal{P}(a_0 a_1 \dots)$. Remark that this mapping is continuous. We will refer to a *sequence* for a finite or infinite word over Σ , and an *observed sequence* for a finite or infinite sequence over Σ_o . Clearly, the application of the mask function onto Σ_o of a sequence yields an observed sequence.

The *observable length* of a run ρ denoted $|\rho|_o \in \mathbb{N} \cup \{\infty\}$, is the number of observable events that occur in it: $|\rho|_o = |\sigma_\rho|_o$. A *signalling* run is a finite run whose last event is observable. Signalling runs are precisely the relevant runs w.r.t. partial observation issues since each observable event provides an additional information about the execution to an external observer. Given states q, q' and an observed sequence $\sigma \in \Sigma_o^+$, we write $q \xrightarrow{\sigma} q'$ if there is a signalling run from q to q' with observed sequence σ .

In the sequel starting from the initial state q_0 , SR denotes the set of signalling runs, and SR_n the set of signalling runs of observable length n . Since we assume that the POpLTS are convergent, for all $n > 0$, SR_n is equipped with a probability distribution defined by assigning measure $\mathbb{P}(\rho)$ to each $\rho \in \text{SR}_n$. Given ρ a finite or infinite run, and $n \leq |\rho|_o$, $\rho_{\downarrow n}$ denotes the

signalling subrun of ρ of observable length n . For convenience, we consider the empty run q_0 to be the single signalling run, of null length.

2.3 Fault diagnosis for POpLTS

To model the problem of fault diagnosis in POpLTS, we assume the event alphabet Σ contains a special event $\mathbf{f} \in \Sigma$ called the *fault*. A run ρ is then said to be *faulty* if its associated sequence of events contains a fault, *i.e.* $\sigma_\rho \in \Sigma^* \mathbf{f} \Sigma^\omega$; otherwise it is *correct*. The set of faulty (resp. correct) runs is denoted \mathbf{F} (resp. \mathbf{C}). For $n \in \mathbb{N}$, we write \mathbf{F}_n for the set of runs ρ such that $\rho_{\downarrow n}$ is faulty and \mathbf{C}_n for the set of runs ρ such that $\rho_{\downarrow n}$ is correct. By definition, for all n , $\Omega = \mathbf{F}_n \uplus \mathbf{C}_n$, $\mathbf{F} = \bigcup_{n \in \mathbb{N}} \mathbf{F}_n$ and $\mathbf{C} = \bigcap_{n \in \mathbb{N}} \mathbf{C}_n$.

In order to reason about faults we partition sequences of observations into three subsets: an observed sequence $\sigma \in \Sigma_o^\omega$ is *surely correct* if $\mathcal{P}^{-1}(\sigma) \cap \mathbf{L}^\omega(\mathcal{M}) \subseteq (\Sigma \setminus \mathbf{f})^\omega$; it is *surely faulty* if $\mathcal{P}^{-1}(\sigma) \cap \mathbf{L}^\omega(\mathcal{M}) \subseteq \Sigma^* \mathbf{f} \Sigma^\omega$; otherwise, it is *ambiguous*. For finite sequences, we need to rely on signalling runs: a finite observed sequence $\sigma \in \Sigma_o^*$ is *surely faulty* (resp. *surely correct*) if for every signalling run ρ with $\mathcal{P}(\sigma_\rho) = \sigma$, ρ is faulty (resp. correct); otherwise it is ambiguous. A (finite signalling or infinite) run ρ is *surely faulty* (resp. *surely correct*, *ambiguous*) if $\mathcal{P}(\rho)$ is surely faulty (resp. surely correct, ambiguous).

In order to specify various requirements for diagnosability we need to refine the notion of ambiguity. Let \mathcal{N} be a POpLTS and $n \in \mathbb{N}$ with $n \geq 1$. Then:

- \mathbf{FAmb}_∞ (resp. \mathbf{CAmb}_∞) is the set of infinite faulty (resp. correct) ambiguous runs of \mathcal{N} ;
- \mathbf{FAmb}_n (resp. \mathbf{CAmb}_n) is the set of infinite runs of \mathcal{N} whose signalling subrun of observable length n is faulty (resp. correct) and ambiguous;

At this point it is interesting to look at the status of the different subsets of runs we have introduced with respect to the Borel hierarchy. The complementary sets \mathbf{F}_n and \mathbf{C}_n are unions of cylinders; so they are open (and by complementation) closed sets. The set of faulty (resp. correct) runs \mathbf{F} (resp. \mathbf{C}) is an open (resp. closed) set as a union (resp. intersection) of open (resp. closed) sets. The sets \mathbf{FAmb}_n and \mathbf{CAmb}_n are unions of cylinders; so they are open. The sets \mathbf{FAmb}_∞ and \mathbf{CAmb}_∞ may be defined as follows. Consider $(\Sigma_o^2)^\omega$ and Ω^2 both equipped with the product topology. $\mathbf{SameObs} = \{(\rho, \rho') \mid \mathcal{P}(\rho) = \mathcal{P}(\rho')\}$ is the inverse image by a continuous mapping of the closed set $\{(\sigma, \sigma) \mid \sigma \in \Sigma_o^\omega\}$. Therefore $\mathbf{SameObs}$ is closed. Thus $\mathbf{C} \times \mathbf{F} \cap \mathbf{SameObs}$ is a Borel set. The first and second projections are exactly \mathbf{CAmb}_∞ and \mathbf{FAmb}_∞ which establishes that these sets are analytic sets (*i.e.* continuous images of Borel sets). The set of analytic sets is a strict superset of Borel sets but every analytic set is still measurable w.r.t. the complete measure [12, 2H8 p.83].

In the context of finite POpLTS, we introduced four possible specifications of diagnosability [3]. There are two discriminating criteria: whether the non ambiguity requirement holds for faulty runs only or for all runs, and whether ambiguity is defined at the infinite run level or for longer and longer finite signalling subruns. Let \mathcal{N} be a POpLTS. Then:

- \mathcal{N} is *IF-diagnosable* if $\mathbb{P}(\mathbf{FAmb}_\infty) = 0$.
- \mathcal{N} is *IA-diagnosable* if $\mathbb{P}(\mathbf{FAmb}_\infty \uplus \mathbf{CAmb}_\infty) = 0$.
- \mathcal{N} is *FF-diagnosable* if $\limsup_{n \rightarrow \infty} \mathbb{P}(\mathbf{FAmb}_n) = 0$.
- \mathcal{N} is *FA-diagnosable* if $\limsup_{n \rightarrow \infty} \mathbb{P}(\mathbf{FAmb}_n \uplus \mathbf{CAmb}_n) = 0$.


We recall in the next theorem all the implications that hold between these definitions. Missing implications do not hold, already for finite-state POpLTS.

► **Theorem 4** ([3]). *Let \mathcal{N} be a POpLTS. Then*

- \mathcal{N} *FA-diagnosable* \Rightarrow \mathcal{N} *IA-diagnosable* and *FF-diagnosable*;
- \mathcal{N} *IA-diagnosable* or *FF-diagnosable* \Rightarrow \mathcal{N} *IF-diagnosable*;

- If \mathcal{N} is finitely branching, then \mathcal{N} is IF-diagnosable iff \mathcal{N} is FF-diagnosable.

In order to illustrate the different kinds of diagnosability, we describe below some discriminating examples.



■ **Figure 2** Left: a POpLTS that is IF-diagnosable but not IA-diagnosable. Right: a POpLTS that is IA-diagnosable but not FA-diagnosable.

Consider the POpLTS \mathcal{N} on the left of Figure 2 where $\{u, \mathbf{f}\}$ is the set of unobservable events (represented by dashed arrows) and \mathcal{P} is the identity over the other events. A faulty run will almost surely produce a b -event that cannot be mimicked by the single correct run. Thus this POpLTS is IF-diagnosable. The unique correct run $\rho = q_0 u q_1 a q_1 \dots$ has probability $\frac{1}{2}$ and its corresponding observed sequence a^ω is ambiguous. Thus the POpLTS is not IA-diagnosable. This simple example shows that, already for finite-state POpLTS, IF-diagnosability does not imply IA-diagnosability.

Similarly, let us look at the POpLTS on the right of Figure 2 where $\{u, \mathbf{f}\}$ is the set of unobservable events and \mathcal{P} is the identity over the other events. Any infinite faulty run will contain a b -event, and cannot be mimicked by a correct run, therefore $\text{FAmb}_\infty = \emptyset$. The two infinite correct runs have a^ω as observed sequence, and cannot be mimicked by a faulty run, thus $\text{CAmb}_\infty = \emptyset$. As a consequence, this POpLTS is IA-diagnosable. Consider now the infinite correct run $\rho = q_0 u q_1 a q_1 \dots$. It has probability $\frac{1}{2}$, and all its finite signalling subruns are ambiguous since their observed sequence is a^n , for some $n \in \mathbb{N}$. Thus for all $n \geq 1$, $\mathbb{P}(\text{CAmb}_n) \geq \frac{1}{2}$, so that this POpLTS is not FA-diagnosable.

3 Characterisation of diagnosability

The aim of this section is to establish “simple” characterisations of the diagnosability notions for a POpLTS $\mathcal{N} = ((\mathcal{L}, \mathbf{P}), \Sigma_o, \mathcal{P})$ and more precisely to study whether one can express it as a Borel set $B \in \mathcal{B}$ only depending on the underlying LTS \mathcal{L} and the mask function \mathcal{P} , such that almost surely a random run belongs to B if and only if \mathcal{N} is diagnosable. Furthermore if possible, one looks for a set B belonging to a low level of the Borel hierarchy. Observe that for all notions, this requires some machinery since the finite runs-based notions FF and FA are expressed by a family of Borel sets and the infinite runs-based notions IF and IA are expressed by a set which is not *a priori* a Borel set.

Pursuing this goal, we introduce a language **pathL** for specifying Borel sets of runs. It is based on *path formulae*. A path formula α is a predicate over finite prefixes of runs. The (pseudo-)syntax of a formula of **pathL** is:

$$\phi ::= \alpha \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \diamond\phi$$

where α is a path formula. In the sequel we use the standard shortcut $\Box\phi \equiv \neg\diamond\neg\phi$.

A formula is evaluated at some position k of a run $\rho = q_0 a_0 q_1 \dots$. The prefix $\rho[0, k]$ of ρ is defined by $\rho[0, k] = q_0 a_0 q_1 \dots q_k$. The semantics of **pathL** is inductively defined by:

- $\rho, k \models \alpha$ if and only if $\alpha(\rho[0, k])$;
- $\rho, k \models \neg\phi$ if and only if $\rho, k \not\models \phi$;

- $\rho, k \models \phi_1 \wedge \phi_2$ if and only if $\rho, k \models \phi_1$ and $\rho, k \models \phi_2$;
- $\rho, k \models \diamond \phi$ if and only if there exists $k' \geq k$ such that $\rho, k' \models \phi$.


Finally $\rho \models \phi$ if and only if $\rho, 0 \models \phi$. Due to the presence of path formulae (with no restriction) this language subsumes LTL and more generally any ω -regular specification language. In order to reason about the probabilistic behaviour of a POpLTS, we introduce qualitative probabilistic formulae $\mathbb{P}^{\bowtie p}(\phi)$ with $\bowtie \in \{<, >, =\}$, $p \in \{0, 1\}$ and $\phi \in \text{pathL}$. The semantics is obvious: $\mathcal{N} \models \mathbb{P}^{\bowtie p}(\phi)$ if and only if $\mathbb{P}_{\mathcal{N}}(\{\rho \in \Omega \mid \rho \models \phi\}) \bowtie p$. Since pathL is closed by complementation the probabilistic formulae can be restricted to $\mathbb{P}^{=0}(\phi)$ and $\mathbb{P}^{>0}(\phi)$.

Let us give some examples of path formulae. Given a finite run $\rho = q_0 a_0 q_1 \dots q_k$, let \mathbf{f} be defined by $\mathbf{f}(\rho) = \text{true}$ if $a_i = \mathbf{f}$ for some index i . This path formula characterises the faulty finite runs. Let \mathfrak{U} be defined by $\mathfrak{U}(\rho) = \text{true}$ if there exists a correct signalling run ρ' with $\mathcal{P}(\rho) = \mathcal{P}(\rho')$. Using the path formulae \mathbf{f} and \mathfrak{U} , we exhibit a formula of pathL that characterises FF-diagnosability.

► **Proposition 5.** *Let \mathcal{N} be a POpLTS. Then \mathcal{N} is FF-diagnosable iff $\mathcal{N} \models \mathbb{P}^{=0}(\diamond \square (\mathbf{f} \wedge \mathfrak{U}))$.*

Due to Theorem 4, in finitely-branching POpLTS the above characterisation also holds for IF-diagnosability. We also need the finitely-branching assumption in order to characterise IA-diagnosability. To this goal, let us introduce a more intricate path formula. For $\sigma \in \Sigma_o^*$, we define $\text{firstf}(\sigma)$ by $\text{firstf}(\sigma) = \min\{k \mid \exists \rho \text{ signalling run } \mathcal{P}(\rho) = \sigma \wedge \rho_{\downarrow k} \text{ is faulty}\}$ with the convention that $\min(\emptyset) = \infty$. Then the path formula \mathfrak{W} is defined by: $\mathfrak{W}(\varepsilon) = \text{false}$ and $\mathfrak{W}(q_0 a_0 \dots q_{n+1}) = \text{true}$ if $\text{firstf}(\mathcal{P}(q_0 a_0 \dots q_{n+1})) = \text{firstf}(\mathcal{P}(q_0 a_0 \dots q_n)) < \infty$.

► **Proposition 6.** *Let \mathcal{N} be a finitely branching POpLTS. Then \mathcal{N} is IA-diagnosable iff $\mathcal{N} \models \mathbb{P}^{=0}(\diamond \square (\mathfrak{U} \wedge \mathfrak{W}))$.*



■ **Figure 3** An infinitely-branching IA-diagnosable POpLTS.

The POpLTS of Figure 3 illustrates the necessity of the finitely-branching requirement in Proposition 6. $\{u, \mathbf{f}\}$ is the set of unobservable events and \mathcal{P} is the identity over the other events. Observation b occurs in every infinite correct run, while the observed sequence of the single infinite faulty run is a^ω . This POpLTS is thus IA-diagnosable. However, it does not satisfy $\mathbb{P}^{=0}(\diamond \square (\mathfrak{U} \wedge \mathfrak{W}))$ since the unique infinite faulty run has probability $\frac{1}{2}$ and satisfies $\square \mathfrak{U}$. Indeed for every $n \in \mathbb{N}$, there is a correct signalling run with observed sequence a^n .

Observe that the sets of runs specified by the characterisations of FF-diagnosability ($\diamond \square (\mathbf{f} \wedge \mathfrak{U})$) and IA-diagnosability ($\diamond \square (\mathfrak{U} \wedge \mathfrak{W})$) are F_σ sets, *i.e.* countable unions of closed sets. Surprisingly, we show that such a characterisation is impossible for FA-diagnosability.

