

Advances on Parameters in Timed Models with Probabilities

ANR PACS Deliverable 6

Benoît Delahaye Etienne André

June 28, 2019

1 Context and objectives

In real-life applications, probabilities are often used as a building block that allows abstracting from physical constraints or unknown environments. While parameters are used in order to *reason* in a formal way on unknown (or purposefully under-specified) aspects of a system, probabilities allow for abstracting some aspects in order to “simplify” the system under study: these aspects are *approximated* using probability distributions, often based on partial information obtained through empirical analysis. In a way, one could say that parameters target aspects of the system that need to be *studied* while probabilities target aspects of the system that are inconsequential or too complex to model in a formal way. By their complementary nature, probabilistic and parametric modelling/analysis can thus be combined in several ways. In Task 6, we focus on timed probabilistic models with parameters that can range on the timing aspects of the system.

In the following, we therefore present our main contributions on this topic. Worth to notice, a postdoctoral student (Paulin Fournier) has been hired to work on this task, and has been involved in one of the contributions presented hereafter. The main publication associated to those contributions is attached at the end of the document.

2 Parametric Interval Probabilistic Timed Automata

Timed automata [AD94] are a widely recognised modelling formalism for reasoning about real-time systems. This modelling formalism, based on finite control automata equipped with clocks, which are real-valued variables which increase uniformly at the same rate, has been extended to the probabilistic framework in [GJ95, KNSS02]. In this context, discrete actions are replaced with probabilistic discrete distributions over discrete actions, allowing to model uncertainties in the system’s behaviour. This formalism has been applied to a number of case studies [KNPS06].

Unfortunately, building a system model based either on imprecise specifications or on imprecise observations often requires to fix arbitrarily a number of constants in the model, which are then calibrated by a fastidious comparison of the model behaviour and the expected behaviour. This is the case for instance for timing constants or transition probability values. In order to incorporate these uncertainties

in the model and to develop automatic calibration, more abstract formalisms have been introduced separately in the timed setting and in the probabilistic setting.

In the timed setting, *parametric timed automata* [AHV93] allow using parameter variables in the guards of timed transitions in order to account for the uncertainty on their values. Parametric probabilistic timed automata were proposed in [AFS13] to answer the following question: given a parameter valuation, what are other valuations preserving the same minimum and maximum probabilities for reachability properties as the reference valuation? Parametric probabilistic timed automata were then given a symbolic semantics in [JK14]; a method has been proposed in that same work to synthesise optimal parameter valuations to maximise or minimise the probability of reaching a discrete location.

In the pure probabilistic setting, Interval Markov Chains (IMCs for short) have been introduced [JL91] to take into account imprecision in the transition probabilities. IMCs extend Markov Chains by allowing to specify intervals of possible probabilities on transitions instead of exact values. Methods have then been developed to decide whether there exists Markov Chains with concrete probability values that match the intervals specified in a given IMC [DLL⁺12].

Contribution In this line of work, we have proposed to combine both abstraction approaches into a single specification theory: Parametric Interval Probabilistic Timed Automata (PIPTAs for short). In this setting, parameters can be used in order to abstract timed constants on transition guards while intervals can be used to abstract imprecise transition probabilities. As for IMCs, it is important to be able to decide whether the probability intervals that are specified in a model allow defining coherent probability distributions (i. e., can be matched in a real-life implementation). This is called the consistency problem. In the context of Interval Probabilistic Timed Automata with no timing parameters (IPTAs for short), we proposed an algorithm that resolves this problem. In our setting, since the behaviour of the system is conditioned by the calibration of parameter values, it is therefore necessary to decide whether there exist parameter values that ensure consistency of the resulting model (and synthesise these values when this is possible). We have shown that the existence of such parameter valuations is undecidable in the general context of PIPTAs. Still, we have proposed a semi-algorithm that synthesises, whenever it terminates, the set of parameter values that ensure consistency of the resulting IPTA.

This work has resulted in a publication in an international conference [AD16] as well as a publication in an international journal [ADF19].

References


- [AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, April 1994.
- [AD16] Étienne André and Benoît Delahaye. Consistency in parametric interval probabilistic timed automata. In *23rd International Symposium on Temporal Representation and Reasoning, TIME 2016, Kongens Lyngby, Denmark, October 17-19, 2016*, pages 110–119, 2016.
- [ADF19] Étienne André, Benoît Delahaye, and Paulin Fournier. Consistency in parametric interval probabilistic timed automata. *Journal of Logical and Algebraic Methods in Programming*, 2019.

- [AFS13] É. André, L. Fribourg, and J. Sproston. An extension of the inverse method to probabilistic timed automata. *Formal Methods in System Design*, (2):119–145, 2013.
- [AHV93] Rajeev Alur, Thomas A. Henzinger, and Moshe Y. Vardi. Parametric real-time reasoning. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, STOC’93, pages 592–601, New York, NY, USA, 1993. ACM.
- [DLL⁺12] Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wasowski. Consistency and refinement for interval markov chains. *J. Log. Algebr. Program.*, 81(3):209–226, 2012.
- [GJ95] H. Gregersen and H. E. Jensen. Formal design of reliable real time systems. Master’s thesis, Department of Mathematics and Computer Science, Aalborg University, 1995.
- [JK14] Aleksandra Jovanović and Marta Z. Kwiatkowska. Parameter synthesis for probabilistic timed automata using stochastic game abstractions. In Joël Ouaknine, Igor Potapov, and James Worrell, editors, *Proceedings of the 8th International Workshop on Reachability Problems (RP 2014)*, volume 8762 of *Lecture Notes in Computer Science*, pages 176–189. Springer, 2014.
- [JL91] B. Jonsson and K.G. Larsen. Specification and refinement of probabilistic processes. In *LICS*, pages 266–277. IEEE Computer, 1991.
- [KNPS06] Marta Z. Kwiatkowska, Gethin Norman, David Parker, and Jeremy Sproston. Performance analysis of probabilistic timed automata using digital clocks. *Formal Methods in System Design*, 29(1):33–78, 2006.
- [KNSS02] Marta Z. Kwiatkowska, Gethin Norman, R. Segala, and Jeremy Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282:101–150, 2002.



Contents lists available at ScienceDirect

Journal of Logical and Algebraic Methods in Programming

www.elsevier.com/locate/jlamp


Consistency in Parametric Interval Probabilistic Timed Automata[☆]

 Étienne André^{a,1}, Benoît Delahaye^{b,*}, Paulin Fournier^b
^a Université Paris 13, LIPN, CNRS, UMR 7030, F-93430, Villetaneuse, France^b Université de Nantes/LS2N UMR CNRS 6004, Nantes, France

ARTICLE INFO

Article history:

Received 17 July 2017

Received in revised form 8 August 2018

Accepted 30 April 2019

Available online xxxx

Keywords:

Parametric verification

Timed probabilistic systems

Parametric probabilistic timed automata

ABSTRACT

We propose a new abstract formalism for probabilistic timed systems, Parametric Interval Probabilistic Timed Automata, based on an extension of Parametric Timed Automata and Interval Markov Chains. In this context, we consider the consistency problem that amounts to deciding whether a given specification admits at least one implementation. In the context of Interval Probabilistic Timed Automata (with no timing parameters), we show that this problem is decidable and propose a constructive algorithm for its resolution. We show that the existence of timing parameter valuations ensuring consistency is undecidable in the general context, but still exhibit a syntactic condition on parameters to ensure decidability. We also propose procedures that resolve both the consistency and the consistent reachability problems when the parametric probabilistic zone graph is finite.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

Motivation. Nowadays, automata-based modeling and verification methods are mainly used in two different ways: for designing digital systems based on (mostly informal) specifications expressed by the end-users of these systems or from the knowledge designers have of their environment; and in order to abstract existing (not necessarily software) systems that are too complex to comprehend in their entirety. In both cases the complexity of the systems being designed calls for increasingly expressive abstraction artifacts such as time and probabilities. Timed automata, introduced in Alur and Dill [1], are a widely recognized modeling formalism for reasoning about real-time systems. This modeling formalism, based on finite control automata equipped with clocks, which are real-valued variables which increase uniformly at the same rate, has been extended to the probabilistic framework in Gregersen and Jensen [20], Kwiatkowska et al. [26]. In this context, discrete actions are replaced with probabilistic discrete distributions over discrete actions, allowing to model uncertainties in the system's behavior. This formalism has been applied to a number of case studies, e.g., in Kwiatkowska et al. [25].

Unfortunately, building a system model based either on imprecise specifications or on imprecise observations often requires to fix arbitrarily a number of constants in the model, which are then calibrated by a fastidious comparison of the

[☆] This work is partially supported by the ANR national research program PACS (ANR-14-CE28-0002).

^{*} Corresponding author.

 E-mail addresses: Etienne.Andre@lipn.univ-paris13.fr (É. André), Benoit.Delahaye@univ-nantes.fr (B. Delahaye), Paulin.Fournier@univ-nantes.fr (P. Fournier).

 URLs: <https://lipn.univ-paris13.fr/~andre/> (É. André), <http://pagesperso.ls2n.fr/~delahaye-b/> (B. Delahaye).

¹ This work was partially done during Étienne André's *délégation CNRS* at École Centrale de Nantes, IRCCyN, CNRS, UMR 6597, France (2015–2016).

model behavior and the expected behavior. This is the case for instance for timing constants or transition probability values. In order to incorporate these uncertainties in the model and to develop automatic calibration, more abstract formalisms have been introduced separately in the timed setting and in the probabilistic setting.

In the timed setting, *parametric timed automata* (PTAs) introduced by Alur et al. [2] allow using parameter variables in the guards of timed transitions in order to account for the uncertainty on their values. The reachability emptiness problem, i.e., the emptiness of the set of valuations for which a given discrete state is reachable, is undecidable for parametric timed automata as shown in Alur et al. [2], even for bounded parameters as shown by Miller [28], for a single integer-valued parameter as shown by Beneš et al. [11], or only when strict inequalities are used as shown by Doyen [19]. Decidable subclasses were exhibited (e.g., Hune et al. [21], Bozzelli and La Torre [13], Jovanović et al. [24], André et al. [8]).

Parametric probabilistic timed automata were proposed in André et al. [5] to answer the following question: given a timing parameter valuation, what are other valuations preserving the same minimum and maximum probabilities for reachability properties as the reference valuation? Parametric probabilistic timed automata were then given a symbolic semantics in Jovanović and Kwiatkowska [23]; a method has been proposed in that same work to synthesize optimal parameter valuations to maximize or minimize the probability of reaching a discrete location.

In the purely probabilistic setting, Interval Markov Chains (IMCs for short) have been introduced by Jonsson and Larsen [22] to take into account imprecision in the transition probabilities. IMCs extend Markov Chains by allowing to specify intervals of possible probabilities on transitions instead of exact values. Methods have then been developed to decide whether there exist Markov Chains with concrete probability values that match the intervals specified in a given IMC (see Delahaye et al. [17]).

Contribution. In this paper, we propose to combine both abstraction approaches into a single specification theory: Parametric Interval Probabilistic Timed Automata (PIPTAs for short). In this setting, parameters can be used in order to abstract timed constants on transition guards while intervals can be used to abstract imprecise transition probabilities. Allowing this higher level of freedom allow for incremental design, where one can first give large sets of values for which the system may be defined, and then further refine them. This refinement will take the form of an instance of a probabilistic interval, or the concrete instance of a timing parameter.

As for IMCs, it is important to be able to decide whether the probability intervals that are specified in a model allow defining consistent probability distributions (i.e., can be matched in a real-life implementation). This is called the consistency problem.

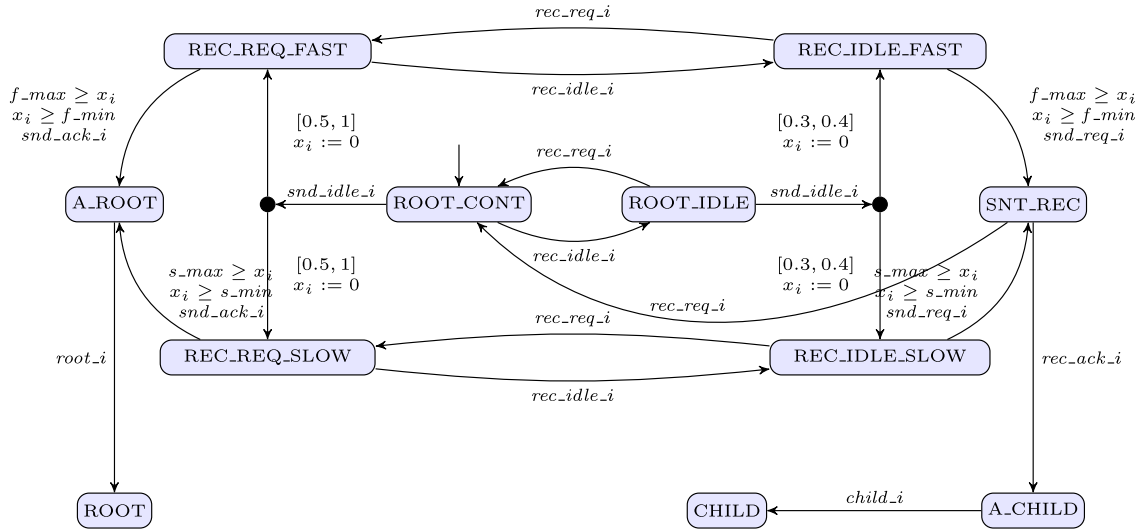
First, in the context of Interval Probabilistic Timed Automata with no timing parameters (IPTAs for short), we propose an algorithm that solves this problem.

Second, in the parametric setting, since the behavior of the system is conditioned by the calibration of parameter values, it is necessary to decide whether there exist parameter values that ensure consistency of the resulting model (and synthesize these values when this is possible). We show that the existence of such parameter valuations is undecidable in the general context of PIPTAs. Still, we exhibit a sufficient syntactic condition on the use of the parameters to ensure decidability, when parameters are partitioned into lower-bound parameters and upper-bound parameters (in their comparisons with clocks). In addition, we propose a construction that characterizes, whenever the parametric probabilistic zone graph is finite, the set of parameter values that ensure consistency of the resulting IPTA. We finally address the problem of parametric consistent reachability, i.e., of synthesizing valuations for which a given state is reachable and the model is consistent.

Example 1. The Root Contention Protocol, used for the election of a leader in the physical layer of the IEEE 1394 standard, consists in first drawing a random number (0 or 1), then waiting for some time according to the result drawn, followed by the sending of a message to the contending neighbor. This is repeated by both nodes until one of them receives a message before sending one, at which point the root is appointed. This protocol was modeled in Collomb-Annichini and Sighireanu [14] using parametric timed automata, in Kwiatkowska et al. [27] with probabilistic timed automata, and in André et al. [5] using parametric probabilistic timed automata, i.e., parametric timed automata extended with (non-parametric) probabilistic distributions.

Fig. 1 shows a PIPTA model of the node i . The wire can be found in Kwiatkowska et al. [27], André et al. [5]. Fig. 1 features one clock x_i and four parameters f_min , f_max , s_min and s_max . In short, the goal of the protocol is that each node reaches either the child status, or the root status. In addition, observe that we use probabilistic interval distributions; they can be seen as an additional design freedom, allowing for incremental design. The one going out from ROOT_IDLE clearly admits no implementation, as no instance of the two intervals $[0.3, 0.4]$ can be such that their sum is equal to 1. This probabilistic interval distribution could be either disabled by setting other probabilities to 0 so that location ROOT_IDLE becomes unreachable; or by tuning the values of the four parameters (or the parameters in the other PIPTAs in parallel) so that the guard going out from ROOT_IDLE becomes unsatisfiable. The rest of this manuscript is dedicated to this problem.

Outline. We start Section 2 with preliminary definitions and then introduce the concepts of IPTAs and PIPTAs. In Section 3, we study the consistency problem for IPTAs and propose a constructive algorithm based on the zone-graph construction that decides whether an IPTA is consistent and produces an implementation if one exists. In Section 4, we move to the general problem of consistency of PIPTAs. We first show that this problem is undecidable in general and then exhibit a

Fig. 1. PIPTA modeling node i in the Root Contention Protocol.

decidable subclass. We then propose a construction that characterizes, whenever the parametric probabilistic zone graph is finite, the set of parameter values ensuring consistency of the resulting IPTA. We also consider the problem of parametric consistent reachability. Finally, Section 5 concludes the paper.

2. Preliminaries

2.1. Clocks, parameters and constraints

Let \mathbb{N} , \mathbb{Z} , \mathbb{Q}_+ and \mathbb{R}_+ denote the sets of non-negative integers, integers, non-negative rational numbers and non-negative real numbers respectively. Given an arbitrary set S , we write $\text{Dist}(S)$ for the set of probabilistic distributions over S .

Throughout this paper, let $X = \{x_1, \dots, x_H\}$ be a set of *clocks*, i.e., real-valued variables that evolve at the same rate, and $\Gamma = \{\gamma_1, \dots, \gamma_M\}$ be a set of *parameters*, i.e., unknown constants used in guards.

A clock valuation is a function $w : X \rightarrow \mathbb{R}_+$. We identify a clock valuation w with the *point* $(w(x_1), \dots, w(x_H))$. We write $\vec{0}$ for the valuation that assigns 0 to each clock. Given $d \in \mathbb{R}_+$, $w + d$ denotes the valuation such that $(w + d)(x) = w(x) + d$, for all $x \in X$. Given $\rho \subseteq X$, we define $[w]_\rho$ as the clock valuation obtained by resetting the clocks in ρ and keeping the other clocks unchanged.

A parameter valuation v is a function $v : \Gamma \rightarrow \mathbb{Q}_+$. We identify a parameter valuation v with the *point* $(v(\gamma_1), \dots, v(\gamma_M))$.

In the following, we assume $\bowtie \in \{<, \leq, \geq, >\}$. Let *aft* range over affine terms over $X \cup \Gamma$, of the form $\sum_{1 \leq i \leq H} \alpha_i x_i + \sum_{1 \leq j \leq M} \beta_j \gamma_j + d$, with $x_i \in X$, $\gamma_j \in \Gamma$, and $\alpha_i, \beta_j, d \in \mathbb{Z}$. Similarly, let *paft* range over parametric affine terms over Γ , that is affine terms without clocks ($\alpha_i = 0$ for all i). A *constraint* C over $X \cup \Gamma$ is a conjunction of inequalities of the form *aft* \bowtie 0 (i.e., a convex polyhedron). Given a parameter valuation v , $v(C)$ denotes the constraint over X obtained by replacing each parameter γ in C with $v(\gamma)$. Likewise, given a clock valuation w , $w(v(C))$ denotes the expression obtained by replacing each clock x in $v(C)$ with $w(x)$. We say that v *satisfies* C , denoted by $v \models C$, if the set of clock valuations satisfying $v(C)$ is nonempty. Given a parameter valuation v and a clock valuation w , we denote by $w|v$ the valuation over $X \cup \Gamma$ such that for all clocks x , $w|v(x) = w(x)$ and for all parameters γ , $w|v(\gamma) = v(\gamma)$. We use the notation $w|v \models C$ to indicate that $w(v(C))$ evaluates to true. We say that C is *satisfiable* if $\exists w, v$ s.t. $w|v \models C$. We define the *time elapsing* of C , denoted by C^\nearrow , as the constraint over X and Γ obtained from C by delaying all clocks by an arbitrary amount of time. Given $\rho \subseteq X$, we define the *reset* of C , written $[C]_\rho$, as the constraint obtained from C by resetting the clocks in ρ , and keeping the other clocks unchanged. We denote by $C \downarrow_\Gamma$ the projection of C onto Γ , i.e., obtained by eliminating the clock variables (e.g., using the Fourier-Motzkin algorithm).

A *guard* g is a constraint over $X \cup \Gamma$ defined by inequalities of the form $x \bowtie z$, where $x \in X$ and z is either a parameter or a constant in \mathbb{Z} .

A *zone* is a polyhedron over a set of clocks in which all constraints on variables are of the form $x \bowtie k$ (rectangular constraints) or $x_i - x_j \bowtie k$ (diagonal constraints), where $x_i \in X$, $x_j \in X$ and k is an integer. Operations on zones are well-documented (see e.g., Bengtsson and Yi [12]).

A *parametric zone* is a convex polyhedron over $X \cup \Gamma$ in which all constraints on variables are of the form $x \bowtie \text{paft}$ (parametric rectangular constraints) or $x_i - x_j \bowtie \text{paft}$ (parametric diagonal constraints), where $x_i \in X$, $x_j \in X$ and *paft* is a parametric affine term over Γ . We denote the set of all parametric zones by \mathcal{Z} .

2.2. Probabilistic timed automata

We start by reviewing the definition of timed probabilistic systems, as defined in Kwiatkowska et al. [26]. A *timed probabilistic system (TPS)* is a tuple $\mathcal{T} = (S, s_0, \Sigma, \Rightarrow)$ where S is a set of *states*, $s_0 \in S$ is the *initial state*, Σ is a finite set of *actions*, and $\Rightarrow \subseteq S \times \mathbb{R}_+ \times \Sigma \times \text{Dist}(S)$ is a *probabilistic transition relation* that associates a probabilistic distribution over S to triples made of a source state in S , a time in \mathbb{R}_+ and an action in Σ .

Probabilistic timed automata (defined by Gregersen and Jensen [20], Kwiatkowska et al. [26]) are an extension of classical timed automata (defined in Alur and Dill [1]) with discrete probability distributions.

2.2.1. Syntax

Definition 1. A Probabilistic Timed Automaton ($\mathbb{P}TA$) \mathcal{P} is a tuple $(\Sigma, L, l_0, X, prob)$, where: i) Σ is a finite set of actions, ii) L is a finite set of locations, iii) $l_0 \in L$ is the initial location, iv) X is a finite set of clocks, v) $prob$ is a *probabilistic edge relation* consisting of elements of the form (l, g, a, ν) , where $l \in L$, g is a zone over the clocks X , $a \in \Sigma$, and $\nu \in \text{Dist}(2^X \times L)$.

Note that we use no invariant; this is an important condition for the correctness of our techniques. However, invariants can be eliminated (moved to the guards prior to the transition), following classical techniques defined for (probabilistic) timed automata.

We use the following conventions for the graphical representation of probabilistic timed automata: locations are represented by nodes, within which name of the location is written; probabilistic edges are represented by arcs from locations, labeled by the associated guard and action, and which split into multiple arcs, each of which leads to a location and which is labeled by a set of clocks to be reset to 0 and a probability (probabilistic edges which correspond to probability 1 are illustrated by a single arc from location to location).