► **Proposition 7.** *There exists a finitely-branching LTS \mathcal{L} and a mask function \mathcal{P} such that for every F_σ set E of runs, there exists a POpLTS $\mathcal{N} = ((\mathcal{L}, \mathbf{P}), \Sigma_o, \mathcal{P})$ such that:*

- either \mathcal{N} is FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) > 0$;
- or \mathcal{N} is not FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) = 0$.

We conjecture that the previous impossibility result also holds for all Borel sets. The next proposition shows that a positive probability condition (instead of a null condition) may not exist whatever the Borel set.

► **Proposition 8.** *There exists a finitely-branching LTS \mathcal{L} and a mask function \mathcal{P} such that for every Borel set E of runs, there exists a POpLTS $\mathcal{N} = ((\mathcal{L}, \mathbf{P}), \Sigma_o, \mathcal{P})$ such that:*

- *either \mathcal{N} is FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) = 0$;*
- *or \mathcal{N} is not FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) > 0$.*

4 Diagnosis for probabilistic pushdown automata

We now turn to a concrete model for infinite-state POpLTS, namely the ones generated by probabilistic pushdown automata, and more specifically by probabilistic visibly pushdown automata. Our goal is to use the characterisations from the previous section to decide the diagnosability of POpLTS generated by partially observable probabilistic visibly pushdown automata (POpVPA). To do so, we face the difficulty that the Borel sets that characterise IF-, IA- and IF-diagnosability are not *a priori* regular, even in the finite branching case. Yet, for POpVPA, we circumvent this problem, and manage to specify these sets by pLTL formula on a determinisation of the model, tagged with the needed atomic propositions. The decidability of the qualitative model checking for recursive probabilistic systems [7] then yields the decidability of the above three diagnosability notions for POpVPA.

4.1 Probabilistic visibly pushdown automata

Among probabilistic infinite-state systems the ones generated by probabilistic pushdown automata [10, 7] support relevant decision procedures. Already in the non-probabilistic case, the subclass of visibly pushdown automata (VPA) [1] is more tractable than the general model. In VPA, the type of events determines whether the operation on the stack is a push, a pop, or possibly changes the top stack symbol, so that the languages defined by VPA enjoy most of the desirable properties regular languages have.

► **Definition 9.** *A probabilistic visibly pushdown automaton (pVPA) is a tuple $\mathcal{A} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ where:*

- Q is a finite set of control states with q_0 the initial state;
- Σ is a finite alphabet of events, partitionned into local, push and pop events $\Sigma = \Sigma_{\#} \uplus \Sigma_{\#} \uplus \Sigma_{\#}$;
- Γ is a finite alphabet of stack symbols including a set of bottom stack symbols Γ_{\perp} with initial symbol $\perp_0 \in \Gamma_{\perp}$;
- $\delta \subseteq Q \times \Gamma \times \Sigma \times Q \times \Gamma^*$ is the set of transitions such that for every $(q, \gamma, a, q', w) \in \delta$, $|w| \leq 2$, $\gamma \in \Gamma_{\perp}$ implies $w \in \Gamma_{\perp}(\Gamma \setminus \Gamma_{\perp})^*$ and $\gamma \notin \Gamma_{\perp}$ implies $w \in (\Gamma \setminus \Gamma_{\perp})^*$;
- \mathbf{P} is the transition probability function fulfilling for every $q \in Q$ and $\gamma \in \Gamma$:

$$\sum_{(q, \gamma, a, q', w) \in \delta} \mathbf{P}[(q, \gamma, a, q', w)] = 1.$$

A transition $t = (q, \gamma, a, q', w) \in \delta$ is said to be a *local* (resp. *push*, *pop*) transition if $|w| = 1$ (resp. $|w| = 2$, $|w| = 0$). We require that for every transition $t = (q, \gamma, a, q', w) \in \delta$, t is a local (resp. push, pop) transition iff a is a local (resp. push, pop) event.


The semantics of a pVPA is an infinite-state pLTS whose states are pairs (q, z) consisting of a control state and a stack contents.

► **Definition 10.** A pVPA $\mathcal{V} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ defines a pLTS $\mathcal{M}_{\mathcal{V}} = (Q_{\mathcal{V}}, (q_0, \perp_0), \Sigma, T_{\mathcal{V}}, \mathbf{P}_{\mathcal{V}})$ where:

- $Q_{\mathcal{V}} = \{(q, z) \mid q \in Q \wedge z \in \Gamma_{\perp}(\Gamma \setminus \Gamma_{\perp})^*\}$;

- $T_{\mathcal{V}} = \{((q, z\gamma), a, (q', zw)) \mid z\gamma \in \Gamma_{\perp}(\Gamma \setminus \Gamma_{\perp})^* \wedge (q, \gamma, a, q', w) \in \delta\}$;
- For every $((q, z\gamma), a, (q', zw)) \in T_{\mathcal{V}}$, $\mathbf{P}_{\mathcal{V}}[((q, z\gamma), a, (q', zw))] = \mathbf{P}[(q, \gamma, a, q', w)]$.

► **Example 11.** Figure 4 gives an example of a pVPA. The event alphabet is composed of local events $\{serve, empty, reset\}$, a push event in and pop events $\{out, \mathbf{f}, abort\}$. A transition $t = (q, \gamma, a, q', w)$ is represented by an edge from state q to state q' and labelled by $\mathbf{P}[t] \cdot \gamma, a, w$. The semantics of this pVPA is precisely the pLTS from Figure 1. Indeed, the stack alphabet consists of two letters $\Gamma = \{\gamma, \perp_0\}$ where the set of bottom stack symbol is $\Gamma_{\perp} = \{\perp_0\}$. Thus one can encode the stack using a counter that gives the number of γ in the stack. For instance, in the pLTS from Figure 1 the configuration $(q_1, \perp_0\gamma^n)$ of the pVPA corresponds to the state q_{1n} .



■ **Figure 4** A pVPA generating the pLTS from Figure 1 with two finite runs.

To define partially observable pVPA, we equip a pVPA with a mask function and require that only local events may be unobservable, and that pushes and pops can still be distinguished. Thus, the observed sequence of a signalling run of a POpVPA still provides the information about the height of the stack since it is equal to the difference of pushes and pops, plus one.

► **Definition 12.** A *partially observable pVPA* (POpVPA) is a tuple $\langle \mathcal{V}, \Sigma_o, \mathcal{P} \rangle$ consisting of a pVPA \mathcal{V} equipped with a mapping $\mathcal{P} : \Sigma \rightarrow \Sigma_o \cup \{\varepsilon\}$ such that:

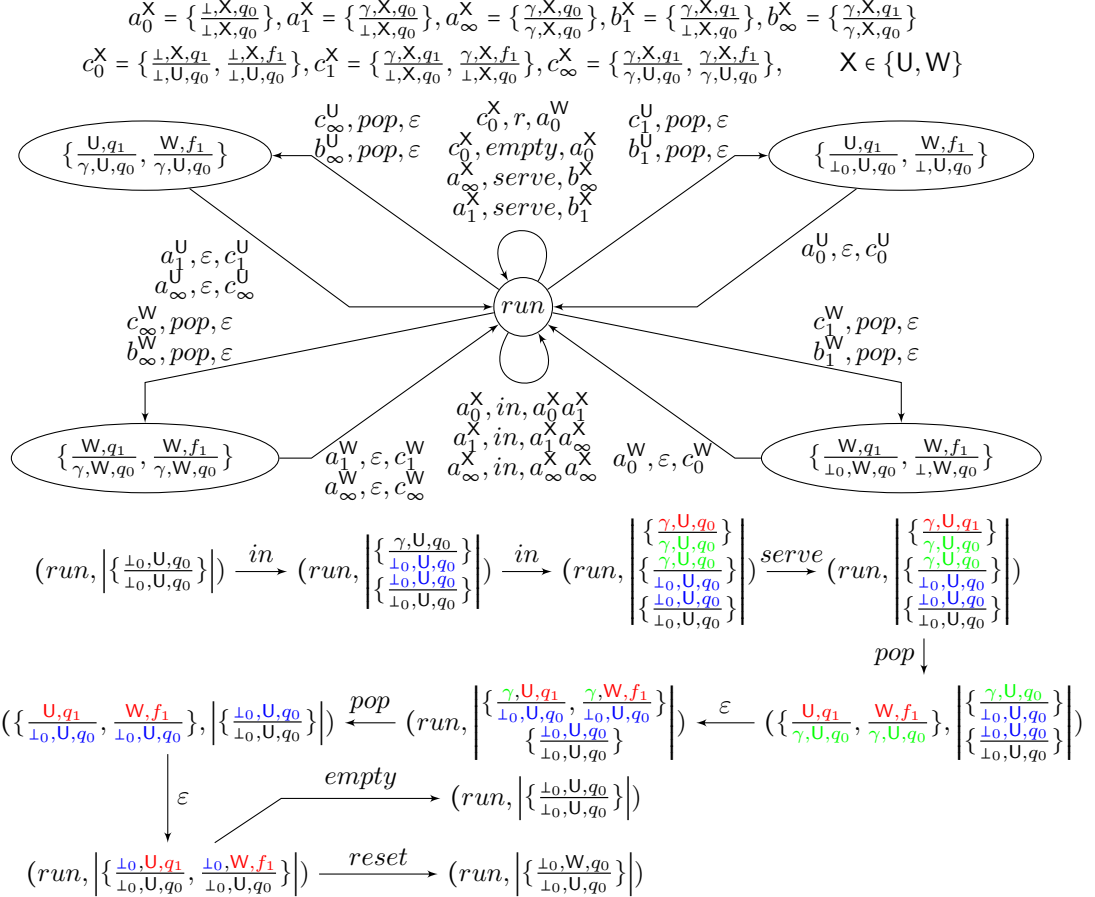
- $\Sigma_o = \Sigma_{o, \natural} \uplus \Sigma_{o, \#} \uplus \Sigma_{o, \flat}$ is the set of observations;
- $\mathcal{P}(\Sigma_{\natural}) \subseteq \Sigma_{o, \natural} \cup \{\varepsilon\}$, $\mathcal{P}(\Sigma_{\#}) \subseteq \Sigma_{o, \#}$ and $\mathcal{P}(\Sigma_{\flat}) \subseteq \Sigma_{o, \flat}$.

In the sequel, we may identify a POpVPA with the POpLTS it generates. In particular, the various concepts of diagnosability are lifted from POpLTS to POpVPA.

4.2 Complexity of diagnosability for POpVPA

To obtain an algorithm for the diagnosability of POpVPA, we follow the finite-state case approach [3]. First, we determinise POpVPA \mathcal{V} into $\mathcal{A}(\mathcal{V})$, with the diagnosis objective in mind, building on the deterministic automaton recognising unambiguous sequences from [8]. We therefore introduce tags that reflect the category of runs (faulty or correct) given an observed sequence with a distinction between “old” and “young” faulty runs. It then suffices to check whether the characterisations hold on the synchronised product $\widehat{\mathcal{V}} \times \mathcal{A}(\mathcal{V})$ where $\widehat{\mathcal{V}}$ enlarges \mathcal{V} by keeping track of a fault occurrence. To reduce to a decidable model checking question, we specify the Borel sets from Section 3 by LTL formulae.

Diagnosis-oriented determinisation. The determinisation of \mathcal{V} (where probabilities are irrelevant for this transformation) into $\mathcal{A}(\mathcal{V})$ exploits some ideas of the original determinisation by Alur and Madhusudan [1], yet, it is customised to diagnosis. In particular, it uses tags that were first defined to construct a deterministic Büchi automaton recognising the unambiguous sequences of a finite LTS [8]. The complete definition of $\mathcal{A}(\mathcal{V})$ is postponed to Appendix B.1. We emphasise here some aspects of the construction and illustrate them on an example.



■ **Figure 5** The VPA $\mathcal{A}(\mathcal{V})$ associated with the POPVPA \mathcal{V} of Figure 4 with two runs.

States and stack symbols. The VPA $\mathcal{A}(\mathcal{V})$ tracks all runs with same observation in parallel memorising their status w.r.t. faults. More precisely to the current set of runs corresponds the symbol on the top of the stack which is a set of tuples where each tuple is written as a fraction $\frac{\gamma, X, q}{\gamma^-, X^-, q^-}$. Let us describe the meaning of this tuple:

- q is the current state of the run and γ is the symbol on the top of its stack;
- $X \in \text{Tg} = \{U, V, W\}$ is the status of the run: U for a correct run, V for a young faulty run and W for an old faulty run;
- The denominator (γ^-, X^-, q^-) , is related to the configuration just after the last push event of the run: γ^- is the stack symbol under the top symbol, while X^- is the status of the run reaching this configuration and q^- the state of this configuration.

A priori, a single state run would be enough. However the simulation of a pop event in the original VPA is performed in two steps requiring some additional states that we explain later.

Illustration. The initial configuration of the VPA $\mathcal{A}(\mathcal{V})$ of Figure 5 ($run, \left\{ \frac{\perp_0, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\}$) corresponds to the empty run represented by a singleton. The denominator of bottom stack symbols is by convention $(\perp_0, \mathbf{U}, q_0)$ and is irrelevant for specifying the transitions of $\mathcal{A}(\mathcal{V})$.

Tag updates. Let us explain how the tag X of an item $\frac{\gamma, X, q}{\gamma, X, q}$ of the current stack symbol is determined. If this item corresponds to a correct run then $X = \mathbf{U}$. When, in a current state, after a transition of $\mathcal{A}(\mathcal{V})$ a (tracked) correct run becomes faulty in the next state, there are two cases. Either there was no tag W in (the numerators of items of) the top stack symbol of the current state then the run is tagged by W . Otherwise it is tagged by V meaning that it is a young faulty run. The tag V (young) becomes W (old) when, in the previous state, there was no tag W in the top stack symbol. A tag W is unchanged along the run.

Push transitions. Given an observed push event $o \in \Sigma_{o, \#}$, from the control state run with top stack symbol bel , there is a looping push transition ($run, bel, o, run, bel' bel''$) in $\mathcal{A}(\mathcal{V})$ that encodes the possible signalling runs with observation o in \mathcal{V} . More precisely for every transition sequence $(q, \alpha) \xrightarrow{o} (r, \beta^- \beta)$ in \mathcal{V} (i.e. a sequence of unobservable local events ending by an event e with $\mathcal{P}(e) = o$) and $\frac{\alpha, X, q}{\alpha, X, q} \in bel$ one inserts $\frac{\beta, Y, r}{\alpha, X, q}$ in bel' and $\frac{\beta, Y, r}{\beta, Y, r}$ in bel'' . The value of Y follows the rules of tag updates.