Example 2. Fig. 2a presents an example of a $\mathbb{P}TA$ with two clocks x and y . For example, l_0 can be exited whenever $y < 2$; then, with probability 0.4 the target location becomes l_2 , resetting x ; or with probability 0.6 the target location is l_1 , resetting y . The transition from l_2 can be explained similarly.

2.2.2. Semantics of $\mathbb{P}TAs$

A $\mathbb{P}TA$ can be interpreted as an infinite TPS. Due to the continuous nature of clocks, the underlying TPS has uncountably many states, and is uncountably branching.

Definition 2 (Concrete semantics of a $\mathbb{P}TA$). Given a $\mathbb{P}TA$ $\mathcal{P} = (\Sigma, L, l_0, X, prob)$, where $H = |X|$, the concrete semantics of \mathcal{P} is given by the timed probabilistic system $\mathcal{T}_{\mathcal{P}} = (S, s_0, \Sigma, \Rightarrow)$, with

- $S = \{(l, w) \in L \times \mathbb{R}_+^H\}$, $s_0 = (l_0, \vec{0})$,
- $((l, w), d, a, \eta) \in \Rightarrow$ if both of the following conditions hold:
 1. time elapse: $\forall d' \in [0, d], (l, w + d') \in S$, and
 2. edge traversal: there exists a probabilistic edge $e = (l, g, a, \nu) \in prob$ such that $w + d \models g$ and, for each $l' \in L$ and $\rho \subseteq X$, $\eta(l', [w + d]_{\rho}) = \nu(\rho, l')$.

Note that, due to the fact that we have no invariants, the first condition (time elapse) is always trivially true.

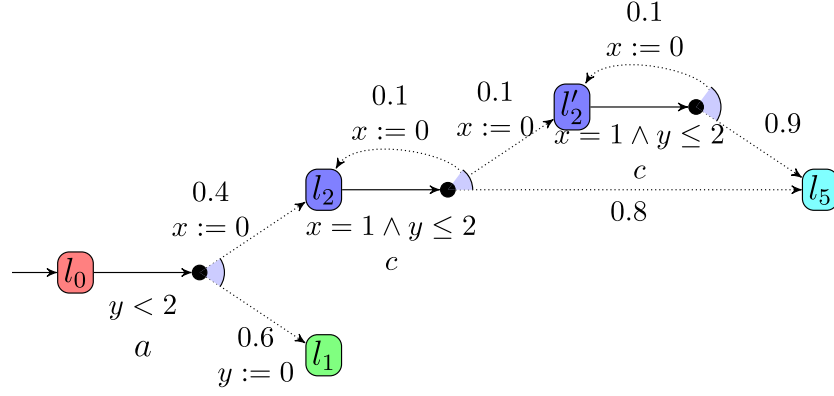
2.3. Parametric interval probabilistic timed automata

In this section, we introduce basic definitions for (*parametric*) *interval probabilistic timed automata*, that extend (parametric) probabilistic timed automata by providing *intervals* for transition probabilities instead of exact probability values. In the spirit of (parametric) Interval Markov Chains defined in Delahaye [15], Delahaye et al. [18], (parametric) interval probabilistic timed automata are used for specifying potentially infinite families (sets) of probabilistic timed automata—those whose exact probability values match the specified intervals—with a finite structure of similar form.

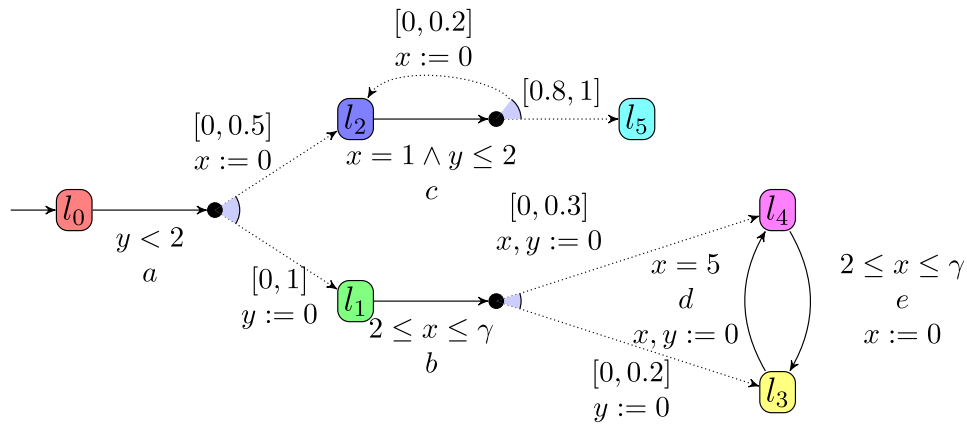
2.3.1. Syntax

Given an arbitrary measurable set S , we call an *interval distribution* over S a function Υ that assigns to each element of S an interval of probabilities $[a, b] \subseteq [0, 1]$. Intuitively, an interval distribution Υ over S represents the set of all distributions $\mu \in \text{Dist}(S)$ that assign to each element $s \in S$ a probability $\mu(s)$ such that $\mu(s) \in \Upsilon(s)$. Formally, let $\text{IntDist}(S)$ denote the set of all interval distributions over S ; we define the implementation of an interval distribution as follows.

Definition 3 (Implementation of an interval distribution). Let S be an arbitrary set. Given an interval distribution $\Upsilon \in \text{IntDist}(S)$, $\nu \in \text{Dist}(S)$ is an *implementation* of Υ , written $\nu \in \Upsilon$ iff, for all $s \in S$, we have $\nu(s) \in \Upsilon(s)$.



(a) A PTA



(b) A PIPTA

Fig. 2. Examples.

We now move to the definition of (parametric) interval probabilistic timed automata.

Definition 4. A Parametric Interval Probabilistic Timed Automaton (PIPTA) \mathcal{PIPTA} is a tuple $(\Sigma, L, l_0, X, \Gamma, \mathbb{I})$, where: i) Σ is a finite set of actions, ii) L is a finite set of locations, iii) $l_0 \in L$ is the initial location, iv) X is a finite set of clocks, v) Γ is a finite set of parameters, vi) \mathbb{I} is an *interval-valued probabilistic edge relation* consisting of elements of the form (l, g, a, Υ) , where $l \in L$, g is a guard, $a \in \Sigma$, and $\Upsilon \in \text{IntDist}(2^X \times L)$ is an interval distribution.

Given a PIPTA $\mathcal{PIPTA} = (\Sigma, L, l_0, X, \Gamma, \mathbb{I})$ and a parameter valuation v , the *valuation* of \mathcal{PIPTA} with v , written $v(\mathcal{PIPTA})$, is an Interval Probabilistic Timed Automaton (IPTA) $\mathcal{IPTA} = (\Sigma, L, l_0, X, \mathbb{I}')$, where \mathbb{I}' is obtained by replacing within \mathbb{I} any occurrence of a parameter γ with $v(\gamma)$ and removing all transitions (l, g, a, Υ) such that $v(g) \equiv \perp$ (technically, this latter part is not strictly speaking necessary, but it syntactically reduces the model a bit).

Remark that IPTAs are very similar to PTAs: the only difference is that probabilistic edges are labeled with intervals instead of exact probability values.

In our graphical representations, when the interval associated with a distribution is reduced to a point (e.g., $[0.5, 0.5]$), we simply represent it using its punctual value (i.e., 0.5). Also, when a distribution is made of a single target location with probability 1, we simply omit the distribution.

Once a parameter valuation is fixed, the resulting IPTA represents a potentially infinite set of PTAs. In order to relate a given IPTA with the PTAs it represents, we use the notion of *implementation* defined hereafter. This notion is similar to the one defined in the context of (parametric) Interval Markov Chains in Delahaye [15], Delahaye et al. [18]. Remark that a PTA implementing an IPTA needs to conserve the exact same clocks, guards and resets.

Definition 5 (Implementation of an IP_{TA}). Let $\mathcal{P} = (\Sigma, L, l_0, X, prob)$ be a IP_{TA} and $\mathcal{IP} = (\Sigma, L', l'_0, X, \mathbb{I})$ be an IP_{TA}.

We say that \mathcal{P} is an implementation of \mathcal{IP} , written $\mathcal{P} \models \mathcal{IP}$, iff there exists a relation $\mathcal{R}_P \subseteq L \times L'$, called an *implementation relation* s.t. $(l_0, l'_0) \in \mathcal{R}_P$ and, whenever $(l, l') \in \mathcal{R}_P$, we have

- $\forall (l, g, a, v) \in prob, \exists (l', g, a, \Upsilon) \in \mathbb{I}$ s.t. $v \preceq_{\mathcal{R}_P} \Upsilon$, and
- $\forall (l', g', a, \Upsilon) \in \mathbb{I}, \exists (l, g', a, v) \in prob$ s.t. $v \preceq_{\mathcal{R}_P} \Upsilon$,

where $v \preceq_{\mathcal{R}_P} \Upsilon$ iff $\exists \delta \in \text{Dist}(L \times L')$ s.t.

- $\forall (\rho, l) \in 2^X \times L, v(\rho, l) > 0 \Rightarrow \sum_{l' \in L'} (\delta(l, l')) = 1$,
- $\forall (\rho', l') \in 2^X \times L', \sum_{l \in L} (v(\rho', l) \cdot \delta(l, l')) \in \Upsilon(\rho', l')$, and
- $\delta(l, l') > 0 \Rightarrow (l, l') \in \mathcal{R}_P$.

In the above definition, the relation \mathcal{R}_P encodes the pairs of states $(l, l') \in L \times L'$ where l is an *implementation* of l' . On the other hand, the relation $\preceq_{\mathcal{R}_P}$ is a lifting of the relation \mathcal{R}_P to distributions over locations (also called a *coupling*), and therefore represents *compatible* distributions w.r.t. \mathcal{R}_P . This notion of satisfaction has been adapted from the notion of “weak weak” satisfaction in the context of Abstract Probabilistic Automata, for which several notions of satisfaction exist. We have chosen this particular notion because it is the most permissive among those presented in Delahaye et al. [16]. For a detailed discussion on this topic, we refer the interested reader to Delahaye et al. [16].

Given an IP_{TA}, deciding whether the family it represents is nonempty is a nontrivial problem. Indeed, the interval distributions used throughout its structure could represent contradictory constraints on the transition probabilities, therefore preventing any IP_{TA} from implementing it.

In the following, we say that a IP_{TA} \mathcal{P} has *the same structure* as an IP_{TA} \mathcal{IP} if and only if the underlying directed graph of \mathcal{P} is a subgraph (up to renaming and removal of unreachable states) of the underlying directed graph of \mathcal{IP} . This notion trivially extends to PIIP_{TA} and other models we use in the rest of the paper.

Definition 6 (Consistency of an IP_{TA}). An IP_{TA} is consistent if it admits at least one implementation.

Example 3. Consider the PIIP_{TA} \mathcal{PIP} given in Fig. 2b, and containing a single parameter γ . Let v_1 be the parameter valuation such that $v_1(\gamma) = 1$. In the IP_{TA} $v_1(\mathcal{PIP})$, the transition outgoing from l_1 can never be taken, as its guard becomes $2 \leq x \leq 1$, which is unsatisfiable. Then, it is clear that the IP_{TA} \mathcal{P} given in Fig. 2a is an implementation of $v_1(\mathcal{PIP})$. We emphasize the fact that location l_2 from $v_1(\mathcal{PIP})$ has been “unfolded” in \mathcal{P} , yielding two locations l_2 and l'_2 . As a consequence, the underlying structure of \mathcal{P} is not identical to the one of $v_1(\mathcal{PIP})$. Nevertheless, both locations l_2 and l'_2 of \mathcal{P} obviously satisfy the original l_2 from $v_1(\mathcal{PIP})$, which allows \mathcal{P} to satisfy $v_1(\mathcal{PIP})$ despite their distinct structures. As a consequence, $v_1(\mathcal{PIP})$ is a consistent IP_{TA}.