Illustration. In Figure 5 several transitions correspond to the transition $(q_0, \perp_0, in, q_0, \perp_0 \gamma)$ of \mathcal{V} , including ($run, \left\{ \frac{\perp_0, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\}, in, run, \left\{ \frac{\perp_0, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\} \left\{ \frac{\gamma, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\}$) and several transitions correspond to the transition $(q_0, \gamma, in, q_0, \gamma \gamma)$ of \mathcal{V} , including ($run, \left\{ \frac{\gamma, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\}, in, run, \left\{ \frac{\gamma, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0} \right\} \left\{ \frac{\gamma, \mathbf{U}, q_0}{\gamma, \mathbf{U}, q_0} \right\}$). Here, the specification of the tag updates is straightforward since it does not involve faulty runs. The runs represented in Figure 5 use these two transitions from the initial state.

Local transitions. Given an observed local event $o \in \Sigma_{o, \#}$, from the control state run with top stack symbol bel , there is a looping local transitions (run, bel, o, run, bel') in $\mathcal{A}(\mathcal{V})$ that encodes the possible signalling runs with observation o in \mathcal{V} . More precisely for every transition sequence $(q, \alpha) \xrightarrow{o} (r, \beta)$ in \mathcal{V} (i.e. a sequence of unobservable local events ended by an event e with $\mathcal{P}(e) = o$) and $\frac{\alpha, X, q}{\alpha, X, q} \in bel$ one inserts $\frac{\beta, Y, r}{\alpha, X, q}$ in bel' . The value of Y follows the rules of tag updates.

Illustration. In the VPA $\mathcal{A}(\mathcal{V})$ of Figure 5 there are several transitions corresponding to transition $(q_0, \gamma, serve, q_1, \gamma)$ of \mathcal{V} including ($run, \left\{ \frac{\gamma, \mathbf{U}, q_0}{\gamma, \mathbf{U}, q_0} \right\}, serve, run, \left\{ \frac{\gamma, \mathbf{U}, q_1}{\gamma, \mathbf{U}, q_0} \right\}$). The runs represented in Figure 5 use this transition.

Pop transitions. Given an observed local event $o \in \Sigma_{o, \#}$, from the control state run with top stack symbol bel , the “pop operation” is performed by a sequence of two transitions: a pop transition labelled by o that keeps in the next state all the information needed by the next (local) transition labelled by ε to move back to state run with a consistent stack symbol. Given an intermediate stack symbol, there is exactly one possible such transition. Thus despite these transitions, $\mathcal{A}(\mathcal{V})$ is still deterministic. The first transition ($run, bel, o, \ell, \varepsilon$) in $\mathcal{A}(\mathcal{V})$ is specified as follows. The next state ℓ is a set of items of the following shape $\frac{X, q}{\alpha, X, q}$. More precisely for every transition sequence $(q, \alpha) \xrightarrow{o} (r, \varepsilon)$ in \mathcal{V} (i.e. a sequence of unobservable local events ended by an event e with $\mathcal{P}(e) = o$) and $\frac{\alpha, X, q}{\alpha, X, q} \in bel$ one inserts $\frac{Y, r}{\alpha, X, q}$ in ℓ . The value of Y follows the rules of tag updates. A transition ($\ell, bel, \varepsilon, run, bel'$) is specified as follows. For every $\frac{X', q'}{\gamma, X, q}$ in ℓ and $\frac{\gamma, X, q}{\gamma, X, q}$ in bel (i.e. the denominator of the first fraction and the numerator of the second fraction match), one inserts $\frac{\gamma, X', q'}{\gamma, X, q}$ in bel' .

Illustration. Let us describe how the *pop* event is performed by two transitions in the runs of the VPA of Figure 5 from the state reached after event *serve*. From q_1 with γ as top of the stack there are two transitions whose observation is *pop*: $(q_1, \gamma, out, q_1, \varepsilon)$ and $(q_1, \gamma, \mathbf{f}, f_1, \varepsilon)$. Thus starting from run with top stack symbol $\left\{ \frac{\gamma, \mathbf{U}, q_1}{\gamma, \mathbf{U}, q_0} \right\}$, one reaches state $\ell = \left\{ \frac{\mathbf{U}, q_1}{\gamma, \mathbf{U}, q_0}, \frac{W, f_1}{\gamma, \mathbf{U}, q_0} \right\}$. The faulty run is tagged with W as there was no tag W in the former top stack symbol. In

the next configuration, the top stack symbol is $\{\frac{\gamma, U, q_0}{\perp_0, U, q_0}\}$. So the transition labelled by ε moves back to state *run* with updated top stack symbol $\{\frac{\gamma, U, q_1}{\perp_0, U, q_0}, \frac{\gamma, W, f_1}{\perp_0, U, q_0}\}$.

Product VPA. We first define $\widehat{\mathcal{V}}$ whose set of states \widehat{Q} is a duplication of Q in correct states Q_c and faulty states Q_f . Given a transition of \mathcal{V} starting from q leading to q' , there is in $\widehat{\mathcal{V}}$ a transition starting from q_f leading to q'_f and a transition starting from q_c leading either to q'_c if the event is not **f** or to q'_f otherwise. We then construct $\mathcal{V}_{\mathcal{A}(\mathcal{V})} = \widehat{\mathcal{V}} \times \mathcal{A}(\mathcal{V})$ the product automaton of $\widehat{\mathcal{V}}$ and $\mathcal{A}(\mathcal{V})$ synchronised on the alphabet of observed events Σ_o . The transitions of $\widehat{\mathcal{V}}$ labelled by unobservable events do not change the second component of the state and the transitions of $\mathcal{A}(\mathcal{V})$ labelled by ε do not change the first component of the state. Due to the determinism of $\mathcal{A}(\mathcal{V})$, $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ has the same probabilistic behaviour as the one of \mathcal{V} except that it memorises additional information along the run. More precisely, let ρ be a run of \mathcal{V} , then $\bar{\rho}$, a run of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$, is obtained from ρ by following the same transitions and adding the single \ominus transition firable after any pop transition. One immediately gets $\mathbb{P}_{\mathcal{V}_{\mathcal{A}(\mathcal{V})}}(\bar{\rho}) = \mathbb{P}_{\mathcal{V}}(\rho)$.

Let us explain how to transform the paths formulae \mathfrak{f} , \mathfrak{U} and \mathfrak{W} into atomic propositions on the pairs $((q, \text{run})(\gamma, \text{bel}))$ consisting of a control state of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ together with a top stack contents. For path formula \mathfrak{f} , we define the corresponding atomic proposition ν_f by $\nu_f((q, \text{run})(\gamma, \text{bel})) = \text{true}$ if and only if $q \in Q_f$. Let $\text{bel} \subseteq (\Gamma \times \text{Tg} \times Q)^2$, we say that \mathfrak{X} occurs in bel if there exists $\frac{\gamma, \mathfrak{X}, q}{\gamma, \mathfrak{X}, q} \in \text{bel}$. We define atomic propositions ν_u and ν_w by: $\nu_u((q, \text{run})(\gamma, \text{bel})) = \text{true}$ if and only if **U** occurs in bel ; and $\nu_w((q, \text{run})(\gamma, \text{bel})) = \text{true}$ if and only if **W** occurs in bel .

Given a run ρ of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$, we write $\text{last}(\rho)$ for the pair formed of the control state and top stack symbol in $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ after ρ . The atomic propositions ν_f and ν_u perfectly reflect the paths formula \mathfrak{f} and \mathfrak{U} , and ν_w is eventually forever true if and only if \mathfrak{W} is.

► **Proposition 13.** *Let ρ be an infinite run of \mathcal{V} . Then:*

- For all $k \in \mathbb{N}$, $\mathfrak{f}(\rho_{\downarrow k}) \Leftrightarrow \nu_f(\text{last}(\bar{\rho}_{\downarrow k}))$ and $\mathfrak{U}(\rho_{\downarrow k}) \Leftrightarrow \nu_u(\text{last}(\bar{\rho}_{\downarrow k}))$;
- $\rho \models \diamond \square \mathfrak{W} \Leftrightarrow \exists K \forall k \geq K. \nu_w(\text{last}(\bar{\rho}_{\downarrow k})) = \text{true}$.

Thanks to the relationships between the paths formulae, and the atomic propositions, and using the characterisations from Section 3, we manage to reduce the FF-, IF- and IA-diagnosis to the model checking of a pLTL formula on the product VPA $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$. Model checking qualitative pLTL for probabilistic pushdown automata is doable in polynomial space in the size of the model [7]. In our case, $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ is exponential in the size of \mathcal{V} . We thus obtain the decidability and a complexity upper-bound for the diagnosability problems for POpVPA.

► **Theorem 14.** *FF-diagnosability, IF-diagnosability and IA-diagnosability are decidable in EXPSpace for POpVPA.*

Reducing the universality problem for VPA, which is known to be EXPTIME-complete [1], we obtain the EXPTIME-hardness of all diagnosability variants for POpVPA.

► **Theorem 15.** *Diagnosability is EXPTIME-hard for POpVPA.*

The restriction to visibly pushdown automata is motivated by the unfeasibility of diagnosis for general probabilistic pushdown automata. The undecidability can be obtained by adapting the proof for diagnosis of *non-probabilistic* pushdown automata [11]. However, in order to show how robust the result is, we rather reduce from the Post Correspondence Problem and prove the undecidability of diagnosability for restricted classes of partially observable probabilistic pushdown automata, see Theorems 23 and 24 in Appendix B.4.

5 Conclusion

We studied the diagnosability problem for infinite-state probabilistic systems, both from a semantical perspective, and from an algorithmic one when considering probabilistic visibly pushdown automata. A natural research aim is to reduce the complexity gap for the diagnosability of POPVPA (currently EXPTIME-hard and in EXPSPACE). We could also investigate the diagnosability problem for other probabilistic extensions infinite state systems, such as lossy channel systems or VASS. Another research direction would be to consider the fault diagnosis problem for continuous-time probabilistic models, starting with CTMC.

References

- 1 R. Alur and P. Madhusudan. Visibly pushdown languages. In *Proc. STOC'04*, pages 202–211. ACM, 2004.
- 2 N. Bertrand, É. Fabre, S. Haar, S. Haddad, and L. Hélouët. Active diagnosis for probabilistic systems. In *Proc. FoSSaCS'14*, volume 8412 of *LNCS*, pages 29–42. Springer, 2014.
- 3 N. Bertrand, S. Haddad, and E. Lefaucheux. Foundation of diagnosis and predictability in probabilistic systems. In *Proc. FSTTCS'14*, volume 29 of *LIPICs*, pages 417–429. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.
- 4 N. Bertrand, S. Haddad, and E. Lefaucheux. Accurate approximate diagnosability of stochastic systems. In *Proc. LATA'16*, volume 9618 of *LNCS*, pages 549–561. Springer, 2016.
- 5 M. P. Cabasino, A. Giua, and C. Seatzu. Diagnosability of discrete-event systems using labeled Petri nets. *IEEE Trans. Automation Science and Engineering*, 11(1):144–153, 2014.
- 6 K. Etessami and M. Yannakakis. Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations. *J. ACM*, 56(1), 2009.
- 7 K. Etessami and M. Yannakakis. Model checking of recursive probabilistic systems. *ACM Trans. Computational Logic*, 13(2):12, 2012.
- 8 S. Haar, S. Haddad, T. Melliti, and S. Schwoon. Optimal constructions for active diagnosis. In *Proc. FSTTCS'13*, volume 24 of *LIPICs*, pages 527–539. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- 9 S. Jiang, Z. Huang, V. Chandra, and R. Kumar. A polynomial algorithm for testing diagnosability of discrete-event systems. *IEEE Trans. Automatic Control*, 46(8):1318–1321, 2001.
- 10 A. Kučera, J. Esparza, and R. Mayr. Model checking probabilistic pushdown automata. *Logical Methods in Computer Science*, 2(1), 2006.
- 11 C. Morvan and S. Pinchinat. Diagnosability of pushdown systems. In *Proceedings of HVC'09*, volume 6405 of *LNCS*, pages 21–33. Springer, 2009.
- 12 Y. N. Moschovakis. *Descriptive Set Theory*. Mathematical Surveys and Monographs. AMS, 2009.
- 13 M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Trans. Automatic Control*, 43(7):908–929, 1998.
- 14 M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE Trans. Automatic Control*, 40(9):1555–1575, 1995.
- 15 D. Thorsley and D. Teneketzis. Diagnosability of stochastic discrete-event systems. *IEEE Trans. Automatic Control*, 50(4):476–492, 2005.

A Proofs for Section 3

► **Proposition 5.** *Let \mathcal{N} be a POPLTS. Then \mathcal{N} is FF-diagnosable iff $\mathcal{N} \models \mathbb{P}^0(\diamond \square (\mathfrak{f} \wedge \mathfrak{U}))$.*

Proof. Consider the set of fault-triggering runs:

$$\mathfrak{R} = \{ \rho = q_0 a_0 q_1 \dots a_{k-1} q_k \mid a_{k-1} = \mathfrak{f} \wedge \forall i < k-1, a_i \neq \mathfrak{f} \} .$$

Write $E = \{ \rho \in \Omega \mid \diamond \square (\mathfrak{f} \wedge \mathfrak{U}) \}$ for the set of runs we are interested in. We further define, for every $\rho \in \mathfrak{R}$, $E_\rho = \{ \rho' \in \Omega \mid \rho \leq \rho' \wedge \rho' \models \square \mathfrak{U} \}$ and for every $n \in \mathbb{N}$, $E_\rho^n = \{ \rho' \in \Omega \mid \rho \leq \rho' \wedge \rho' \models \square^n \mathfrak{U} \}$ where $\rho \models \square^n \phi$ if for every $k \leq n$, $\rho, k \models \phi$. Observe that $E = \bigcup_{\rho \in \mathfrak{R}} E_\rho$ and that $E_\rho = \bigcap_{n \in \mathbb{N}} E_\rho^n$. Thus $\mathbb{P}(E) = \sum_{\rho \in \mathfrak{R}} \mathbb{P}(E_\rho)$ and $\lim_{n \rightarrow \infty} \mathbb{P}(E_\rho^n) = \mathbb{P}(E_\rho)$.

• Assume first that $\mathbb{P}(E) > 0$. Then, there exists $\rho \in \mathfrak{R}$ such that $\mathbb{P}(E_\rho) > 0$. By definition, for every $n > |\rho|_o$ $\mathbb{P}(\text{FAmb}_n) \geq \mathbb{P}(E_\rho)$. Thus, \mathcal{N} is not FF-diagnosable.