An important problem is therefore to decide whether a given IP_{TA} is consistent, which we address in the next section.

3. The consistency problem for IP_{TAs}

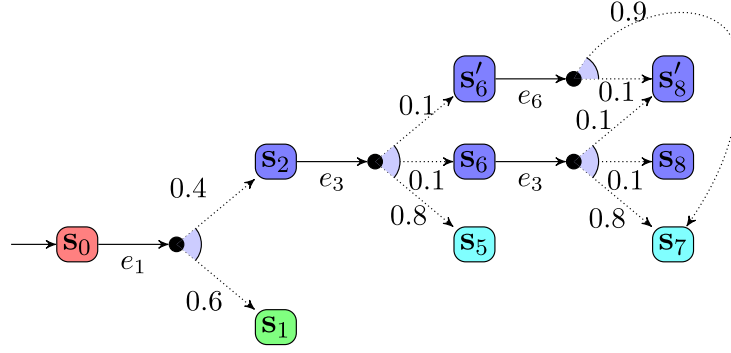
In this section, we address the problem of deciding whether a given IP_{TA} is consistent. Unlike in the context of IMCs, where it is proven that a given IMC is consistent iff it admits an implementation with the same structure, a given IP_{TA} can be consistent but still not admit any implementation that respects its structure. Indeed, the structure of implementations depends on the structure of the zone graph rather than on the structure of the IP_{TA} itself which can be different. Algorithms such as those proposed for deciding consistency of (p)IMCs in Delahaye et al. [18] therefore cannot be directly adapted to the IP_{TAs} setting as they are dependent on this property.

Fortunately, the operational semantics of IP_{TAs} can be expressed in terms of Interval Markov Decision Processes (IMDPs), which are similar to IMCs and satisfy the same structural properties regarding consistency. We therefore propose an algorithm for deciding consistency of IP_{TAs} based on the consistency of their symbolic IMDP semantics. An alternative solution would be to “normalize” IP_{TAs} into special IP_{TAs} where all edges can fire, via a region construction. For the sake of simplicity, we only explore the first solution. We start with preliminary definitions on IMDPs, then formally define the symbolic semantics of IP_{TAs} and finally propose an algorithm for deciding whether a given IP_{TA} is consistent.

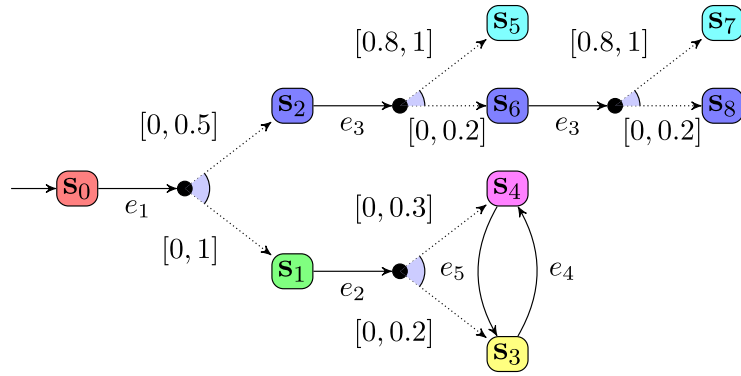
3.1. Preliminary definitions

An IMDP is a tuple (S, s_0, Σ, T) where S is a set of states, $s_0 \in S$ is the initial state, Σ is a finite set of actions and $T \subseteq S \times \Sigma \times \text{IntDist}(S)$ is a probabilistic (interval) transition relation.

Example 4. Fig. 3b depicts an example of an IMDP. Just as for IP_{TAs}, when the interval associated with a distribution is reduced to a point (which is not the case here), we simply represent it using its punctual value. When a distribution is made of a single target location with probability 1, we simply omit the distribution (e.g., from s_3 to s_4).



(a) An example of an MDP



(b) An example of an IMDP

Fig. 3. Examples.

Definition 7 (MDP). An MDP is an IMDP such that for each $(s, a, I) \in T$, and for all $s' \in S$, we have $I(s') = [m, m]$ is a singleton. In addition, for each $(s, a, I) \in T$, we have

$$\sum_{s' \in S} I(s') = 1.$$

Example 5. Fig. 3a depicts an example of an MDP.

Definition 8 (Implementation of an IMDP). Let $\mathcal{IM} = (S, s_0, \Sigma, T)$ be an IMDP. Let $\mathcal{M} = (S', s'_0, \Sigma, T')$ be an MDP. We say that \mathcal{M} is an implementation of \mathcal{IM} , written $\mathcal{M} \models \mathcal{IM}$, if $\exists \mathcal{R}_M \subseteq S' \times S$ s.t. $(s'_0, s_0) \in \mathcal{R}_M$ and $(s', s) \in \mathcal{R}_M$ if

- $\forall (s', a, \iota) \in T', \exists (s, a, I) \in T$ s.t. $\iota \preceq_{\mathcal{R}_M} I$, and
- $\forall (s, a, I) \in T, \exists (s', a, \iota) \in T'$ s.t. $\iota \preceq_{\mathcal{R}_M} I$,

where $\iota \preceq_{\mathcal{R}_M} I$ iff $\exists \delta \in \text{Dist}(S' \times S)$ s.t.

- $\forall s' \in S', \iota(s') > 0 \Rightarrow \sum_{s \in S} (\delta(s', s)) = 1$,
- $\forall s \in S, \sum_{s' \in S'} (\iota(s') \cdot \delta(s', s)) \in I(s)$, and
- $\delta(s', s) > 0 \Rightarrow (s', s) \in \mathcal{R}_M$.

As for IPTAs, we say that an IMDP is *consistent* iff it admits at least one implementation.

Example 6. The IMDP given in Fig. 3b admits no implementation: indeed, on the (single) transition labeled with e_2 , no valuation of the two intervals $[0, 0.3]$ and $[0, 0.2]$ is such that the sum of both valuations is equal to 1. Nevertheless, it could be that the IMDP is still consistent if one assigns a 0-probability on the transition from s_0 to s_1 . However, although this would be compatible with the interval $(0 \in [0, 1])$, the second interval (to s_2) does not accept a 1-probability since its probability must be within $[0, 0.5]$.

As said above IMDPs satisfy the same structural property as IMCs concerning implementations: they are consistent iff they admit at least one implementation that respects their structure. This result is formalized in the following lemma.

Lemma 1 (*Structure of an implementation*). *An IMDP \mathcal{IM} is consistent iff there exists an MDP \mathcal{M} with the same structure s.t. $\mathcal{M} \models \mathcal{IM}$.*

Proof. Let $\mathcal{IM} = (S, s_0, \Sigma, T)$ be an IMDP.

One direction of this result is trivial: if there exists an MDP \mathcal{M} with the same structure as \mathcal{IM} s.t. $\mathcal{M} \models \mathcal{IM}$, then \mathcal{IM} is clearly consistent.

The reverse implication is more involved. Assume that \mathcal{IM} is consistent, i.e., there exists an MDP $\mathcal{M} = (S', s'_0, \Sigma, T')$, with no assumption on its structure, such that $\mathcal{M} \models \mathcal{IM}$. We then have to build an MDP $\mathcal{M}^* = (S, s_0, \Sigma, T^*)$ such that $\mathcal{M}^* \models \mathcal{IM}$. Observe that S and s_0 must be identical to that of \mathcal{IM} because they have the same structure.

Let \mathcal{R}_M be the relation witnessing that $\mathcal{M} \models \mathcal{IM}$ and let $f : S \rightarrow S' \cup \{\perp\}$ be a function that associates to all states in \mathcal{IM} one of the states from \mathcal{M} that contributes to its implementation, if there is any, and \perp otherwise. Formally, for all $s \in S$, if $f(s) \neq \perp$ then $(f(s), s) \in \mathcal{R}_M$, and whenever there exists $s' \in S'$ such that $(s', s) \in \mathcal{R}_M$, we have $f(s) \neq \perp$.

The transition relation T^* of \mathcal{M}^* is constructed as follows: For each state s that is implemented, i.e., such that $f(s) \neq \perp$, and probabilistic interval transition $(s, a, I) \in T$ in \mathcal{IM} , we build a corresponding transition (s, a, I') in \mathcal{M}^* from the transitions in \mathcal{M} that implement (s, a, I) . In other words, we pick one of the states that satisfy s (using function f) and mimic its outgoing transitions in \mathcal{M}^* . All the other states that satisfy s are simply removed. States that are not implemented do not serve for consistency and are therefore not considered.

Formally, let $(s_1, a, I) \in T$ be a probabilistic interval transition in \mathcal{IM} . From Definition 8, we know that there exists $(f(s_1), a, \iota) \in T'$ s.t. $\iota \preceq_{\mathcal{R}_M} I$. According to the definition, there exists at least one function δ that witnesses $\iota \preceq_{\mathcal{R}_M} I$. In the following we pick one such function and name it $\delta_{(\iota, I)}$. The distribution ι' is then constructed as follows: for all $s_2 \in S$, let $\iota'(s_2) = \sum_{s' \in S'} \iota(s') \cdot \delta_{(\iota, I)}(s', s_2)$.

By definition of $\delta_{(\iota, I)}$, observe that $\iota'(s_2) \in I(s_2)$ for all $s_2 \in S$ and that, whenever $\iota'(s_2) > 0$, $f(s_2) \neq \perp$.

Clearly, \mathcal{M}^* is therefore an implementation of \mathcal{IM} , with witnessing relation \mathcal{R}_M^* defined as the identity relation on the set of states $s \in S$ such that $f(s) \neq \perp$. \square

3.2. A symbolic semantics for IPTAs

We equip IPTAs with a symbolic semantics, defined below. Basically, it is inline with the symbolic semantics defined for timed automata in the form of a zone graph, with the addition of probabilistic intervals on the edges; as a consequence, the semantics becomes not an LTS, but an IMDP.

Definition 9 (*Symbolic semantics of an IPTA*). Given an IPTA $\mathcal{IP} = (\Sigma, L, l_0, X, \mathbb{I})$, the symbolic semantics of \mathcal{IP} is given by the IMDP $(\mathbf{S}, \mathbf{s}_0, \mathbb{I}, T)$, with

- $\mathbf{S} = \{(l, C) \in L \times \mathcal{Z}\}$, $\mathbf{s}_0 = (l_0, (\bigwedge_{1 \leq i \leq H} x_i = 0)^\nearrow)$, where l is the location and C the associated zone,
- $((l, C), e, I) \in T$ if $e = (l, g, a, \Upsilon) \in \mathbb{I}$ and for all $l' \in L$, for all $\rho \subseteq X$ such that $\Upsilon(\rho, l') > 0$, $C' = ([C \wedge g]_\rho)^\nearrow$, and $I((l', C')) = \Upsilon(\rho, l')$.

Given a symbolic state $\mathbf{s} = (l, C)$, we denote by $\mathbf{s}.l$ and $\mathbf{s}.C$ its location and its associated zone (symbolic constraint), respectively.

Observe that, whenever an IPTA has no probabilistic choice, then the IMDP becomes a labeled transition system, and the symbolic semantics matches that of timed automata given in the form of a zone graph (see e.g., Bengtsson and Yi [12]). It is well-known that the zone graph of a timed automaton can have an infinite number of states; however, applying the classical k -extrapolation (that basically splits zones between a part where the clock constraints are smaller or equal to k and a part where constraints are larger than k , where k is the largest integer-constant in the timed automaton) yields finiteness (see, e.g., Behrmann et al. [10]). In the following, we apply the classical k -extrapolation to the symbolic constraints of the semantics of an IPTA \mathcal{IP} , and therefore the number of states in the IMDP described in Definition 9 is finite. We refer to the symbolic semantics of \mathcal{IP} as the *probabilistic zone graph* of \mathcal{IP} .

Remark that the probabilistic zone graph is defined for IPTAs in the form of an IMDP; a PTA can be understood as an IPTA, and its associated zone graph becomes an MDP.

Example 7. The probabilistic zone graph of the PTA in Fig. 2a is the MDP given in Fig. 3a. The symbolic states $\mathbf{s}_i = (l_i, C_i)$ are expanded in Table 1.

Table 1
Description of the states in Fig. 3a.

State	Location	C
s_0	l_0	$x = y \wedge x \geq 0$
s_1	l_1	$0 \leq x - y < 2 \wedge y \geq 0$
s_2	l_2	$0 \leq y - x < 2 \wedge x \geq 0$
s_5	l_5	$0 \leq y - x \leq 1 \wedge x \geq 1$
s_6	l_2	$1 \leq y - x \leq 2 \wedge x \geq 0$
s'_6	l'_2	$1 \leq y - x \leq 2 \wedge x \geq 0$
s_7	l_5	$y \geq 2 \wedge y = x + 1$
s_8	l_2	$y \geq 2 \wedge y = x + 2$
s'_8	l'_2	$y \geq 2 \wedge y = x + 2$

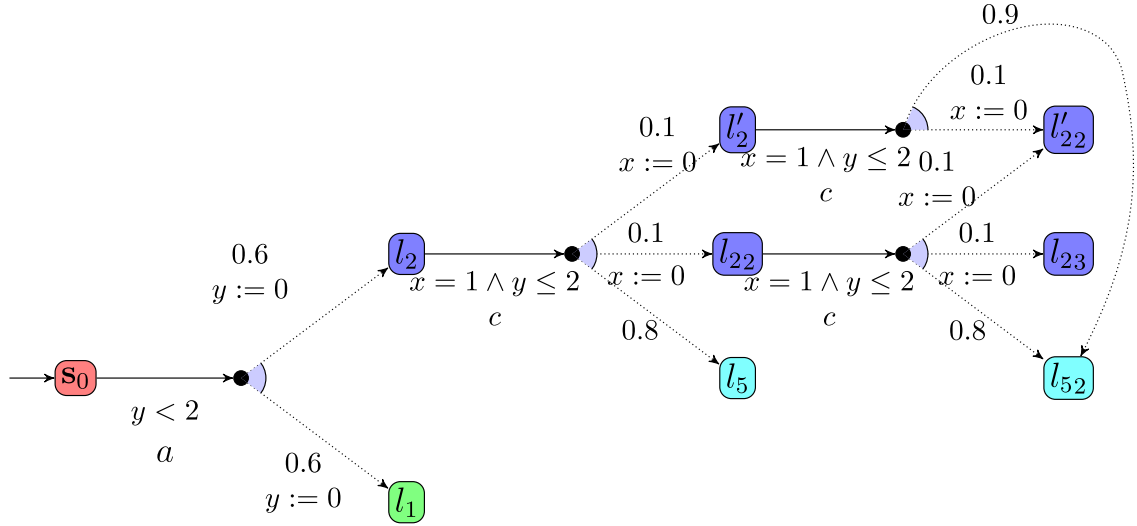


Fig. 4. A PTA reconstructed from the probabilistic zone graph in Fig. 3a.

3.3. Reconstructing an IPPTA from a probabilistic zone graph

It is well-known that, given a timed automata \mathcal{A} and its zone graph, a second timed automaton \mathcal{A}' can be reconstructed from the zone graph, with the same structure as the zone graph, and such that the zone graph of \mathcal{A}' is the same as that of \mathcal{A} . We extend this technique here to IPPTAs.

The construction. Let $\mathcal{IP} = (\Sigma, L, l_0, X, \mathbb{I})$ be an IPPTA; let $\mathcal{IM} = (\mathbf{S}, s_0, \mathbb{I}, T)$ be its probabilistic zone graph. Let us build a second IPPTA $\mathcal{IP}' = (\Sigma, L', l'_0, X, \mathbb{I}')$ as follows.

First, each state of \mathcal{IM} is translated into a location of \mathcal{IP}' , i.e., we have $L' = \mathbf{S}$.

Second, the initial location of \mathcal{IP}' is the initial state of \mathcal{IM} , i.e., we have $l'_0 = s_0$.

Third, for each transition $(s, e, I) \in T$ in \mathcal{IM} , with $e = (l, g, a, \Upsilon)$, we create in \mathcal{IP}' a transition (s, g, a, Υ') , where Υ' is defined as follows: for each s' such that $I(s') > 0$, then $\Upsilon'(\rho', s') = I(s')$, where ρ' is the set of clocks to be reset from $s.l$ to $s'.l$ via edge e in \mathcal{IP} .

Given an IPPTA \mathcal{IP} with probabilistic zone graph \mathcal{IM} . We denote by $\text{Reconstruct}(\mathcal{IM})$ the IPPTA \mathcal{IP}' reconstructed from \mathcal{IM} following the above technique. Obviously, this construction also applies to PTA, which are just IPPTA where intervals are reduced to single points.

An equivalence result. As should be expected, the probabilistic zone graph \mathcal{IM}' of the IPPTA \mathcal{IP}' reconstructed from the probabilistic zone graph \mathcal{IM} of a IPPTA \mathcal{IP} is equivalent to \mathcal{IM} .

Proposition 1. Let \mathcal{IP} be an IPPTA and \mathcal{IM} be its probabilistic zone graph. Let $\mathcal{IP}' = \text{Reconstruct}(\mathcal{IM})$. Let \mathcal{IM}' be the probabilistic zone graph of \mathcal{IP}' .

Then \mathcal{IM}' is equivalent to \mathcal{IM} up to location renaming.

Example 8. We apply the above procedure to the probabilistic zone graphs from Fig. 3a and Fig. 3b. The PTA and IPPTA reconstructed from these zone graphs are given in Fig. 4 and Fig. 5, respectively. Remark that their probabilistic zone graphs are again that of Fig. 3a and Fig. 3b.

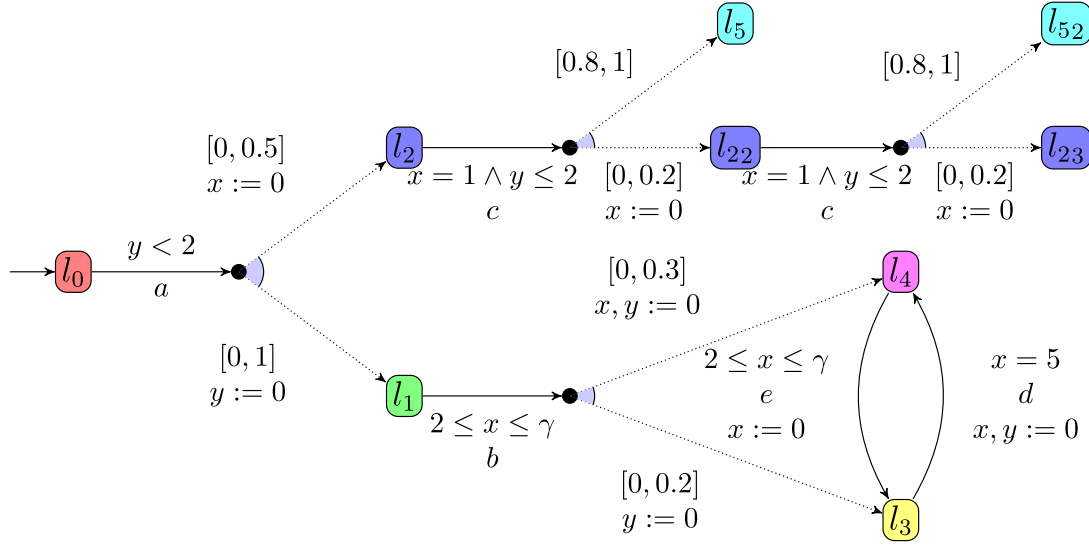


Fig. 5. An IPPTA reconstructed from the probabilistic zone graph in Fig. 3b.

3.4. An algorithm for the consistency of IPPTAs

We start with the following observation: by construction, the purpose of the symbolic semantics of IPPTAs is to represent, at a lower level of abstraction, the same set of objects. Intuitively, the symbolic IMDP semantics of a given IPPTA should therefore be consistent iff the original IPPTA is itself consistent. This result is formally proven in Proposition 2.

Proposition 2. An IPPTA \mathcal{IP} is consistent iff its probabilistic zone graph is consistent.

Proof. Let $\mathcal{IP} = (\Sigma, L, l_0, X, \mathbb{I})$ be an IPPTA. Let $\mathcal{IM} = (\mathbf{S}, \mathbf{s}_0, \mathbb{I}, T)$ be the probabilistic zone graph of \mathcal{IP} .

\Rightarrow Assume \mathcal{IP} is consistent, and let us show that its probabilistic zone graph is consistent. From the definition of consistency, there exists a PTA $\mathcal{P} = (\Sigma, L', l'_0, X, prob)$ such that $\mathcal{P} \models \mathcal{IP}$, with implementation relation \mathcal{R}_P . Let $\mathcal{M} = (\mathbf{S}', \mathbf{s}'_0, prob, T')$ be the probabilistic zone graph of \mathcal{P} . Let us show that $\mathcal{M} \models \mathcal{IM}$.

We therefore define a relation \mathcal{R}_M , and show that it is an implementation relation. We define $\mathcal{R}_M = \{((l, C), (l', C')) \mid (l, l') \in \mathcal{R}_P \wedge C = C'\}$.

– From Definition 9, the initial state of \mathcal{M} is $\mathbf{s}'_0 = (l'_0, (\bigwedge_{1 \leq i \leq H} x_i = 0)^\nearrow)$; the initial state of \mathcal{IM} is $\mathbf{s}_0 = (l_0, (\bigwedge_{1 \leq i \leq H} x_i = 0)^\nearrow)$. Since $\mathcal{P} \models \mathcal{IP}$ then from Definition 5 we have $(l'_0, l_0) \in \mathcal{R}_P$, and therefore $(\mathbf{s}'_0, \mathbf{s}_0) \in \mathcal{R}_M$.

– Let $((l, C), (l', C)) \in \mathcal{R}_M$.

* Let $((l', C), a, \iota) \in T'$. Since Definition 9, there exists an edge $e' = (l', g, a, \gamma) \in prob$. Therefore, by \mathcal{R}_P , there exists an edge $e = (l, g, a, \gamma) \in \mathbb{I}$ such that $\gamma \leq_{\mathcal{R}_P} \gamma$.

As a consequence, by Definition 9 and since the guards are the same in \mathcal{P} and \mathcal{IP} , there exists $((l, C), a, I) \in T$.

Moreover, by Definition 9, we have $\iota \leq_{\mathcal{R}_M} I$, with $\delta_{\mathcal{R}_M}((l', C'), (l, C)) = \delta_{\mathcal{R}_P}(l', l)$ if $C = C'$ and 0 otherwise.