• Assume now that $\mathbb{P}(E) = 0$. So, for every $\rho \in \mathfrak{R}$, $\mathbb{P}(E_\rho) = 0$. Let us pick some $\varepsilon > 0$. Since $F = \bigcup_{n \in \mathbb{N}} F_n$, there exists n_0 such that for every $n \geq n_0$, $\mathbb{P}(F \setminus F_n) \leq \frac{\varepsilon}{3}$. Let $\mathfrak{R}' = \{ \rho \in \mathfrak{R} \mid |\rho|_o < n_0 \}$. Pick a finite subset \mathfrak{R}'' of \mathfrak{R}' such that $\sum_{\rho \in \mathfrak{R}' \setminus \mathfrak{R}''} \mathbb{P}(\rho) \leq \frac{\varepsilon}{3}$. Define $K = |\mathfrak{R}''|$. Let n_1 be such that for every $n \geq n_1$ and every $\rho \in \mathfrak{R}''$, $\mathbb{P}(E_\rho^n) \leq \frac{\varepsilon}{3K}$. Observe now that for every $n \geq n_0$, $\text{FAmb}_n \subseteq (F \setminus F_n) \cup \bigcup_{\rho \in \mathfrak{R}' \setminus \mathfrak{R}''} C(\rho) \cup \bigcup_{\rho \in \mathfrak{R}''} E_\rho^n$. Thus, for every $n \geq n_1$, $\mathbb{P}(\text{FAmb}_n) \leq \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + K \frac{\varepsilon}{3K} = \varepsilon$. Since ε is arbitrary, \mathcal{N} is FF-diagnosable. ◀

► **Proposition 6.** *Let \mathcal{N} be a finitely branching POPLTS. Then \mathcal{N} is IA-diagnosable iff $\mathcal{N} \models \mathbb{P}^0(\diamond \square (\mathfrak{U} \wedge \mathfrak{W}))$.*

Proof. It is enough to show that $\rho \in \Omega$ is ambiguous if and only if $\rho \models \diamond \square (\mathfrak{U} \wedge \mathfrak{W})$. We focus below on correct runs; the case of faulty runs is similar and even simpler.

• Let $\rho \in \text{CAmb}_\infty$. Since ρ is ambiguous, there exists a faulty run ρ' such that $\mathcal{P}(\rho') = \mathcal{P}(\rho)$. Let k_0 be such that $\rho'_{\downarrow k_0}$ is faulty. Thus for all $k \geq k_0$, $\text{firstf}(\mathcal{P}(\rho_{\downarrow k})) \leq k_0$ and in addition it is non decreasing. So there exists some $k_1 \geq k_0$ such that for all $k \geq k_1$, $\text{firstf}(\mathcal{P}(\rho_{\downarrow k}))$ is constant. We thus obtain $\rho \models \diamond \square \mathfrak{W}$. Moreover, since $\rho \models \square \mathfrak{U}$, we conclude that $\rho \models \diamond \square (\mathfrak{U} \wedge \mathfrak{W})$.

• Let ρ be a correct run such that $\rho \models \diamond \square (\mathfrak{U} \wedge \mathfrak{W})$. Thus there is a position k_0 such that for all $k \geq k_0$, $\rho, k \models \mathfrak{W}$. In particular, by definition of \mathfrak{W} , for all $k \geq k_0$, there is a finite signalling run $\rho^{(k)}$ such that $\mathcal{P}(\rho^{(k)}) = \mathcal{P}(\rho_{\downarrow k})$ and $\rho'_{\downarrow k_0}$ is faulty. Consider the tree of these runs $\rho^{(k)}$ by merging the common prefixes. This tree is finitely branching and infinite. By König's lemma, it must admit an infinite branch, corresponding to a run ρ' with $\mathcal{P}(\rho') = \mathcal{P}(\rho)$ and $\rho'_{\downarrow k_0}$ faulty. We deduce that ρ is ambiguous. ◀


Let us recall some standard facts about Borel sets and measures. A set F is closed if and only if $F = \bigcap_{n \in \mathbb{N}} O_n$ where O_n is a union of cylinders defined by $O_n = \{ C(\rho) \mid |\rho| = n \wedge \exists \rho' \in F, \rho \leq \rho' \}$. Thus an F_σ set F can be written as $F = \bigcup_{m \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} O_{m,n}$ where $O_{m,n}$ is a union of cylinders whose associated paths have length n . Without loss of generality, the sequence of closed sets may be chosen as a non decreasing sequence. The measures we have defined in the core of the paper are *regular*. In particular, for every measurable set E such that $\mathbb{P}(E) > 0$, there exists a closed set $F \subseteq E$ such that $\mathbb{P}(F) > 0$.

► **Proposition 7.** *There exists a finitely-branching LTS \mathcal{L} and a mask function \mathcal{P} such that for every F_σ set E of runs, there exists a POPLTS $\mathcal{N} = ((\mathcal{L}, \mathbf{P}), \Sigma_o, \mathcal{P})$ such that:*

- either \mathcal{N} is FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) > 0$;
- or \mathcal{N} is not FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) = 0$.

Proof. Consider the LTS $\mathcal{L} = \langle Q, q_0, \Sigma, T \rangle$ defined as follows and let the mask function be defined by: $\mathcal{P}(u) = \mathcal{P}(\mathfrak{f}) = \varepsilon$ and \mathcal{P} is the identity over the other events.

- $Q = \{f_1, q_f\} \cup \{q_i \mid i \in \mathbb{N}\}$;
- $\Sigma = \{a, b, c, u, \mathbf{f}\}$;
- $T = \{(q_0, u, q_f), (q_0, u, q_1), (q_f, a, q_f), (q_f, b, q_f), (q_f, \mathbf{f}, f_1), (f_1, b, f_1), (f_1, c, f_1)\} \cup \{(q_i, a, q_{i+1}), (q_i, b, q_{i+1})\}_{i \geq 1}$.



■ **Figure 6** A family of POPLTS whose underlying LTS has no appropriate characterisation of FA-diagnosability.

We consider a family of POPLTS, represented in Figure 6, with underlying LTS \mathcal{L} . For $\mathbf{p} = (p_n)_{n \geq 1}$ a sequence of probabilities, we define the POPLTS $\mathcal{N}_{\mathbf{p}} = ((\mathcal{L}, \mathbf{P}_{\mathbf{p}}), \Sigma_o, \mathcal{P})$ in which for every $n \geq 1$ the probability that b occurs from state q_n is $\mathbf{P}_{\mathbf{p}}(q_n, b, q_{n+1}) = p_n$, and all other probabilities are independent of \mathbf{p} : $\mathbf{P}_{\mathbf{p}}(q_0, u, q_f) = \mathbf{P}_{\mathbf{p}}(q_0, u, q_1) = \mathbf{P}_{\mathbf{p}}(f_1, b, f_1) = \mathbf{P}_{\mathbf{p}}(f_1, c, f_1) = \frac{1}{2}$, $\mathbf{P}_{\mathbf{p}}(q_f, a, q_f) = \mathbf{P}_{\mathbf{p}}(q_f, b, q_f) = \mathbf{P}_{\mathbf{p}}(q_f, \mathbf{f}, f_1) = \frac{1}{3}$.

Observe that $\lim_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n) = 0$ and $\mathbb{P}(\text{CAmb}_{n-1}) = p_n + \frac{2^{n-1}}{3^n}$. Therefore, $\mathcal{N}_{\mathbf{p}}$ is FA-diagnosable iff $\lim_{n \rightarrow \infty} p_n = 0$.

Let E be an arbitrary F_{σ} set. Pick some FA-diagnosable $\mathcal{N}_{\mathbf{p}}$ *i.e.* with $\lim_{n \rightarrow \infty} p_n = 0$. If $\mathbb{P}_{\mathbf{p}}(E) > 0$ where $\mathbb{P}_{\mathbf{p}}$ is the probability measure of this POPLTS, we are done. Assume thus that $\mathbb{P}_{\mathbf{p}}(E) = 0$. In order to define a second POPLTS, via \mathbf{p}' , consider an infinite increasing sequence $\{n_j\}_{j \leq 1}$ and let for $n \notin \{n_j\}_{j \leq 1}$, $p'_n = p_n$ and for $n \in \{n_j\}_{j \geq 1}$, $p'_n = \frac{1}{2}$. Due to the sub-sequence $p'_{n_j} = \frac{1}{2}$, $\mathcal{N}_{\mathbf{p}'}$ is not FA-diagnosable. The sequence $\{n_j\}_{j \leq 1}$ depends on $\mathbb{P}_{\mathbf{p}}$ and will be defined after some preliminary observations.

Let $F = \{\rho \mid q_0 u q_1 \leq \rho\}$. Denoting $\mathbb{P}_{\mathbf{p}'}$ the probability measure of the second POPLTS, observe that $\mathbb{P}_{\mathbf{p}'}(E \setminus F) = \mathbb{P}_{\mathbf{p}}(E \setminus F) = 0$. Using the above discussion, the F_{σ} set $E \cap F = \bigcup_{m \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} O_{m,n}$ where for all m, n , $O_{m,n}$ is a disjoint union of cylinders $C(\rho)$ with $|\rho| = n$, $O_{m,n+1} \subseteq O_{m,n}$ and $O_{m,n} \subseteq O_{m+1,n}$. Denote $F_m = \bigcap_{n \in \mathbb{N}} O_{m,n}$. For all m , $\lim_{n \rightarrow \infty} \mathbb{P}_{\mathbf{p}}(O_{m,n}) = \mathbb{P}_{\mathbf{p}}(E \cap F_m) \leq \mathbb{P}_{\mathbf{p}}(E \cap F) = 0$.

- n_1 is chosen such that for all $n \geq n_1$, $p_n \leq \frac{1}{2}$. Observe now that for all n_j ,

$$p'_{n_j} = \frac{1}{2} = \frac{1}{2p_{n_j}} p_{n_j} \quad \text{and} \quad 1 - p'_{n_j} = \frac{1}{2} \leq 1 - p_{n_j} \leq \frac{1}{2p_{n_j}} (1 - p_{n_j})$$

By definition of $\mathbf{P}_{\mathbf{p}'}$, since $O_{m,n}$ is a disjoint union of cylinders $C(\rho)$ with $|\rho| = n$, applying inductively the previous inequalities, for all n such that $n_k < n \leq n_{k+1}$ (denoting $n_0 = 0$):

$$\mathbb{P}_{\mathbf{p}'}(O_{m,n}) \leq \frac{\mathbb{P}_{\mathbf{p}}(O_{m,n})}{2^k \prod_{1 \leq j < k} p_{n_j}}. \quad (1)$$

- Assume that we have chosen n_1, \dots, n_k . Since $\lim_{n \rightarrow \infty} \mathbb{P}_{\mathbf{p}}(O_{k,n}) = 0$, there exists $n_{k+1} > n_k$ such that $\mathbb{P}_{\mathbf{p}}(O_{k,n_{k+1}}) \leq \prod_{1 \leq j \leq k} p_{n_j}$. We choose such an index.


Equation 1 now implies that for all $m \leq k$, $\mathbb{P}_{\mathbf{p}'}(O_{m,n_{k+1}}) \leq \mathbb{P}_{\mathbf{p}'}(O_{k,n_{k+1}}) \leq \frac{1}{2^k}$. Thus for all m , $\mathbb{P}_{\mathbf{p}'}(F_m) = \lim_{k \rightarrow \infty} \mathbb{P}_{\mathbf{p}'}(O_{m,n_{k+1}}) = 0$. Since $E \cap F = \bigcup_{m \in \mathbb{N}} F_m$, $\mathbb{P}_{\mathbf{p}'}(E \cap F) = 0$ and so $\mathbb{P}_{\mathbf{p}'}(E) = 0$. ◀

► **Proposition 8.** *There exists a finitely-branching LTS \mathcal{L} and a mask function \mathcal{P} such that for every Borel set E of runs, there exists a POpLTS $\mathcal{N} = ((\mathcal{L}, \mathbf{P}), \Sigma_o, \mathcal{P})$ such that:*

- *either \mathcal{N} is FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) = 0$;*
- *or \mathcal{N} is not FA-diagnosable and $\mathbb{P}_{\mathcal{N}}(E) > 0$.*

Proof. Consider the LTS $\mathcal{L} = \langle Q, q_0, \Sigma, T \rangle$ defined as follows, and let the mask function be defined by: $\mathcal{P}(u) = \mathcal{P}(\mathbf{f}) = \varepsilon$ and \mathcal{P} is the identity over the other events.

- $Q = \{f_1, q_f, q_0\} \cup \{q_w \mid w \in (a+b)^*\}$;
- $\Sigma = \{a, b, c, u, \mathbf{f}\}$;
- $T = \{(q_0, u, q_f), (q_0, u, q_1), (q_f, a, q_f), (q_f, b, q_f), (q_f, \mathbf{f}, f_1), (f_1, b, f_1), (f_1, c, f_1)\} \cup \{(q_w, a, q_{wa}), (q_w, b, q_{wb})\}_{w \in (a+b)^*}$.



■ **Figure 7** Another family of POpLTS whose underlying LTS has no appropriate characterisation of FA-diagnosability.

We consider a family of POpLTS, represented in Figure 7, with underlying LTS \mathcal{L} , parameterised by a mapping $\mathbf{p} : (a+b)^* \rightarrow (0, 1)$. Let $\mathcal{N}_{\mathbf{p}} = ((\mathcal{L}, \mathbf{P}_{\mathbf{p}}), \Sigma_o, \mathcal{P})$ be the POpLTS such that the probability that b occurs from state q_w is $\mathbf{P}_{\mathbf{p}}(q_w, b, q_{wb}) = \mathbf{p}(w)$, and all other probabilities are independent from \mathbf{p} : $\mathbf{P}_{\mathbf{p}}(q_0, u, q_f) = \mathbf{P}_{\mathbf{p}}(q_0, u, q_1) = \mathbf{P}_{\mathbf{p}}(f_1, b, f_1) = \mathbf{P}_{\mathbf{p}}(f_1, c, f_1) = \frac{1}{2}$, $\mathbf{P}_{\mathbf{p}}(q_f, a, q_f) = \mathbf{P}_{\mathbf{p}}(q_f, b, q_f) = \mathbf{P}_{\mathbf{p}}(q_f, \mathbf{f}, f_1) = \frac{1}{3}$. In the sequel, for convenience, we also write $\mathbf{p}(w, b)$ for $\mathbf{p}(w)$, and define $\mathbf{p}(w, a) = 1 - \mathbf{p}(w)$, so that $\mathbf{P}_{\mathbf{p}}(q_w, a, q_{wa}) = \mathbf{p}(w, a)$.

Word w can be decomposed into letters $w = w[1] \dots w[n]$, and we give notations for factors: $w[1, k] = w[1] \dots w[k]$ with the convention that $w[1, 0] = \varepsilon$. Finally we define $p_{\mathbf{p}}(w) = \prod_{1 \leq k \leq n} \mathbf{p}(w[1, k-1], w[k])$, as the probability to read w from q_{ε} . Since $\lim_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n) = 0$ and $\mathbb{P}(\text{CAmb}_{n-1}) = \sum_{|w|=n-1} \mathbf{p}(w, b) + \frac{2^{n-1}}{3^n}$, we deduce that $\mathcal{N}_{\mathbf{p}}$ is FA-diagnosable iff $\lim_{n \rightarrow \infty} \sum_{|w|=n-1} \mathbf{p}(w, b) = 0$.

Let E be an arbitrary measurable set. Pick some POpLTS $\mathcal{N}_{\mathbf{p}}$ which is FA-diagnosable, *i.e.* with $\lim_{n \rightarrow \infty} \sum_{|w|=n-1} \mathbf{p}(w, b) = 0$. If $\mathbb{P}_{\mathbf{p}}(E) = 0$ where $\mathbb{P}_{\mathbf{p}}$ is the probability of this POpLTS, we are done. Assume therefore that $\mathbb{P}_{\mathbf{p}}(E) > 0$. Let $F = \{\rho \mid q_0 u q_{\varepsilon} \sqsubseteq \rho\}$ be the set of runs starting with a u -transition to q_{ε} . Denoting $\mathbb{P}_{\mathbf{p}'}$ the probability measure of any other POpLTS $\mathcal{N}_{\mathbf{p}'}$, observe that $\mathbb{P}_{\mathbf{p}'}(E \setminus F) = \mathbb{P}_{\mathbf{p}}(E \setminus F)$. So, if $\mathbb{P}_{\mathbf{p}'}(E \setminus F) > 0$, then by picking any non FA-diagnosable $(\mathcal{L}, \mathbf{P}_{\mathbf{p}'})$, we are done. So assume $\mathbb{P}_{\mathbf{p}'}(E \setminus F) = 0$ which implies $\mathbb{P}_{\mathbf{p}'}(E \cap F) > 0$. Using our recalls, there exists a closed set $G \subseteq E \cap F$ with $\mathbb{P}_{\mathbf{p}'}(G) > 0$.

If $G = F$ then $\mathbb{P}_{\mathbf{p}'}(G) = \mathbb{P}_{\mathbf{p}}(G) = \frac{1}{2}$. In this case, we can therefore conclude by picking any non FA-diagnosable POpLTS $\mathcal{N}_{\mathbf{p}'}$.

Assuming $G \subsetneq F$, since G is closed, there is some cylinder $C(\rho)$ with $\rho = q_0 u q_{\varepsilon} \dots q_w$ such that $G \cap C(\rho) = \emptyset$. Then we define the POpLTS $\mathcal{N}_{\mathbf{p}'}$ as the POpLTS $\mathcal{N}_{\mathbf{p}}$ except that for every $w \leq w'$ and every $x \in \{a, b\}$, $\mathbf{p}'(w', x) = \frac{1}{2}$. Thus for every $n \geq |w|$, $\sum_{|w'|=n} \mathbf{p}'(w', b) \geq \frac{\mathbf{p}_{\mathbf{p}}(\rho)}{2}$. So $\mathcal{N}_{\mathbf{p}'}$ is not FA-diagnosable. On the other hand, $\mathbb{P}_{\mathbf{p}'}(E \cap F) \geq \mathbb{P}_{\mathbf{p}'}(G) = \mathbb{P}_{\mathbf{p}}(G) > 0$. ◀

B

 Details and proofs for Section 4

B.1 Formal definitions

Here we give formal definitions omitted in the core of the paper due to space constraints. More precisely given a POpVPA \mathcal{V} , we define its estimate VPA $\mathcal{A}(\mathcal{V})$, its enlarged VPA $\widehat{\mathcal{V}}$ and their synchronised product.

Let $\mu \in \{g, c, f\}$ we write $(q, \gamma) \xrightarrow{o}_\mu (q', w)$ with $o \in \Sigma_o$ if when $\mu = g$ (resp. c, f), there exists a general (resp. correct, faulty) run of transitions starting from (q, γ) to (q', w) such that all transitions are unobservable except the last one labelled by e with $\mathcal{P}(e) = o$. Let ρ be such a run then we also write $(q, \gamma) \xrightarrow{\rho}_\mu (q', w)$. All transitions of such runs are local except the last one whose type depends on the type of o .

► **Definition 16.** Given $(\mathcal{V}, \mathcal{P}, \Sigma_o)$ a POpVPA with $\mathcal{V} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$, its *estimate VPA* is the deterministic VPA $\mathcal{A}(\mathcal{V}) = (Q^e, \Sigma_o, \Gamma^e, \delta^e)$ defined by:

- $Q^e = \{run\} \uplus (2^{\Gamma \times (\mathbf{Tg} \times Q)^2} \setminus \emptyset)$ is the set of states with initial state $q_0^e = run$;
- $\Gamma^e = 2^{(\Gamma \times \mathbf{Tg} \times Q)^2} \setminus \emptyset$ is the stack alphabet with set of bottom stack symbols $\Gamma_\perp^e = 2^{Init} \setminus \emptyset$ where $Init = \{ \frac{\perp_0, X, q}{\perp_0, U, q_0} \mid (X, q) \in \mathbf{Tg} \times Q \}$ and initial stack symbol $\perp_0^e = \{ \frac{q_0, U, \perp_0}{q_0, U, \perp_0} \}$;
- The transition relation δ^e is defined as follows.

local transitions $(run, bel, o, run, bel') \in \delta^e$ if:

- $\frac{\beta, U, r}{\alpha^-, U, q^-} \in bel'$ iff there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_c (r, \beta)$.
- If W occurs in bel , $\frac{\beta, W, r}{\alpha^-, X, q^-} \in bel'$ iff there exists $\frac{\alpha, W, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta)$.
- If W occurs in bel , $\frac{\beta, V, r}{\alpha^-, X, q^-} \in bel'$ iff
 - (1) there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \beta)$ or
 - (2) there exists $\frac{\alpha, V, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta)$.
- If W does not occur in bel , $\frac{\beta, W, r}{\alpha^-, X, q^-} \in bel'$ iff
 - (1) there exists $\frac{\alpha, U, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \beta)$ or
 - (2) there exists $\frac{\alpha, V, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta)$.

push transitions $(run, bel, o, run, bel' bel'') \in \delta^e$ if:

- $\frac{\beta^-, U, r}{\alpha^-, U, q^-} \in bel'$ and $\frac{\beta, U, r}{\beta^-, U, r} \in bel''$ iff there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_c (r, \beta^- \beta)$.
- If W occurs in bel , $\frac{\beta^-, W, r}{\alpha^-, X, q^-} \in bel'$ and $\frac{\beta, W, r}{\beta^-, W, r} \in bel''$ iff there exists $\frac{\alpha, W, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta^- \beta)$.
- If W occurs in bel , $\frac{\beta^-, V, r}{\alpha^-, X, q^-} \in bel'$ and $\frac{\beta, V, r}{\beta^-, V, r} \in bel''$ iff
 - (1) there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \beta^- \beta)$ or
 - (2) there exists $\frac{\alpha, V, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta^- \beta)$.
- If W does not occur in bel , $\frac{\beta^-, W, r}{\alpha^-, X, q^-} \in bel'$ and $\frac{\beta, W, r}{\beta^-, W, r} \in bel''$ iff
 - (1) there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \beta^- \beta)$ or
 - (2) there exists $\frac{\alpha, V, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta^- \beta)$.

pop transitions $(run, bel, o, \ell, \varepsilon) \in \delta^e$ with $\ell \in Q^e \setminus \{run\}$ if:

- $\frac{U, r}{\alpha^-, U, q^-} \in \ell$ iff $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_c (r, \varepsilon)$.
- If W occurs in bel , $\frac{W, r}{\alpha^-, X, q^-} \in \ell$ iff there exists $\frac{\alpha, W, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \varepsilon)$.
- If W occurs in bel , $\frac{V, r}{\alpha^-, X, q^-} \in \ell$ iff
 - (1) there exists $\frac{\alpha, U, q}{\alpha^-, U, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \varepsilon)$ or
 - (2) there exists $\frac{\alpha, V, q}{\alpha^-, X, q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \varepsilon)$.

- If W does not occur in bel , $\frac{W,r}{\alpha^-,X^-,q^-} \in \ell$ iff
 - (1) there exists $\frac{\alpha,U,q}{\alpha^-,U,q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_f (r, \varepsilon)$ or
 - (2) there exists $\frac{\alpha,V,q}{\alpha^-,X^-,q^-} \in bel$ and $(q, \alpha) \xrightarrow{o}_g (r, \beta^- \beta)$.

ε -transitions $(\ell, bel, \varepsilon, run, bel') \in \delta^\varepsilon$ if:

$$\frac{\alpha,X',r}{\alpha^-,X^-,q^-} \in bel' \text{ iff there exists } \frac{\alpha,X,q}{\alpha^-,X^-,q^-} \in bel \text{ and } \frac{X',r}{\alpha^-,X^-,q^-} \in \ell.$$

While $\mathcal{A}(\mathcal{V})$ contains ε -transitions it is deterministic: from any configuration, either a single ε -transition is enabled or for all event o , there is at most one o -transition enabled. We say that a configuration is *stable* if its associated state is *run*.

Illustration. Let us look at the run given in the example of Figure 5. It starts in the initial configuration $(run, \left\{ \frac{\perp_0, U, q_0}{\perp_0, U, q_0} \right\})$ which represents the empty run.

From q_0 there exists only one path of observation *in* the POpVPA. As this path is correct, by reading *in* on the estimate VPA we reach $(run, \left\{ \frac{\gamma, U, q_0}{\perp_0, U, q_0} \right\})$. The new element of the stack

$\left\{ \frac{\gamma, U, q_0}{\perp_0, U, q_0} \right\}$ signifies that the real stack has head γ and is in q_0 after a correct run, moreover the run entered q_0 when it pushed this γ and it does not have a second non-terminal element

in our stack. Reading a second *in* is still doable by a single run, we reach $(run, \left\{ \frac{\gamma, U, q_0}{\gamma, U, q_0} \right\})$

which modifies one information compared to before: we know from the bottom part of the head stack that the stack has at least a second γ .

Reading a *serve* then is possible as there exists a correct signalling run from q_0 to q_1 with only observable *serve*. The estimate VPA modifies the head stack so as to represent that the run we follow is now in q_1 but without modifying anything else.

Reading a *pop* event raises a complication: from q_1 with head of stack γ , reading a *pop* can be done by a correct run staying in q_1 or by a faulty run going in f_1 . To represent this and the popping of the stack, we go in two steps. In the first step, we go to the state $\left\{ \frac{U, q_1}{\gamma, U, q_0}, \frac{W, f_1}{\gamma, U, q_0} \right\}$ which keeps the information of the two possibilities of current configuration and we pop the stack. In the second step, we deterministically read an ε transition that transfer this information from the state to the stack. In order to transfer the information, the estimate VPA checks which of the current possible runs (represented by $\frac{U, q_1}{\gamma, U, q_0}$ and $\frac{W, f_1}{\gamma, U, q_0}$) corresponds to each of the new head of stack. This is done by comparing the bottom part of the run with the top part of the head of stack, here γ, U, q_0 in every cases. Reading a second *pop* realises a similar process reaching $(run, \left\{ \frac{\perp_0, U, q_1}{\perp_0, U, q_0}, \frac{\perp_0, W, f_1}{\perp_0, U, q_0} \right\})$. An *empty* would lead to $(run, \left\{ \frac{\perp_0, U, q_0}{\perp_0, U, q_0} \right\})$ as there is a correct run from q_1 to q_0 labelled by *empty* but no run from f_1 with such label. Conversely a *reset* can not be read from q_1 but it can be read from f_1 , thus we reach $(run, \left\{ \frac{\perp_0, W, q_0}{\perp_0, U, q_0} \right\})$.

The estimate VPA manages information in order to evaluate ν_u and ν_w . In order to evaluate ν_f , the enlarged POpVPA keeps within its states the status (correct/faulty) of the run.

► **Definition 17.** Let $\mathcal{V} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ be a pVPA. Then the pVPA $\widehat{\mathcal{V}} = (\widehat{Q}, \Sigma, \Gamma, \widehat{\delta}, \widehat{\mathbf{P}})$ is defined by:

- $\widehat{Q} = Q_c \uplus Q_f$ where $Q_c = \{q_c \mid q \in Q\}$ and $Q_f = \{q_f \mid q \in Q\}$ with initial state $q_{0,c}$;
- For all $(q, \gamma, a, q', w) \in \delta$ with $a \neq \mathbf{f}$ and all $g \in \{c, f\}$, $(q_g, \gamma, a, q'_g, w) \in \widehat{\delta}$;
- For all $(q, \gamma, \mathbf{f}, q', w) \in \delta$ and all $g \in \{c, f\}$, $(q_g, \gamma, \mathbf{f}, q'_g, w) \in \widehat{\delta}$;
- For all $(q_g, \gamma, a, q'_g, w) \in \widehat{\delta}$, $\widehat{\mathbf{P}}(q_g, \gamma, a, q'_g, w) = \mathbf{P}(q, \gamma, a, q', w)$.

We now define the product $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ between $\widehat{\mathcal{V}}$ and $\mathcal{A}(\mathcal{V})$ that keeps all the information we need along a run.

► **Definition 18.** Given $(\mathcal{V}, \mathcal{P}, \Sigma_o)$ a POpVPA with $\widehat{\mathcal{V}} = (\widehat{Q}, \Sigma, \Gamma, \widehat{\delta}, \widehat{\mathbf{P}})$ and $\mathcal{A}(\mathcal{V}) = (Q^e, run, \Sigma_o, \Gamma^e, \delta^e)$, their synchronised product is the pVPA $\mathcal{V}_{\mathcal{A}(\mathcal{V})} = (Q^A, \Sigma \cup \{\Theta\}, \Gamma^A, \delta^A, \mathbf{P}^A)$ where:

- $Q^A = \widehat{Q} \times Q^e$ is the set of control states with initial state $q_0^A = (q_{0,c}, run)$;
- $\Gamma^A = \Gamma \times \Gamma^e$ is the stack alphabet with $\Gamma_{\perp} \times \Gamma_{\perp}^e$ the set of bottom stack symbols and $\perp_0^A = (\perp_0, \frac{\perp_0, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0})$ the initial symbol;
- The transition relation δ^A consists of:

local transitions.

- For all $(q, \gamma, a, q', \gamma') \in \widehat{\delta}$ with a unobservable and $bel \in \Gamma^e$,
 $((q, run), (\gamma, bel), a, (q', run), (\gamma', bel)) \in \delta^A$;
- For all $(q, \gamma, a, q', \gamma') \in \widehat{\delta}$ and $(run, bel, o, run, bel') \in \delta^e$ with $\mathcal{P}(a) = o$;
 $((q, run), (\gamma, bel), a, (q', run), (\gamma', bel')) \in \delta^A$;
- For all $(\ell, bel, \varepsilon, run, bel') \in \delta^e$, $q \in \widehat{Q}$ and $\gamma \in \Gamma$,
 $((q, \ell), (\gamma, bel), \Theta, (q, run), (\gamma, bel')) \in \delta^A$;

push transitions.

- For all $(q, \gamma, a, q', \gamma' \gamma'') \in \widehat{\delta}$ and $(run, bel, o, run, bel' bel'') \in \delta^e$ with $\mathcal{P}(a) = o$;
 $((q, run), (\gamma, bel), a, (q', run), (\gamma', bel')(\gamma'', bel'')) \in \delta^A$;

pop transitions.