* Similarly, for all $((l, C), a, I) \in T$, there exists $((l', C), a, \iota) \in T'$ such that $\iota \leq_{\mathcal{R}_M} I$ by \mathcal{R}_P and Definition 9.

Therefore $\mathcal{M} \models \mathcal{IM}$.

\Leftarrow Assume the probabilistic zone graph \mathcal{IM} of \mathcal{IP} is consistent, and let us show that \mathcal{IP} is consistent. From Lemma 1, there exists $\mathcal{M} = (\mathbf{S}, \mathbf{s}_0, \mathbb{I}, T')$ such that $\mathcal{M} \models \mathcal{IM}$ with the same structure as \mathcal{IM} , and with implementation relation \mathcal{R}_M (note that \mathcal{R}_M is the identity because they have the same structure).

Let us reconstruct an IPPTA $\mathcal{IP}' = (\Sigma, L', l'_0, X, \mathbb{I}')$ from the probabilistic zone graph \mathcal{IM} , using the procedure Reconstruct from Section 3.3. Now, let $\mathcal{P} = (\Sigma, L', l'_0, X, prob)$ with the same structure as \mathcal{IP}' , and where $prob$ is obtained by replacing every occurrence of I in \mathbb{I}' by ι taken from T in \mathcal{M} . Note that there is a one-to-one correspondence between T' and T since they have the same structure, which is a key point here.

Recall that, during Reconstruct, for each transition $(\mathbf{s}, e, I) \in T$ in \mathcal{IM} , with $e = (l, g, a, \gamma)$, we create in \mathcal{IP}' a transition $(\mathbf{s}, g, a, \gamma')$, where γ' is defined as follows: for each \mathbf{s}' such that $I(\mathbf{s}') > 0$, then $\gamma'(\rho', \mathbf{s}') = I(\mathbf{s}')$, where ρ' is the set of clocks to be reset from $\mathbf{s}.l$ to $\mathbf{s}'.l$ via edge e in \mathcal{IP} . Here, we simply replace I with ι , where ι is the distribution corresponding to I in \mathcal{M} .

Now, let us show that $\mathcal{P} \models \mathcal{IP}$. Recall from Section 3.3 that the locations in \mathcal{P} are of the form (l, C) . We thus define $\mathcal{R}_P = \{(l', C'), I \mid l' = I\}$.

– From the reconstruction Reconstruct, the initial location of \mathcal{P} is (l_0, C_0) . Therefore, $((l_0, C_0), l_0) \in \mathcal{R}_P$.

– Let $((l, C), l') \in \mathcal{R}_P$.

- * Let $e = ((l, C), g, a, v) \in \text{prob}$. By Reconstruct, there must exist $((l, C), e, \iota) \in T'$ in \mathcal{M} with $v(\rho, (l', C')) = \iota((l', C'))$, for all (l', C') , where ρ is the set of clocks to be reset from l to l' via edge e in \mathcal{P} . Therefore, from \mathcal{R}_M , there exists $((l, C), e, I) \in T$ of \mathcal{IM} s.t. $\iota \preceq_{\mathcal{R}_M} I$.

As a consequence, by the zone graph construction (Definition 9), there exists a transition $(l, g, a, \Upsilon) \in \mathbb{I}$ in \mathcal{IP} such that $I((l', C')) = \Upsilon(\rho, l')$, for all ρ, l', C' .

Let $\delta_{\mathcal{R}_p}$ be such that $\delta_{\mathcal{R}_p}((l', C'), l'') = 1$ if $l' = l''$ and 0 otherwise. Finally, by \mathcal{R}_M , we obtain $v \preceq_{\mathcal{R}_p} \Upsilon$.

- * Similarly, for all $(l, g, a, \Upsilon) \in \mathbb{I}$ in \mathcal{IP} , there exists $((l, C), g, a, v) \in \text{prob}$ such that $v \preceq_{\mathcal{R}_p} \Upsilon$.

Therefore $\mathcal{P} \models \mathcal{IP}$. \square

Given the results presented in Lemma 1 and Proposition 2, deciding whether a given \mathbb{IPTA} \mathcal{IP} is consistent can be done by deciding whether its probabilistic zone graph admits at least one implementation that preserves its structure.

Such an algorithm was provided in Delahaye [15] in the context of IMCs instead of IMDPs. We show how this algorithm can be adapted to our context. As for IMCs, we say that a state is *locally inconsistent* in a given IMDP iff one of its outgoing probabilistic (interval) transitions cannot be implemented, i.e., if there is no distribution that matches the specified intervals. Let $\mathcal{IM} = (S, s_0, \mathbb{I}, T)$ be the IMDP symbolic semantics of a given \mathbb{IPTA} . Our algorithm is given in Algorithm 1.

Algorithm 1: Consistency of IMDPs.

```

1 Let Inc be the set of locally inconsistent states in  $\mathcal{IM}$  and Passed =  $\emptyset$ .
2 while  $s_0 \notin \text{Passed}$  and  $\text{Inc} \neq \emptyset$  do
3   Let  $s \in \text{Inc}$  and  $\text{Passed} = \text{Passed} \cup \{s\}$ .
4   Replace all transitions  $(s', a, I)$  such that  $I(s) \neq [0, 0]$  with  $(s', a, I')$  where
      •  $I'(s'') = I(s'')$  for all  $s'' \neq s$ , and
      •  $I'(s) = I(s) \cap [0, 0]$ 
   Update  $\text{Inc} \subseteq (S \setminus \text{Passed})$ .
```

Algorithm 1 is based on the following principle: as soon as a locally inconsistent state is detected, it is made unreachable by modifying the incoming interval probabilities to $I(s) \cap [0, 0]$. Remark that if 0 is not an admissible transition probability, the inconsistency is transferred to the predecessor states because $I(s) \cap [0, 0] = \emptyset$.

In the context of IMCs, it is proven in Delahaye [15] that this algorithm converges and that the original IMC is consistent iff the initial state is not locally inconsistent in the resulting IMC. The proof from Delahaye [15] can be trivially adapted to the context of IMDPs.

Proposition 2 together with Algorithm 1 and the above discussion on termination give the following theorem:

Theorem 1. *The consistency problem for \mathbb{IPTA} s is decidable.*

4. Consistency-emptiness and synthesis for \mathbb{PIPTA} s

We now move to the parametric setting and consider the following two problems:

Consistency-emptiness problem:

INPUT: A \mathbb{PIPTA} \mathcal{PIP}

PROBLEM: does there exist a parameter valuation v such that $v(\mathcal{PIP})$ is consistent?

Consistency-synthesis problem:

INPUT: A \mathbb{PIPTA} \mathcal{PIP}

PROBLEM: find all parameter valuations v for which $v(\mathcal{PIP})$ is consistent.

In the following, we first address the consistency-emptiness problem and show that, while this problem is undecidable in the general context of \mathbb{PIPTA} s, it becomes decidable for a syntactic subclass (Section 4.1). We then introduce an adaptation of the parametric zone-graph construction for parametric timed automata (Section 4.2), and propose in Section 4.3 a construction for the consistency-synthesis problem. This construction can only be applied when the parametric probabilistic zone-graph construction of the original \mathbb{PIPTA} is finite. When this is the case, the set of parameter values that are synthesized is exactly those that ensure consistency of the resulting \mathbb{IPTA} . We finally address the more general problem of consistent reachability in Section 4.4.

4.1. The emptiness problem

4.1.1. Undecidability in the general case

The undecidability of the consistency-emptiness for PIPPTAs follows from the undecidability of the reachability emptiness for parametric timed automata.

Theorem 2. *The consistency-emptiness for PIPPTAs is undecidable.*

Proof. The reachability emptiness for parametric timed automata (i.e., the existence of at least one parameter valuation for which a given location is reachable) is undecidable (see e.g., Alur et al. [2], Jovanović et al. [24], André and Markey [9], Beneš et al. [11], and André [3] for a complete survey). In particular, it is undecidable even without invariant as shown in Alur et al. [2], Beneš et al. [11], which is inline with our setting.

We prove our result by reducing from the reachability emptiness for parametric timed automata. Assume a PTA (without probability), with a special goal location. From that goal location, let us add an unguarded transition to a new location for which no implementation exists (for example a single transition labeled with $[0.5, 0.5]$). Hence there exists a parameter valuation for which the underlying IPTA admits no implementation iff there exists a parameter valuation for which the goal location is reachable—which is undecidable. \square

The undecidability of the consistency-emptiness problem rules out the possibility to, in general, compute a solution to the consistency-synthesis problem. In the following, we will still address this computation problem by proposing a synthesis procedure that can be used when the parametric probabilistic zone graph is finite.

4.1.2. A decidability result

Despite the negative result of Theorem 2, we can exhibit a decidability result for a syntactic subclass of PIPPTAs. In Hune et al. [21], a syntactic subclass, namely lower-bound/upper-bound parametric timed automata (L/U-PTAs) is introduced that restricts the use of parameters in parametric timed automata. Basically, in an L/U-PTA, any parameter must be either always used as an upper-bound (a constraint $x \leq \gamma$ or $x < \gamma$) or always as a lower-bound ($x \geq \gamma$ or $x > \gamma$) in the guards and invariants of the parametric timed automaton. L/U-PTAs benefit from several main decidability results: the EF-emptiness, or reachability-emptiness, problem (“is the set of parameter valuations for which a given location is reachable empty?”) is shown to be decidable in Hune et al. [21]. Then, the infinite acceptance emptiness (“is the set of parameter valuations for which a given set of locations is visited infinitely often along some run empty?”) and universality (“is a given set of locations visited infinitely often along some run for all parameter valuations?”) have been proved to be decidable for L/U-PTAs with integer-valued parameters in Bozzelli and La Torre [13]. Unavoidability was studied in Jovanović et al. [24] while liveness and deadlocks were studied in André and Lime [6] with a thin frontier between decidability and undecidability. Finally, the EF-universality problem was shown to be decidable for L/U-PTAs over rational parameters by André [3].

In the following, we reuse the concept of lower-bound and upper-bound parameters in the setting of PIPPTAs.

Definition 10 (L/U-PIPTA). An L/U-PIPTA is a PIPPTA whose set of parameters is partitioned into lower-bound parameters and upper-bound parameters, where an upper-bound (resp. lower-bound) parameter γ_i is such that, for every guard constraint $x \bowtie z$, we have: $z = \gamma_i$ implies $\bowtie \in \{\leq, <\}$ (resp. $\bowtie \in \{\geq, >\}$).

Example 9. Consider the PIPPTA in Fig. 7a. Then it is an L/U-PIPTA with one upper-bound parameter γ_1 and one lower-bound parameter γ_2 .

L/U-PTAs enjoy a well-known monotonicity property recalled in the following lemma (that corresponds to a reformulation of [21, Prop 4.2]), stating that increasing upper-bound parameters or decreasing lower-bound parameters can only add behaviors.