- For all $(q, \gamma, a, q', \varepsilon) \in \widehat{\delta}$ and $(run, bel, o, \ell, \varepsilon) \in \delta^e$ with $\mathcal{P}(a) = o$;
 $((q, run), (\gamma, bel), a, (q', \ell), \varepsilon) \in \delta^A$;

- The transition probability function \mathbf{P}^A is defined by:

- $\mathbf{P}^A((q, run), (\gamma, bel), a, (q', run), (\gamma', bel')) = \widehat{\mathbf{P}}(q, \gamma, a, q', \gamma')$;
- $\mathbf{P}^A((q, run), (\gamma, bel), a, (q', run), (\gamma', bel')(\gamma'', bel'')) = \widehat{\mathbf{P}}(q, \gamma, a, q', \gamma' \gamma'')$;
- $\mathbf{P}^A((q, run), (\gamma, bel), a, (q', \ell), \varepsilon) = \widehat{\mathbf{P}}(q, \gamma, a, q', \varepsilon)$;
- for $\ell \in Q^e \setminus \{run\}$, $\mathbf{P}^A((q, \ell), (\gamma, bel), \Theta, (q, run), (\gamma, bel')) = 1$.

Illustration. The product POpVPA contains the current run of the POpVPA, information on the correctness of the run and the information given by the estimate POpVPA. If we look at the faulty run given in the example of Figure 4, after reading *in*, we are in state $(q_{0,c}, run)$ meaning our real state is q_0 , it was reached by a correct run and our estimate

VPA is in state *run*, the head of stack is $(\gamma, \left[\begin{array}{c} \{\frac{\gamma, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0}\} \\ \{\frac{\perp_0, \mathbf{U}, q_0}{\perp_0, \mathbf{U}, q_0}\} \end{array} \right])$, meaning our real head is γ and the

rest is the head of the estimate VPA. If we follow the faulty run until after the first pop, we reach the state $(f_{1,f}, \{\frac{\mathbf{U}, q_1}{\gamma, \mathbf{U}, q_0}, \frac{\mathbf{W}, f_1}{\gamma, \mathbf{U}, q_0}\})$, we are thus in f_1 with a faulty run and the estimate VPA is in one of the temporary states. In order to leave this state, we read a Θ which leads to the state $(f_{1,f}, run)$. Θ is an event affecting only the part of the POpVPA corresponding to the estimate VPA, making it realises the ε transition.

Given a finite run ρ of \mathcal{V} , we inductively define the run $\bar{\rho}$ of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ as follows. First $(q_0, \perp_0) = (q_0^A, \perp_0^A)$. Let ρ of length $n \geq 1$, $a \in \Sigma$ and $q \in Q$ and $\gamma_1, \dots, \gamma_h \in \Gamma$ such that $\rho = \rho' a(q, \gamma_1 \dots \gamma_h)$. If $a \notin \Sigma_b$ then $\bar{\rho} = \bar{\rho}' a((q_g, run), (\gamma_1, bel_1) \dots (\gamma_h, bel_h))$ where $g = c$ iff ρ is correct and $(run, bel_1 \dots bel_h)$ is the configuration reached by $\mathcal{P}(\rho)$ in $\mathcal{A}(\mathcal{V})$. If $a \in \Sigma_b$ then $\bar{\rho} = \bar{\rho}' a((q_g, \ell), (\gamma_1, bel_1) \dots (\gamma_h, bel_h)) \Theta ((q_g, run), (\gamma_1, bel_1) \dots (\gamma_{h-1}, bel_{h-1})(\gamma_h, bel'_h))$ where $g = c$ iff ρ is correct, $(\ell, bel_1 \dots bel_h)$ is the configuration reached by $\mathcal{P}(\rho)$ in $\mathcal{A}(\mathcal{V})$ and $(run, bel_1 \dots bel_{h-1} bel'_h)$ is the single next configuration reached by an ε transition. As previously observed, $\mathbb{P}(\rho) = \mathbb{P}(\bar{\rho})$.

B.2 Decidability of diagnosability for POpVPA

In order to prove decidability of diagnosability for a POpVPA \mathcal{V} , one wants to check whether the formulae characterising diagnosability hold on \mathcal{V} . To do so, we transform the **pathL** formulae of Section 3 into pLTL properties that are checked on $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$. These **pathL** formulae use three paths formulae \mathfrak{f} , \mathfrak{U} and \mathfrak{W} . In the core of the paper, we explained how to define alternative pLTL formulae, relying on atomic propositions ν_f , ν_u and ν_w that only depend on the current control state and top of stack symbol of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$. Proposition 19 links runs of \mathcal{V} and observed sequences of $\mathcal{A}(\mathcal{V})$ and Proposition 13 establishes the correctness of the ν 's with respect to the paths formulae \mathfrak{f} , \mathfrak{U} and \mathfrak{W} .

► **Proposition 19.** *Let σ be an observed sequence of $\mathcal{A}(\mathcal{V})$ and ρ^* be its corresponding finite run with successive stable configurations $(run, w_0) \dots (run, w_n)$. Let $w_n = bel_1 \dots bel_h$ and for $i < n$, $bel^{(i)}$ be the top stack symbol of w_i . Then:*

- For all $\frac{\gamma_h, X_h, q_h}{\gamma_{h-1}, X_{h-1}, q_{h-1}} \in bel_h$, there exists a sequence $(\frac{\gamma_i, X_i, q_i}{\gamma_{i-1}, X_{i-1}, q_{i-1}})_{0 < i < h}$ such that for all i , $\frac{\gamma_i, X_i, q_i}{\gamma_{i-1}, X_{i-1}, q_{i-1}} \in bel_i$ and a signalling run ρ of \mathcal{V} such that $\mathcal{P}(\rho) = \sigma$ that reaches configuration $(q_h, \gamma_1 \dots \gamma_h)$. In addition:
 - if $X_h = U$ then ρ may be chosen correct;
 - if $X_h \neq U$ then ρ may be chosen faulty;
 - if $X_h = W$ then there exists $0 < k \leq n$, such that $\rho_{\downarrow k}$ is faulty and W does not occur in $bel^{(k-1)}$.
- Conversely, let ρ be a signalling run of \mathcal{V} such that $\mathcal{P}(\rho) = \sigma$ reaching configuration $(q_h, \gamma_1 \dots \gamma_h)$, there exists a sequence $(\frac{\gamma_i, X_i, q_i}{\gamma_{i-1}, X_{i-1}, q_{i-1}})_{0 < i \leq h}$ such that for all i , $\frac{\gamma_i, X_i, q_i}{\gamma_{i-1}, X_{i-1}, q_{i-1}} \in bel_i$. In addition:
 - if ρ is correct then $X_h = U$;
 - if ρ is faulty then $X_h \neq U$;
 - if there exists $0 < k \leq n$, such that $\rho_{\downarrow k}$ is faulty and W does not occur in $bel^{(k-1)}$ then $X_h = W$.

Proof. We prove it by induction on $|\sigma|$. The basis case is straightforward. For the inductive step, we only detail the most involved case: $\sigma[n] \in \Sigma_{o,b}$. For the properties related to tags, we only detail the ones related to W . Denote $\sigma' = \sigma[1] \dots \sigma[n-1]$ and $w_{n-1} = bel'_1 \dots bel'_h bel'_{h+1}$.

- Let $\frac{\gamma_h, X_h, q_h}{\gamma_{h-1}, X_{h-1}, q_{h-1}} \in bel_h$. By construction, there exists $\frac{\gamma'_{h+1}, X'_{h+1}, q'_{h+1}}{\gamma'_h, X'_h, q'_h} \in bel'_{h+1}$ with $\gamma'_h = \gamma_h$, a signalling run $(q'_{h+1}, \gamma'_{h+1}) \xrightarrow{\rho''} (q_h, \varepsilon)$ with $proj(\rho'') = \sigma[n]$, $\frac{\gamma'_h, X'_h, q'_h}{\gamma'_{h-1}, X'_{h-1}, q'_{h-1}} \in bel'_h$ where $(\gamma'_{h-1}, X'_{h-1}, q'_{h-1}) = (\gamma_{h-1}, X_{h-1}, q_{h-1})$ and X_h is obtained by updating X'_{h+1} w.r.t. bel'_{h+1} and ρ'' . In particular if $X_h = W$ then

- (1) $X'_{h+1} = W$, or
- (2) W does not occurs in bel'_{h+1} and (a) $X'_{h+1} = V$ or (b) $X'_{h+1} = U$ and ρ'' is faulty.

By inductive hypothesis, there exists a sequence $(\frac{\gamma'_i, X'_i, q'_i}{\gamma'_{i-1}, X'_{i-1}, q'_{i-1}})_{0 < i \leq h}$ such that for all i , $\frac{\gamma'_i, X'_i, q'_i}{\gamma'_{i-1}, X'_{i-1}, q'_{i-1}} \in bel'_i$ and a signalling run ρ' of \mathcal{V} such that $\mathcal{P}(\rho') = \sigma'$ reaching configuration $(q'_{h+1}, \gamma'_1 \dots \gamma'_{h+1})$. Consider the signalling run $\rho = \rho' \rho''$; it reaches configuration $(q_h, \gamma_1 \dots \gamma_h)$. Since for all $i < h$, $bel'_i = bel_i$, the sequence $(\frac{\gamma'_i, X'_i, q'_i}{\gamma'_{i-1}, X'_{i-1}, q'_{i-1}})_{0 < i < h}$ and the run ρ are appropriate. The three additional properties follow from the rules of tag updates.

In particular, if $X_h = W$, then

- the assertion (1) holds and then the property comes from the inductive hypothesis, or
- the assertion (2) holds which implies that W does not occur in bel'_{h+1} and ρ is faulty.
- Let ρ be a signalling run of \mathcal{V} such that $\mathcal{P}(\rho) = \sigma$ which reaches configuration $(q_h, \gamma_1 \dots \gamma_h)$. Let us write $\rho = \rho_{\downarrow n-1} \rho''$ with $(q'_{h+1}, \gamma'_{h+1}) \xrightarrow{\rho''} (q_h, \varepsilon)$. By the inductive hypothesis, there

exists a sequence $(\frac{\gamma'_i, X'_i, q'_i}{\gamma'_{i-1}, X'_{i-1}, q'_{i-1}})_{0 < i \leq h+1}$ such that for all i , $\frac{\gamma'_i, X'_i, q'_i}{\gamma'_{i-1}, X'_{i-1}, q'_{i-1}} \in bel'_i$ and for all $i \leq h$, $\gamma'_i = \gamma_i$. By construction, $\frac{\gamma_h, X_h, q_h}{\gamma'_{h-1}, X'_{h-1}, q'_{h-1}} \in bel_h$ for some X_h . Since $bel_i = bel'_i$ for all $i < h$, we obtain the required sequence of items.

The three additional properties follow from the rules of tag updates. In particular, assume there exists $0 < k \leq n$, such that $\rho_{\downarrow k}$ is faulty and W does not occur in bel^{k-1} .

◦ If $\rho_{\downarrow n-1}$ is correct then, as ρ is faulty, ρ'' is faulty and W does not occur in $bel^{n-1} = bel'_{h+1}$. So by construction $X_h = W$.

◦ If $\rho_{\downarrow n-1}$ is faulty then

- either $X'_{h+1} = W$ and by construction $X_h = W$,
- or $X'_{h+1} = V$. By induction hypothesis there does not exist $0 < k \leq n-1$, such that $\rho_{\downarrow k}$ is faulty and W does not occur in bel^{k-1} . So W does not occur in $bel^{n-1} = bel'_{h+1}$. Therefore $X_h = W$.

◀

► **Proposition 13.** *Let ρ be an infinite run of \mathcal{V} . Then:*

- For all $k \in \mathbb{N}$, $f(\rho_{\downarrow k}) \Leftrightarrow \nu_f(\text{last}(\bar{\rho}_{\downarrow k}))$ and $\mathfrak{U}(\rho_{\downarrow k}) \Leftrightarrow \nu_u(\text{last}(\bar{\rho}_{\downarrow k}))$;
- $\rho \models \diamond \square \mathfrak{W} \Leftrightarrow \exists K \forall k \geq K. \nu_w(\text{last}(\bar{\rho}_{\downarrow k})) = \text{true}$.

Proof. First, remark that f and ν_f obviously coincide: they both express that a fault occurred.

To prove the second item, about \mathfrak{U} and ν_u , we use the link from between observed sequences and the tag U in $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$. Let σ be an observed sequence triggered by a run of \mathcal{V} . Then bel_σ is the top stack symbol of the stable configuration in $\mathcal{A}(\mathcal{V})$ reached by the run accepting σ (so ending by an ε -transition if the last event is a pop event). Due to Proposition 19, U occurs in bel_σ iff there is a correct signalling run of \mathcal{V} with observed sequence σ . According to the definition of ν_u , we thus deduce that for any finite signalling run ρ of \mathcal{V} , $\nu_u(\text{last}(\rho)) = \text{true}$ iff $\mathfrak{U}(\rho) = \text{true}$.

We now establish the link between \mathfrak{W} and ν_w . To show the left-to-right implication, let $\rho \in \Omega$ and $K_0 \in \mathbb{N}$ be such that $\rho, K_0 \models \square \mathfrak{W}$. By definition of \mathfrak{W} , $\text{firstf}(\mathcal{P}(\rho_{\downarrow k}))$ is constant and bounded by K_0 for $k \geq K_0$. For all $k \in \mathbb{N}$, let bel_k be the top stack symbol reached in $\mathcal{A}(\mathcal{V})$ after reading the observed sequence $\mathcal{P}(\rho_{\downarrow k})$. If for all $k \geq K_0$, W occurs in bel_k , then for all $k \geq K_0$, $\nu_w(\text{last}(\bar{\rho}_{\downarrow k})) = \text{true}$. Otherwise there exists $K_1 \geq K_0$ such that W does not occur in bel_{K_1} . Let $k > K_1$, as $\text{firstf}(\mathcal{P}(\rho_{\downarrow k})) \leq K_0$, there exists a faulty run ρ' of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ such that $\mathcal{P}(\rho') = \mathcal{P}(\bar{\rho}_{\downarrow n})$ and $\rho'_{\downarrow K_0}$ is faulty. W does not occur in bel_{K_1} and $\rho'_{\downarrow K_1+1}$ is faulty. Thus by Proposition 19, W occurs in bel_k . Therefore for all $n > K_1$, $\nu_w(\text{last}(\bar{\rho}_{\downarrow n})) = \text{true}$.