Lemma 2 (Hune et al. [21]). *Let \mathcal{A} be an L/U-PTA and v be a parameter valuation. Let v' be a valuation such that for each upper-bound parameter γ^+ , $v'(\gamma^+) \geq v(\gamma^+)$ and for each lower-bound parameter γ^- , $v'(\gamma^-) \leq v(\gamma^-)$. Then any run of $v(\mathcal{A})$ is a run of $v'(\mathcal{A})$.*

Remark 1. Unfortunately, we cannot directly use the monotonicity property of L/U-PTAs to prove the decidability of the consistency-emptiness for L/U-PIPTAs. It could have been helpful to consider the IPTA, say $\mathcal{IP}_{\infty,0}$, obtained by replacing every lower-bound parameter (resp. upper-bound parameter) in the guards of a PIPPTA \mathcal{PIP} with ∞ (resp. 0).² $\mathcal{IP}_{\infty,0}$ is the most restrictive IPTA obtained from \mathcal{PIP} (as any other valuation will have more behaviors thanks to Lemma 2). If

² Valuating a parameter with ∞ is achieved as follows: for each upper-bound parameter γ for which $v(\gamma) = \infty$, we delete any comparison of a clock with γ (i.e., the clock constraint becomes the most permissive); for each lower-bound parameter γ for which $v(\gamma) = \infty$, we replace any constraint in which γ appears by false (i.e., the transition labeled by the guard is deleted). Therefore the result of this valuation is an IPTA as expected.




Fig. 6. An L/U-PIPTA inconsistent for all parameter valuations.

$\mathcal{IP}_{\infty,0}$ is inconsistent, then any other parameter valuation is clearly inconsistent as well thanks to the monotonicity of L/U-PTAs; however, if $\mathcal{IP}_{\infty,0}$ is consistent, it is not possible to conclude that a (non-infinite) parameter valuation is consistent. In fact, it is easy to exhibit a counter-example for which $\mathcal{IP}_{\infty,0}$ is consistent, but for which $v(\mathcal{PIP})$ is inconsistent for any (non-infinite) valuation v . This is the case of the PIPTA depicted in Fig. 6: when γ is replaced with ∞ , the system is stuck forever in l_0 , and is therefore consistent. For any (non-infinite) parameter valuation, the system can take the transition labeled with a , which is inconsistent due to the sum of the probabilities.

Still, we show in Theorem 3 below that the consistency-emptiness problem is decidable in the context of L/U-PIPTAs.

To prove decidability, we use the following reasoning. An L/U-PIPTA is consistent if we can “block” the inconsistent edges, i.e., those who cannot admit any implementation because the sum of their probabilities cannot be 1. There are two ways of achieving this goal: either set to 0 some of the probabilities on all paths leading to a given inconsistent edge, or tune the timing parameters so as to forbid this edge because the guard can never be satisfied. The first way can be achieved by enumerating all possible combinations to set to 0 some probabilities. The second way can be achieved by parametric model checking: for a given combination of probabilities set to 0, if we can find at least one valuation for which none of the inconsistent edges is reachable, then we can answer false to the consistency-emptiness problem for L/U-PIPTAs. Finding at least one such valuation is equivalent to answering no to the EF-universality problem—which is decidable for L/U-PTAs as shown in André [3]. In the following, we explain this reasoning step by step.

We first need a notation, used in the proof of Theorem 3, and later in Section 4.3.

Definition 11 (*Feasible supports*). Given an interval distribution I in an IMDP, let $FS(I)$ denote the *feasible supports* of I i.e., set of sets of target states for which a consistent distribution can be assigned. Formally, $FS(I) = \{S' \subseteq S \mid \exists \mu \in \text{Dist}(S) \text{ s.t. } \forall s \in S : \mu(s) \in I(s) \text{ and } \forall s \in S : \mu(s) > 0 \text{ iff } s \in S'\}$.

That is, each set of states S' is such that there exists a distribution I' for which the probability of reaching each state in S' is not zero, such that this distribution is a punctual distribution, is an implementation of I and is consistent.

Definition 12. An interval distribution I is *inconsistent* if $FS(I) = \emptyset$.

By extension, we say that an edge is inconsistent if its interval distribution is inconsistent.

We also use the same notions for interval distributions in PIPTAs.

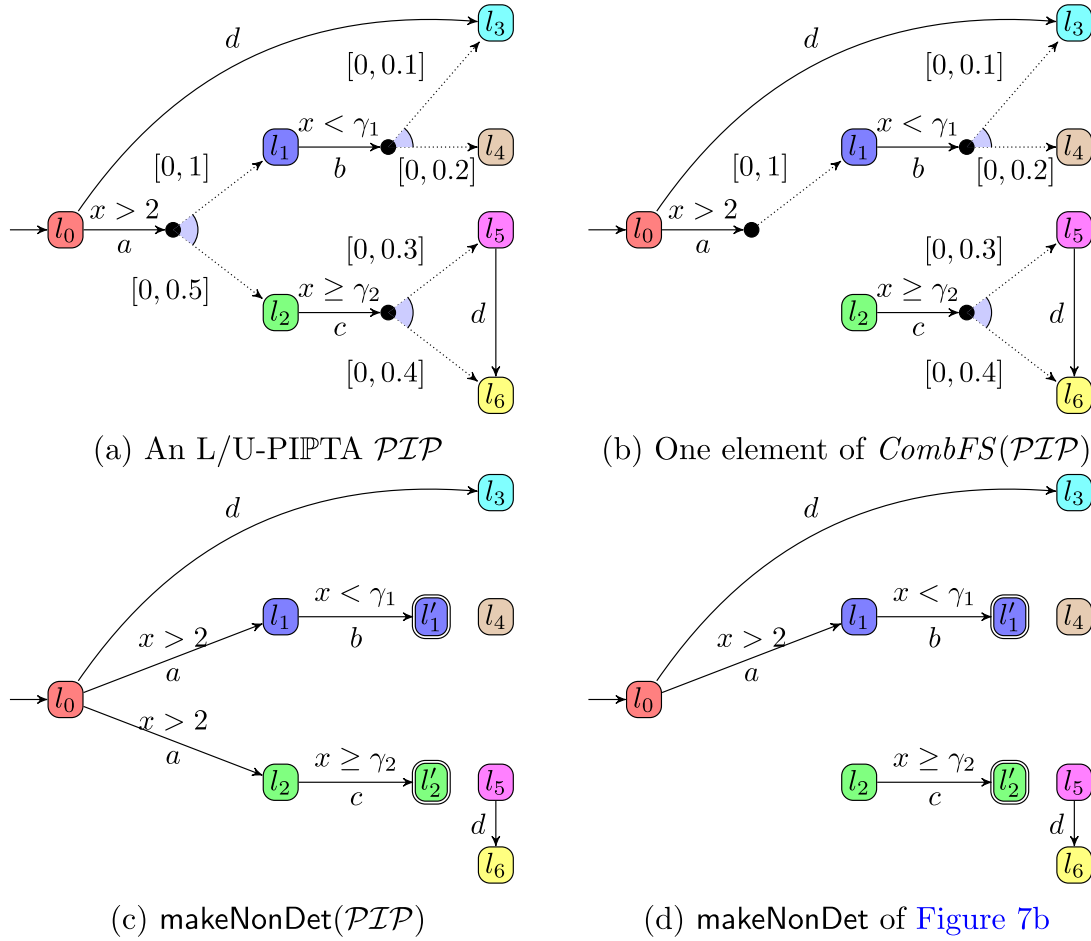
Example 10. In Fig. 2b, let Υ_1 be the (unique) interval distribution outgoing from l_1 . We have $FS(\Upsilon_1) = \{\}$, as no implementation can make this distribution consistent. That is, Υ_1 is an inconsistent edge. Let Υ_2 be the (unique) interval distribution outgoing from l_2 . We have $FS(\Upsilon_2) = \{\{l_5\}, \{l_2, l_5\}\}$.

We now define the set of PIPTAs obtained by taking all possible combinations of feasible supports.

Definition 13. Given a PIPTA \mathcal{PIP} , let $\text{CombFS}(\mathcal{PIP})$ denote the set of all possible PIPTAs obtained from \mathcal{PIP} by selecting for each interval distribution Υ exactly one element from $FS(\Upsilon)$ when $FS(\Upsilon) \neq \emptyset$, or by keeping the original distribution if $FS(\Upsilon) = \emptyset$.

Intuitively, $\text{CombFS}(\mathcal{PIP})$ contains all possible ways to remove transitions by setting some probabilities to 0 while keeping the sum of the other probabilities possibly equal to 1.

Example 11. Consider the PIPTA \mathcal{PIP} in Fig. 2b. $\text{CombFS}(\mathcal{PIP})$ contains 4 PIPTAs. All are such that l_1 has the same outgoing distribution to l_4 and l_3 as in Fig. 2b (as FS is empty for this distribution). Two of these 4 PIPTAs (say 1 and 2) are such that the distribution outgoing from l_0 goes to both l_1 and l_2 (with the same probabilities as in Fig. 2b), while two others (say 3 and 4) are such that this distribution is only going to l_1 . In addition, two of these 4 PIPTAs (say 1 and 3) are such that the distribution outgoing from l_2 goes to both l_2 and l_5 , while two others (say 2 and 4) are such that this distribution is only going to l_5 .

Fig. 7. Exemplifying CombFS , makeNonDet and makeAcc .

Example 12. Consider now the PIPTA \mathcal{PIP} in Fig. 7a. Then $\text{CombFS}(\mathcal{PIP})$ contains the 2 PIPTAs in Fig. 7a (i.e., itself) and in Fig. 7b.

Given an L/U-PTA \mathcal{A} and a subset G of its locations, let us denote by $\text{EFuniv}(\mathcal{A}, G)$ the result of EF-universality for locations G in \mathcal{A} , i.e., the answer to the following question: “is the set of valuations v such that at least one location of G is reachable in $v(\mathcal{A})$ universal?” Or put differently, do *all* valuations reach at least one location of G ? Recall that EF-universality is decidable for L/U-PTAs as shown in André [3].

We need two additional notations before introducing our decision procedure. First, given a PIPTA \mathcal{PIP} , let $\text{makeNonDet}(\mathcal{PIP})$ denote the PTA obtained from \mathcal{PIP} by performing the following three operations:

1. for each location l , for each inconsistent edge Υ from l , add a new non-probabilistic transition from l to a fresh location, with the same guard as on Υ ;
2. remove all inconsistent edges;
3. replace all probabilistic distributions with non-determinism (see e.g., André et al. [5]).

Observe that, if \mathcal{PIP} is an L/U-PIPTA, then $\text{makeNonDet}(\mathcal{PIP})$ is an L/U-PTA. Second, $\text{makeAcc}(\mathcal{PIP})$ returns the set of locations added at step 1.

Beyond transforming the L/U-PIPTA into a non-probabilistic L/U-PTA, the rationale behind makeNonDet is that we need to test for reachability of an *edge*, which is not possible natively in L/U-PTAs; therefore, we add new locations target of these edges. In addition, we cannot test for the reachability of an arbitrary existing location target of these edges, as they may be reached via other paths too. In Fig. 7a, l_3 is reachable via the inconsistent edge outgoing from l_1 , but also directly from l_0 : only reaching l_3 via l_1 should be avoided, which justifies the creation of l'_1 in Fig. 7c.

Example 13. Consider again the PIPTA \mathcal{PIP} in Fig. 7a. Then the L/U-PTA result of $\text{makeNonDet}(\mathcal{PIP})$ is given in Fig. 7c, while $\text{makeAcc}(\mathcal{PIP})$ is $\{l'_1, l'_2\}$.