Let us show the right-to-left implication. Let $\rho \in \Omega$ and $K \in \mathbb{N}$ be such that for all $k \geq K$, $\nu_w(\text{last}(\bar{\rho}_{\downarrow k})) = \text{true}$. By definition of ν_w for all $k \geq K$, W occurs in bel_k (defined as above). Let $k \geq K$, by Proposition 19, there exists a run ρ' of $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ such that $\mathcal{P}(\rho') = \mathcal{P}(\bar{\rho}_{\downarrow k})$ and there exists $n \leq k$ such that $\rho'_{\downarrow n}$ is faulty and W does not occur in bel_{n-1} . Thus $n \leq K$. Therefore for all $k \geq K$, $\text{firstf}(\mathcal{P}(\rho_{\downarrow k})) \leq K$. Since beyond K firstf is bounded, it is non decreasing and then eventually constant. Let K' such that for all $k \geq K'$, $\text{firstf}(\mathcal{P}(\rho_{\downarrow k})) = \text{firstf}(\mathcal{P}(\rho_{\downarrow k-1}))$. So $\rho, K' \models \square \mathfrak{W}$ and thus $\rho \models \diamond \square \mathfrak{W}$. ◀

We extend ν_f , ν_u and ν_w over configurations $cf = ((q, \ell), w)$ with $\ell \neq \text{run}$ by $\nu_f(cf) = \nu_u(cf) = \nu_w(cf) = \text{true}$.

► **Theorem 14.** *FF-diagnosability, IF-diagnosability and IA-diagnosability are decidable in EXPSPACE for POpVPA.*

Proof. The above lemmas allows us to derive pLTL characterisations of diagnosability for POpVPA. Namely, for \mathcal{V} a POpVPA, as \mathcal{V} and $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ have the same probabilistic behaviour,

- \mathcal{V} is FF-diagnosable iff $\mathcal{V}_{\mathcal{A}(\mathcal{V})} \models \mathbb{P}^=0(\diamond \square (\nu_f \wedge \nu_u))$;
- \mathcal{V} is IA-diagnosable iff $\mathcal{V}_{\mathcal{A}(\mathcal{V})} \models \mathbb{P}^=0(\diamond \square (\nu_u \wedge \nu_w))$.

Moreover, since the POpLTS generated by POpPDA are finitely-branching, IF-diagnosability coincides with FF-diagnosability [3] (See also 4). The two above qualitative pLTL formulae can be checked on general probabilistic pushdown automata (beyond visibly pushdown ones) thanks to [6]. More precisely, one can transform $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ into a recursive Markov chain (the transformation is linear) [7]. Then, the model checking of qualitative pLTL on recursive Markov chains is doable in PSPACE in the size of the Recursive Markov Chain and EXPTIME in the size of the formulae [6]. In our case, the product VPA $\mathcal{V}_{\mathcal{A}(\mathcal{V})}$ is exponential in the size of \mathcal{V} and the size of the formulae is constant. This yields an EXPSPACE algorithm for checking diagnosability of POpVPA. ◀

B.3 EXPTIME-hardness of the diagnosability for POpVPA

We prove here Theorem 15, stating the EXPTIME-hardness of diagnosability for POpVPA. Let us restate it below more precisely.

► **Theorem 20.** *FF-diagnosability, FA-diagnosability and IA-diagnosability are EXPTIME-hard for POpVPA.*


Proof. Let us start with FF-diagnosability. The proof is by reduction from the universality problem for VPA, which is known to be EXPTIME-hard [1].

From a VPA $\mathcal{V} = (Q, \Sigma, \Gamma, \delta)$ and a subset of accepting control states $Q_f \subseteq Q$, we build a pVPA $\mathcal{V}' = (Q', \Sigma', \Gamma', \delta', \mathbf{P}')$ as follows:

- $Q' = Q \cup \{f_0, f_b, q'_0, q_b\}$ and q'_0 is the initial state;
- $\Sigma' = \Sigma \uplus \{\mathbf{f}, u, b, \natural\}$;
- $\Gamma' = \Gamma \uplus \{B\}$ and $\Gamma'_\perp = \Gamma_\perp$;
- Writing δ_\natural , resp. $\delta_\#$ and δ_b for the set of local resp. push and pop transitions of \mathcal{V} , δ' consists of the following transitions:
 - local** $\delta_\natural \cup \{(q'_0, \perp_0, u, \perp_0, q_0), (q'_0, \perp_0, \mathbf{f}, \perp_0, f_0), (f_0, \gamma, \natural, \gamma, f_b) \mid \gamma \in \Gamma \cup \{\perp_0\}\}$
 $\cup \{(q, \gamma, \natural, \gamma, q_b) \mid q \in Q_f, \gamma \in \Gamma \cup \{\perp_0\}\}$
 $\cup \{(f_0, \gamma, a, \gamma, f_0) \mid a \in \Sigma_\natural, \gamma \in \{B, \perp_0\}\}$
 $\cup \{(q_b, \perp_0, \natural, \perp_0, q_0), (f_b, \perp_0, \natural, \perp_0, f_0)\}$;
 - push** $\delta_\# \cup \{(f_0, \gamma, a, \gamma B, f_0) \mid a \in \Sigma_\#, \gamma \in \{B, \perp_0\}\}$;
 - pop** $\delta_b \cup \{(f_0, B, a, \varepsilon, f_0) \mid a \in \Sigma_b\} \cup \{(f_b, B, b, \varepsilon, f_b)\} \cup \{(q_b, \gamma, b, \varepsilon, q_b) \mid \gamma \in \Gamma\}$;
- \mathbf{P}' is such that for every $\gamma \in \Gamma$, $\mathbf{P}'(f_0, \gamma, \natural, \gamma, f_b) = \frac{1}{2}$, and assigns arbitrary positive probabilities to the other transitions in δ' .

We further consider the POpVPA $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ with $\Sigma_o = \Sigma \cup \{b, \natural\}$ and the masking function satisfies $\mathcal{P}(u) = \mathcal{P}(\mathbf{f}) = \varepsilon$ and $\mathcal{P}(x) = x$ for any other event $x \in \Sigma'$. This construction is illustrated in Figure 8. The figure uses the following shortcuts: $a_b \in \Sigma_b$, $a_\natural \in \Sigma_\natural$, $a_\# \in \Sigma_\#$, $\gamma \in \Gamma$, $\gamma' \in \{B, \perp_0\}$ and $z \in \Gamma \setminus \{\perp_0\}$.

The correct observed sequences in $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ are either of the form $w_1 \natural b^{k_1} \natural w_2 \dots \natural b^{k_{n-1}} \natural w_n$ or of the form $w_1 \natural b^{k_1} \natural w_2 \dots \natural w_{n-1} \natural b^m$. In these decompositions, w_i , for $i < n$, is a sequence corresponding to a run of \mathcal{V} starting in q_0 and ending in some accepting state $q_f \in Q_f$, k_i is the number of elements in the stack after reading w_i in \mathcal{V} and also in \mathcal{V}' (apart from the bottom stack symbol \perp_0), w_n is the sequence associated to a run of \mathcal{V} starting in q_0 , and m



■ **Figure 8** A POpVPA for the EXPTIME-hardness of FF-diagnosability.

is at most the number of elements in the stack after reading w_{n-1} in \mathcal{V} . Note that k_i only depends on w_i , and does not depend on the exact run over w_i , since \mathcal{V} is a VPA.

Now, the faulty observed sequences in $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ are either of the form $w_1 \natural b^{k_1} \natural w_2 \dots \natural b^{k_{n-1}} \natural w_n$ or of the form $w_1 \natural b^{k_1} \natural w_2 \dots \natural w_{n-1} \natural b^m$. In these decompositions, $w_i \in \Sigma^*$, k_i is the size of the stack in \mathcal{V}' (apart from the bottom stack symbol \perp_0) after reading w_i and m is at most the number of elements in the stack of \mathcal{V}' after reading w_{n-1} .

Let us show that \mathcal{V} is not universal if and only if $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ is FF-diagnosable.


First assume that \mathcal{V} is not universal. Then there exists a word $w \in \Sigma^*$ such that no run of \mathcal{V} reading w ends in an accepting state q_f . However, the observed sequence of any faulty run almost-surely contains the factor $\natural w \natural$. Indeed, faulty runs almost surely visit infinitely often the configuration (f_b, \perp_0) , and from there, the probability λ to read $\natural w \natural$ is positive. Let ρ be an infinite faulty run. Its observed sequence is of the form $\mathcal{P}(\rho) = w_1 \natural b^{k_1} \natural w_2 \natural b^{k_2} \natural w_3 \dots$ with $k_i \leq |w_i|$ for every i . If there exists $i \leq n$ such that $w_i = w$ then ρ is surely faulty, since it has no corresponding correct run. The latter statement can be refined. For $n \geq |w|$, if, for every $i \leq n$, $|w_i| \leq n$ and there exists $i \leq n$ such that $w_i = w$ then ρ_{12n^2+n} is surely faulty. Indeed, $|w_i \natural b^{k_i}| \leq 2n + 1$, w occurs at the latest for $i = n$, and once it occurs the prefix is surely faulty. Let us therefore consider faulty runs that do not satisfy this property. We let $\text{Avoid}_n = \{\rho \in \mathcal{F} \mid \mathcal{P}(\rho) = w_1 \natural b^{k_1} \natural w_2 \natural b^{k_2} \natural w_3 \dots \wedge (\forall i \leq n \ w_i \neq w \vee \exists i \leq n \ |w_i| > n)\}$. By construction, $\text{FAmb}_{2n^2+n} \subseteq \text{Avoid}_n$. Moreover, using standard union-sum inequalities, $\mathbb{P}(\text{Avoid}_n) \leq (1 - \lambda)^n + \frac{n}{2^n}$ (recall that λ is the probability to read $\natural w \natural$ from (f_0, \perp_0)). Thus $\lim_{n \rightarrow \infty} \mathbb{P}(\text{Avoid}_n) = 0$ and hence $\lim_{n \rightarrow \infty} \mathbb{P}(\text{FAmb}_n) = 0$ so that $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ is FF-diagnosable. Assume now that \mathcal{V} is universal. Let ρ be an infinite surely faulty run of $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$. We write ρ' for the greatest ambiguous prefix of ρ and $a \in \Sigma_o \cup \{\natural, b\}$ such that $\rho'a$ is again a prefix of ρ . Observe that a cannot be b since the number of b 's between two \natural 's, whether on the left or right-hand-side of \mathcal{V}' , is entirely determined by the word of Σ_o^* read before the first \natural . For the same reason, if $a = \natural$, $\mathcal{P}(\rho')$ ends with a word of Σ_o^* (i.e. the number of \natural 's in $\mathcal{P}(\rho')$ is even). Let w be the greatest suffix of $\mathcal{P}(\rho')$ contained in Σ_o^* . If $a = \natural$, we deduce that there is no run starting in q_0 with observed sequence w and ending in an accepting state of \mathcal{V} . Therefore, \mathcal{V} is not universal. Similarly, if $a \in \Sigma_o$, then there is no run starting in q_0 and with observed sequence wa . In that case also, \mathcal{V} is not universal. We hence conclude that there is no infinite surely faulty run in $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$. As the probability to generate faulty runs is positive, this implies that $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ is not IF-diagnosable. Now, IF-diagnosability is equivalent to FF-diagnosability for finitely branching POpLTS (see Theorem 4), and so $\langle \mathcal{V}', \Sigma_o, \mathcal{P} \rangle$ is not FF-diagnosable.

Let us now argue for the EXPTIME-hardness of FA-diagnosability and IA-diagnosability. From the VPA $\mathcal{V} = (Q, \Sigma, \Gamma, \delta)$ and pVPA $\mathcal{V}' = (Q', \Sigma', \Gamma', \delta')$ defined above, we construct a

pVPA $\mathcal{V}'' = (Q'', \Sigma'', \Gamma'', \delta'', \mathbf{P}'')$ such that

- $Q'' = Q' \cup \{q_c\}$ and q'_0 is the initial state;
- $\Sigma'' = \Sigma \cup \{\mathbf{f}, u, \sharp, \alpha\}$;
- $\Gamma'' = \Gamma$;
- $\delta'' = \delta' \cup \{(q, \alpha, \gamma, q_c) \mid \gamma \in \Gamma \cup \{\perp_0\}, q \in Q \cup \{q_c\}\}$;
- \mathbf{P}'' assigns arbitrary positive probabilities to transitions in δ'' .

We further consider the POpVPA $\langle \mathcal{V}'', \Sigma_o, \mathcal{P} \rangle$ with $\Sigma_o = \Sigma'' \setminus \{\mathbf{f}, u\}$, and the masking function satisfies $\mathcal{P}(\mathbf{f}) = \mathcal{P}(u) = \varepsilon$ and $\mathcal{P}(x) = x$ for any other event x . The construction is illustrated in Figure 9, where we use the shortcuts: $a_b \in \Sigma_b$, $a_{\natural} \in \Sigma_{\natural}$, $a_{\sharp} \in \Sigma_{\sharp}$, $\gamma \in \Gamma$, $\gamma' \in \{B, \perp_0\}$ and $z \in \Gamma \setminus \{\perp_0\}$.



■ **Figure 9** A POpVPA for EXPTIME-hardness of FA-diagnosability and IA-diagnosability.

\mathcal{V}'' is a slight modification of \mathcal{V}' : from any state of \mathcal{V} (accepting or not), reading the new letter α leads to the sink state q_c . As a consequence, for any correct run of $\langle \mathcal{V}'', \Sigma_o, \mathcal{P} \rangle$, there is a positive probability at each step to perform event α and become surely correct. This implies $\lim_{n \rightarrow \infty} \mathbb{P}(\text{CAmb}_n)_{n \in \mathbb{N}} = 0$. Observe that the above proof for \mathcal{V}' also applies to \mathcal{V}'' : \mathcal{V} is not universal if and only if $\langle \mathcal{V}'', \Sigma_o, \mathcal{P} \rangle$ is FF-diagnosable. Now, since $\lim_{n \rightarrow \infty} \mathbb{P}(\text{CAmb}_n)_{n \in \mathbb{N}} = 0$, FF-diagnosability, FA-diagnosability and IA-diagnosability coincide for $\langle \mathcal{V}'', \Sigma_o, \mathcal{P} \rangle$. We conclude that \mathcal{V} is not universal if and only if $\langle \mathcal{V}'', \Sigma_o, \mathcal{P} \rangle$ is diagnosable (for any notion of diagnosability). ◀

B.4 Undecidability of diagnosability for POpPDA

As stated in the core of the paper, diagnosability is undecidable for partially observable probabilistic pushdown automata (POpPDA). Let us first give the definition of pPDA and POpPDA. Contrary to VPA, in PDA, the action does not determine the operation (push, pop, local) on the stack.

► **Definition 21.** A *probabilistic pushdown automaton* (pPDA) is a tuple $\mathcal{A} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ where:

- Q is a finite set of control states with q_0 the initial state;
- Σ is a finite alphabet of events;
- Γ is a finite alphabet of stack symbols including a set of bottom stack symbols Γ_{\perp} with initial symbol $\perp_0 \in \Gamma_{\perp}$;
- $\delta \subseteq Q \times \Gamma \times \Sigma \times Q \times \Gamma^*$ is the set of transitions such that for every $(q, \gamma, a, q', w) \in \delta$, $|w| \leq 2$, $\gamma \in \Gamma_{\perp}$ implies $w \in \Gamma_{\perp}(\Gamma \setminus \Gamma_{\perp})^*$ and $\gamma \notin \Gamma_{\perp}$ implies $w \in (\Gamma \setminus \Gamma_{\perp})^*$;

- \mathbf{P} is the transition probability function fulfilling for every $q \in Q$ and $\gamma \in \Gamma$:

$$\sum_{(q,\gamma,a,q',w) \in \delta} \mathbf{P}[(q,\gamma,a,q',w)] = 1.$$

► **Definition 22.** A *partially observable pPDA* (POpPDA) is a tuple $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ consisting of a pPDA \mathcal{A} equipped with a mapping $\mathcal{P} : \Sigma \rightarrow \Sigma_o \cup \{\varepsilon\}$ where Σ_o is the set of observations.

The undecidability of diagnosability for POpPDA can be derived from the undecidability of diagnosability for non-probabilistic PDA [11]. However, to show how robust the result is, we refine the statement into Theorems 23 and 24: undecidability already holds for two (incomparable) subclasses of POpPDA with restriction on what is observable and on the number of phases of any run. A *phase* is a portion of run in which the stack either never decreases or never increases.

► **Theorem 23.** *The diagnosability problems are undecidable for POpPDA even when (1) the top of the stack is not updated, (2) every event labelling a push transition is fully observable and corresponds to the pushed symbol, and (3) every run consists of at most two phases.*

Proof. The proof is by reduction from the Post correspondence problem (PCP). An instance of PCP is given by an integer $n \in \mathbb{N}$ and two families of non-empty words $\{v_i\}_{i \leq n}$ and $\{w_i\}_{i \leq n}$ on the alphabet $\{a, b\}$. The following question is undecidable: does there exist $k > 0$ and $i_1, \dots, i_k \in \{1, \dots, n\}$ such that $w_{i_1} \dots w_{i_k} = v_{i_1} \dots v_{i_k}$?

In this proof, we let ℓ_i (resp. m_i) be the length of v_i (resp. w_i). Also, given a word w and $k \leq |w|$ we use $w[k]$ to denote the k^{th} -letter of w .

From an instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ of PCP, we build a pPDA $\mathcal{A} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ as follows:

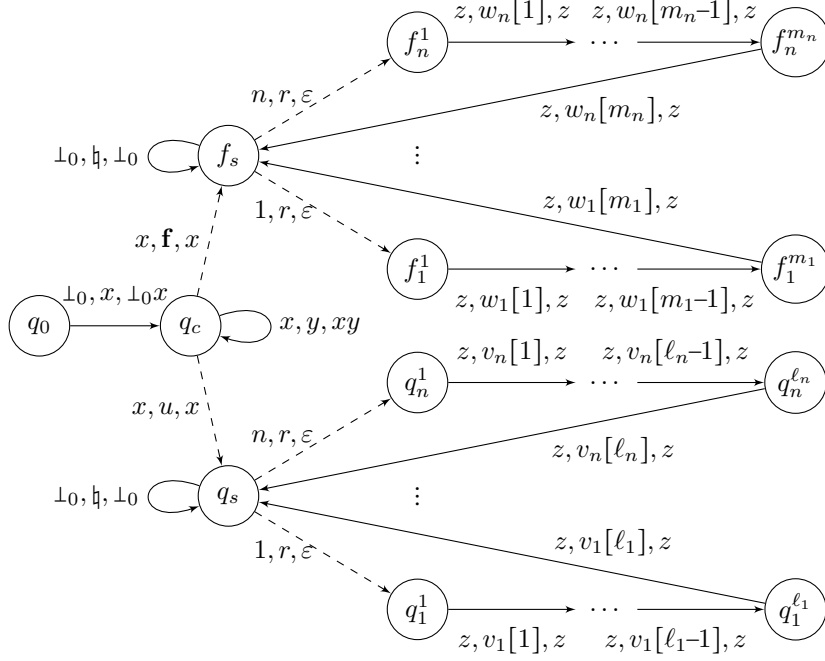
- $Q = \{q_0, q_c, q_s, f_s\} \cup \{q_i^k \mid 1 \leq i \leq n, 1 \leq k \leq \ell_i\} \cup \{f_i^k \mid 1 \leq i \leq n, 1 \leq k \leq m_i\}$;
- $\Sigma = \{1, \dots, n, \natural, u, r, \mathbf{f}, a, b\}$;
- $\Gamma = \{1, \dots, n, \perp_0\}$ with $\Gamma_\perp = \{\perp_0\}$;
- δ consists of the following transitions:
 - $\{(q_0, \perp_0, x, \perp_0 x, q_c) \mid 1 \leq x \leq n\}$
 - $\{(q_c, x, y, xy, q_c) \mid 1 \leq x, y \leq n\}$
 - $\{(q_i^k, z, v_i[k], z, q_i^{k+1}) \mid 1 \leq i \leq n, 1 \leq k < \ell_i, z \in \{\perp_0, 1, \dots, n\}\}$
 - $\{(f_i^k, z, w_i[k], z, f_i^{k+1}) \mid 1 \leq i \leq n, 1 \leq k < m_i, z \in \{\perp_0, 1, \dots, n\}\}$
 - $\{(q_i^{\ell_i}, z, v_i[\ell_i], z, q_s) \mid 1 \leq i \leq n, z \in \{\perp_0, 1, \dots, n\}\}$
 - $\{(f_i^{m_i}, z, w_i[m_i], z, f_s) \mid 1 \leq i \leq n, z \in \{\perp_0, 1, \dots, n\}\}$
 - $\{(q_s, x, r, \varepsilon, q_x^1) \mid 1 \leq x \leq n\}$
 - $\{(f_s, x, r, \varepsilon, f_x^1) \mid 1 \leq x \leq n\}$
 - $\{(q_c, x, u, x, q_s), (q_c, x, \mathbf{f}, x, f_s) \mid 1 \leq x \leq n\}$
 - $\{(q_s, \perp_0, \natural, \perp_0, q_s), (f_s, \perp_0, \natural, \perp_0, f_s)\}$.
- \mathbf{P} assigns arbitrary positive probabilities to transitions in δ :
 - $\mathbf{P}(q, \gamma, a, q', w) > 0 \Leftrightarrow (q, \gamma, a, q', w) \in \delta$ and $\sum_{(q,\gamma,a,q',w) \in \delta} \mathbf{P}[(q,\gamma,a,q',w)] = 1$.

We further consider the POpPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ with $\Sigma_o = \Sigma \setminus \{r, u, \mathbf{f}\}$, and the masking function satisfies $\mathcal{P}(u) = \mathcal{P}(r) = \mathcal{P}(\mathbf{f}) = \varepsilon$ and $\mathcal{P}(x) = x$ for any other event x . This POpPDA is represented in Figure 10.

Let us prove that the instance of the PCP is positive if and only if the POpPDA is IF-, IA- and FA-diagnosable.

Assume first that there exists a solution i_1, \dots, i_k to the PCP instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$. Consider in the POpPDA the faulty run:

$$\rho_f = q_0(i_j q_c)_{j \leq k} \mathbf{f}(f_s r(f_{i_j}^p w_{i_j}[p]))_{p \leq m_{i_j}} (f_s \natural)^{\omega},$$



■ **Figure 10** A POPDA for the proof of Theorem 23.

and the correct run:

$$\rho_c = q_0(i_j q_c)_{j \leq k} u(q_s r(q_{i_j}^p v_{i_j}[p]))_{p \leq \ell_{i_j}}(q_s \uparrow)_{j \leq k} (q_s \downarrow)^\omega.$$

These two runs have the same observed sequence: $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c) = i_1 \dots i_k w \uparrow^\omega$ with $w = w_{i_1} \dots w_{i_k} = v_{i_1} \dots v_{i_k}$. Therefore, ρ_f is an infinite ambiguous faulty run. Given that $\mathbb{P}(\rho_f) > 0$, we deduce that the POPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ is not IF-diagnosable. From Theorem 4, it is also neither IA-diagnosable nor FA-diagnosable.

Conversely, assume that the PCP instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ has no solution. Independently of that, observe that \uparrow almost surely occurs in an infinite run of the pPDA \mathcal{A} . Thus, for any $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that the measure of signalling runs with observable length N that reach configurations (q_s, \perp_0) or (f_s, \perp_0) by an event \uparrow is at least $1 - \varepsilon$. Consider a correct run ρ_c with observable length N , ending in (q_s, \perp_0) and containing at least an occurrence of \uparrow . Its observed sequence is of the form $\mathcal{P}(\rho_c) = i_1 \dots i_k v_{i_1} \dots v_{i_k} \uparrow^m$ for some i_1, \dots, i_k, m . Due to the fact that $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ has no solution, no faulty run can have the same observed sequence. Therefore, ρ_c is surely correct. Symmetrically, any faulty run ending in (f_s, \perp_0) after an occurrence of \uparrow is surely faulty. We thus conclude that, for any $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that $\mathbb{P}(\text{FAmb}_N \uplus \text{CAmb}_N) \leq \varepsilon$. As a consequence, the POPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ is FA-diagnosable. By Theorem 4 it is also IA-diagnosable and IF-diagnosable. ◀

A similar undecidability result holds for a classe of POPDA in which pop events are fully observable, and the number of phases is constant:


► **Theorem 24.** *The diagnosability problems are undecidable for POPDA even when (1) the top of the stack is not updated, (2) every event labelling a pop transition is fully observable and corresponds to the popped symbol, and (3) every run consists of at most two phases.*

Proof. The proof follows the same line as the one for Theorem 23.

From an instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ of PCP, let us define a pPDA $\mathcal{A} = (Q, \Sigma, \Gamma, \delta, \mathbf{P})$ where:

- $Q = \{q_0, q_s, f_s, q_e, f_e\} \cup \{q_i^k \mid 1 \leq i \leq n, 1 \leq k \leq \ell_i\} \cup \{f_i^k \mid 1 \leq i \leq n, 1 \leq k \leq m_i\}$;
- $\Sigma = \{1, \dots, n, \natural, u, c, \mathbf{f}, a, b\}$;
- $\Gamma = \{1, \dots, n, \perp_0\}$ with $\Gamma_\perp = \{\perp_0\}$;
- δ consists of the following transitions:
 - $\{(q_0, \perp, u, \perp, q_s), (q_0, \perp, \mathbf{f}, \perp, f_s), (q_e, \perp, \natural, \perp, q_e), (f_e, \perp, \natural, \perp, f_e)\}$
 - $\cup \{(q_i^k, z, v_i[k], z, q_i^{k+1}) \mid 1 \leq i \leq n, 1 \leq k < \ell_i, z \in \{\perp, 1, \dots, n\}\}$
 - $\cup \{(f_i^k, z, w_i[k], z, f_i^{k+1}) \mid 1 \leq i \leq n, 1 \leq k < m_i, z \in \{\perp, 1, \dots, n\}\}$
 - $\cup \{(q_i^{\ell_i}, z, v_i[\ell_i], z, q_s) \mid 1 \leq i \leq n, z \in \{\perp, 1, \dots, n\}\}$
 - $\cup \{(f_i^{m_i}, z, w_i[m_i], z, f_s) \mid 1 \leq i \leq n, z \in \{\perp, 1, \dots, n\}\}$
 - $\cup \{(q_s, z, c, zx, q_x^1) \mid z \in \{\perp, 1, \dots, n\}, x \in \{1, \dots, n\}\}$
 - $\cup \{(f_s, z, c, zx, f_x^1) \mid z \in \{\perp, 1, \dots, n\}, x \in \{1, \dots, n\}\}$
 - $\cup \{(q_s, x, x, \varepsilon, q_e) \mid x \in \{1, \dots, n\}\}$
 - $\cup \{(f_s, x, x, \varepsilon, f_e) \mid x \in \{1, \dots, n\}\}$
 - $\cup \{(q_e, x, x, \varepsilon, q_e) \mid x \in \{1, \dots, n\}\}$
 - $\cup \{(f_e, x, x, \varepsilon, f_e) \mid x \in \{1, \dots, n\}\}$.
- \mathbf{P} assigns arbitrary positive probabilities to transitions in δ .

We further consider the POpPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ with $\Sigma_o = \Sigma \setminus \{c, u, \mathbf{f}\}$, and the masking function satisfies $\mathcal{P}(u) = \mathcal{P}(c) = \mathcal{P}(\mathbf{f}) = \varepsilon$ and $\mathcal{P}(x) = x$ for any other event x . This POpPDA is represented in Figure 11.



■ **Figure 11** A POpPDA for the proof of Theorem 24.

Let us prove that the instance of the PCP is positive if and only if the POpPDA is IF-, IA- and FA-diagnosable.

Assume first that there exists a solution i_1, \dots, i_k to the PCP instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$. Consider the faulty run:

$$\rho_f = q_0 \mathbf{f} f_s (c(f_{i_j}^p w_{i_j} [p]))_{p \leq m_{i_j}} f_s)_{j \leq k} (i_j f_e)_{j \leq k} (\mathfrak{h} f_e)^\omega ,$$

and the correct run:

$$\rho_c = q_0 u q_s (c(q_{i_j}^p v_{i_j} [p]))_{p \leq \ell_{i_j}} q_s)_{j \leq k} (i_j q_e)_{j \leq k} (\mathfrak{h} q_e)^\omega .$$

These two runs have the same observed sequence: $\mathcal{P}(\rho_f) = \mathcal{P}(\rho_c) = w_{i_1} \dots i_k \mathfrak{h}^\omega$ with $w = w_{i_1} \dots w_{i_k} = v_{i_1} \dots v_{i_k}$. Therefore, ρ_f is an infinite ambiguous faulty run. Given that $\mathbb{P}(\rho_f) > 0$, we deduce that the POpPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ is not IF-diagnosable. From Theorem 4, it is also neither IA-diagnosable nor FA-diagnosable.

Conversely, assume that the PCP instance $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ has no solution.

Independently of that, observe that \mathfrak{h} almost surely occurs in an infinite run of the pPDA \mathcal{A} . Thus, for any $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that the measure of signalling runs with observable length N that reach configurations (q_e, \perp_0) or (f_e, \perp_0) by an event \mathfrak{h} is at least $1 - \varepsilon$. Consider a correct run ρ_c with observable length N ending in (q_e, \perp_0) and with an occurrence of \mathfrak{h} . Its observed sequence is of the form $v_{i_1} \dots v_{i_k} i_1 \dots i_k \mathfrak{h}^m$ for some i_1, \dots, i_k, m . Due to the fact that $(n, \{v_i\}_{i \leq n}, \{w_i\}_{i \leq n})$ has no solution, no faulty run can have the same observed sequence. Therefore, ρ_c is surely correct. Symmetrically, any faulty run ending in (f_e, \perp_0) by an occurrence of \mathfrak{h} is surely faulty. We thus conclude that, for any $\varepsilon > 0$, there exists $N \in \mathbb{N}$ such that $\mathbb{P}(\text{FAmb}_N \uplus \text{CAmb}_N) \leq \varepsilon$. As a consequence, the POpPDA $\langle \mathcal{A}, \Sigma_o, \mathcal{P} \rangle$ is FA-diagnosable. By Theorem 4 it is also IA-diagnosable and IF-diagnosable. \blacktriangleleft