We can now give below the main equation to solve consistency-emptiness for L/U-PIPTAs.

$$\bigwedge_{\mathcal{P}IP' \in \text{CombFS}(\mathcal{P}IP)} \text{EFuniv}(\text{makeNonDet}(\mathcal{P}IP'), \text{makeAcc}(\mathcal{P}IP')) \quad (1)$$

The idea is that consistency-emptiness holds for an L/U-PIPTA if, for each combination of probabilities set to 0 (*CombFS*), for all parameter valuations (*EFuniv*), some of the locations target of an inconsistent distribution (*makeAcc*) are always reachable. In other words, there is no way to set some probabilities to 0 and to exhibit some parameter valuations that would avoid an inconsistent distribution.

Note that this procedure can be easily implemented by enumerating all PIPTAs in *CombFS*(*PIPT*), replacing probabilities with non-determinism as in *makeNonDet*, and testing *EFuniv* on each resulting L/U-PTA using the procedures given in Bozzelli and La Torre [13], André [3].

Example 14. Consider again the PIPTA *PIPT* in Fig. 7a. Recall that *CombFS*(*PIPT*) is given in Figs. 7a and 7b, and *makeAcc*(*PIPT*) = {*l*'₁, *l*'₂}. Therefore, checking consistency-emptiness for *PIPT* amounts to checking EF-universality of locations {*l*'₁, *l*'₂} in the L/U-PTAs in Figs. 7c and 7d. For the L/U-PTA in Fig. 7c, *EFuniv* gives true, as *l*'₂ can be reached for any valuation of γ_2 and regardless of γ_1 (it suffices to wait enough time in *l*₂ so that the guard $x \geq \gamma_2$ becomes enabled). However, for the L/U-PTA in Fig. 7d, *EFuniv* gives false: indeed, while *l*'₂ is clearly unreachable, *l*'₁ can only be reached if $\gamma_1 > 2$. Therefore, there exist valuations (typically $\gamma_1 \in [0, 2]$) for which locations {*l*'₁, *l*'₂} are unreachable.

In fact, it can be shown that, for the PIPTA *PIPT* in Fig. 7a, the set of valuations for which the PIPTA is consistent is $\gamma_1 \in [0, 2] \wedge \gamma_2 \geq 0$. The idea is to disable the transition to *l*₂ using probabilities (i.e., assigning 1 to *l*₁ and 0 to *l*₂), and to disable the transitions to *l*₃ and *l*₄ by tuning γ_1 .

Theorem 3. *The consistency-emptiness for L/U-PIPTAs is decidable.*

Proof. Given an L/U-PIPTA *PIPT*, we show that Equation (1) holds iff the consistency-emptiness holds for *PIPT*, i.e., no parameter valuation v is such that $v(\mathcal{P}IP)$ is consistent.

⇒ Assume Equation (1) holds. Then EF-universality is true for all possible combinations of probabilities set to 0 (given by *CombFS*(*PIPT*)). That is, for each of these potentially consistent models, for any valuation v , it is always possible to reach at least one of the new locations added by *makeNonDet*, and therefore one of the inconsistent edges in the original model. Therefore, from Definition 6, $v(\mathcal{P}IP)$ is inconsistent for all v . Therefore the consistency-emptiness holds for *PIPT*.

⇐ Assume consistency-emptiness holds for *PIPT*, i.e., no parameter valuation v is such that $v(\mathcal{P}IP)$ is consistent. Then there is no way to tune the probabilities and to tune the timing parameters to avoid the inconsistent edges, and therefore to avoid the new locations added by *makeNonDet*. Then for any $\mathcal{P}IP' \in \text{CombFS}(\mathcal{P}IP)$, we have that

$$\text{EFuniv}(\text{makeNonDet}(\mathcal{P}IP'), \text{makeAcc}(\mathcal{P}IP'))$$

holds. Then Equation (1) holds.

The result then follows from the decidability of the EF-universality problem for L/U-PTAs proved in André [3]. □

The decidability of the emptiness in Theorem 3 does not necessarily mean that the exact synthesis can be achieved. In fact, we show in the following result that the consistency-synthesis for L/U-PIPTAs is intractable in practice, as the set of valuations cannot be represented using, e.g., a finite union of polyhedra.

Proposition 3. *The result of the consistency-synthesis for L/U-PIPTAs cannot be represented using any formalism for which the emptiness of the intersection is decidable.*

Proof. We adapt a reasoning from Jovanović et al. [24] originally showing that the synthesis for L/U-PTAs is intractable. Let us assume an arbitrary PIPTA (not necessarily L/U). For each parameter γ_i , let us create a parameter γ_i^l and a parameter γ_i^u (and delete γ_i). Then, let us replace each constraint $x \leq \gamma_i$ with $x \leq \gamma_i^u$, each constraint $x < \gamma_i$ with $x < \gamma_i^u$, each constraint $x \geq \gamma_i$ with $x \geq \gamma_i^l$, each constraint $x > \gamma_i$ with $x > \gamma_i^l$, and each constraint $x = \gamma_i$ with $x \geq \gamma_i^l \wedge x \leq \gamma_i^u$. We obtain an L/U-PIPTA. Clearly, if $\gamma_i^l = \gamma_i^u$ for all i , then the behavior of the L/U-PIPTA is identical to that of the original PIPTA.

Now, assume that the result of the consistency-synthesis for L/U-PIPTAs can be represented using a formalism for which the emptiness of the intersection is decidable. We can therefore synthesize all valuations for which the L/U-PIPTAs is consistent using such a formalism. Then, let us intersect this result with $\bigwedge_{1 \leq i \leq M} \gamma_i^l = \gamma_i^u$. Finally, let us check whether this intersection is empty. We are thus able to test the consistency-emptiness of the original PIPTA—which contradicts Theorem 2. □

Table 2
Description of the states in Fig. 3b.

State	Location	C	$C \downarrow_{\Gamma}$
s_0	l_0	$x = y \wedge x \geq 0 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_1	l_1	$0 \leq x - y < 2 \wedge y \geq 0 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_2	l_2	$0 \leq y - x < 2 \wedge x \geq 0 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_3	l_3	$2 \leq x - y \leq \gamma \wedge y \geq 0$	$\gamma \geq 2$
s_4	l_4	$x = y \wedge x \geq 0 \wedge \gamma \geq 2$	$\gamma \geq 2$
s_5	l_5	$0 \leq y - x \leq 1 \wedge x \geq 1 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_6	l_2	$1 \leq y - x \leq 2 \wedge x \geq 0 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_7	l_5	$y \geq 2 \wedge y = x + 1 \wedge \gamma \geq 0$	$\gamma \geq 0$
s_8	l_2	$y \geq 2 \wedge y = x + 2 \wedge \gamma \geq 0$	$\gamma \geq 0$

In the rest of the section, despite the negative results of Theorem 2 and Proposition 3, we will still attempt to address synthesis for the full class of PIPTAs.

4.2. A symbolic semantics for PIPTAs

We equip PIPTAs with a symbolic semantics, defined below. Basically, it is inline with the symbolic semantics defined for parametric timed automata (see e.g., André et al. [4], Jovanović et al. [24]), with the addition of probabilistic intervals on the edges; as a consequence, the semantics becomes not an LTS, but an IMDP. Remark that this is a conservative extension of the symbolic semantics of IPTA presented in Definition 9.

Definition 14 (Symbolic semantics of a PIPTA). Given a PIPTA $\mathcal{PIP} = (\Sigma, L, l_0, X, \Gamma, \mathbb{I})$, the symbolic semantics of \mathcal{PIP} is given by the IMDP (S, s_0, \mathbb{I}, T) , with

- $S = \{(l, C) \in L \times \mathcal{Z}\}$, $s_0 = (l_0, (\bigwedge_{1 \leq i \leq H} x_i = 0)^{\nearrow})$,
- $((l, C), e, \Upsilon') \in T$ if there exists $e = (l, g, a, \Upsilon) \in \mathbb{I}$ such that for all $l' \in L$, for all $\rho \subseteq X$ such that $\Upsilon(\rho, l') > 0$, $C' = [C \wedge g]_{\rho}^{\nearrow}$, and $\Upsilon'((l', C')) = \Upsilon(\rho, l')$.

Observe that, whenever a PIPTA has no probabilistic choice (i.e., is a PTA), then the IMDP becomes a labeled transition system, and the symbolic semantics matches that of parametric timed automata. We refer to the symbolic semantics of \mathcal{PIP} as the *parametric probabilistic zone graph* of \mathcal{PIP} .

Just as in parametric timed automata, the number of symbolic states in a PIPTA can be infinite in general.

In parametric timed automata, the *reachability condition* is the projection onto the parameters of a parametric zone (see Jovanović et al. [24]). It is well-known that, given a symbolic run of a parametric timed automaton leading to a symbolic state (l, C) , there exists an equivalent concrete run iff $\gamma \models C \downarrow_{\Gamma}$ (see e.g., Hune et al. [21]). Since our definition of zones matches that of Hune et al. [21], this result extends to PIPTAs in a straightforward manner.

Lemma 3. Let \mathcal{PIP} be a PIPTA. Consider a run in the parametric probabilistic zone graph of \mathcal{PIP} reaching state (l, C) . Let v be a parameter valuation. Then, there exists an equivalent run in $v(\mathcal{PIP})$ iff $v \models C \downarrow_{\Gamma}$.

By equivalent run, we mean (just as for parametric timed automata) an identical discrete structure (locations and edges).

Example 15. The parametric probabilistic zone graph of the PIPTA in Fig. 2b is the IMDP given in Fig. 3b. The symbolic states $s_i = (l_i, C_i)$ are expanded in Table 2. In addition, we also give the reachability condition of each state, i.e., the projection onto the parameters of the zone ($C \downarrow_{\Gamma}$).

4.3. A construction for consistency-synthesis for PIPTAs

Unlike for IPTAs/IMDPs where inconsistent states can only be avoided by enforcing their incoming probabilities to 0, there are two ways of avoiding inconsistent states in PIPTAs. Indeed, while imposing a 0 probability to all transitions going to inconsistent states is a safe choice, it is also possible to avoid inconsistent states by cleverly choosing parameter values such that the guards of transitions potentially going to these states are never satisfied.

The construction we propose for synthesizing parameter valuations ensuring consistency of a given PIPTA is based on the following observation: Since parameters only occur in transition guards, the choice of parameter values cannot interfere with the choice of probability distributions matching (or not) the specified intervals. That comes from the fact that, given a state s , all successors of this state via a given transition have the same parameter constraint (this would not hold with invariants). As a consequence, states that can be made unreachable through probabilistic choice can be made so regardless of the choice of parameter values.

Notations. We first introduce a few notations to make our construction more compact. Given an IMDP $\mathcal{IM} = (S, s_0, \mathbb{I}, T)$ (representing the semantics of a PIPTA \mathcal{PIP}), let $T_{out}(\mathbf{s})$ denote the set of transitions of source \mathbf{s} , i.e., $T_{out}(\mathbf{s}) = \{(\mathbf{s}, e, I) \in T\}$.

Given a transition $(\mathbf{s}, e, I) \in T$, we may want to forbid this transition; recall that the guard (in the original PIPTA) is the same for all targets, as there is a single guard per interval distribution. As we have no invariants, all target states of a given transition have the same reachability condition (i.e., $C' \downarrow_\Gamma$, for a target $\mathbf{s}' = (I', C')$). Therefore, in order to forbid a transition, it suffices to negate the reachability condition of any of the target states of this transition. Let $ForbidD(I)$ denote this result, i.e., $ForbidD(I) = \neg C' \downarrow_\Gamma$, where (I', C') is an (arbitrary) target state of I .

Finally, recall that a disjunction over an empty set of clauses is by definition false. Therefore, we use $\bigvee_i K_i$ to denote the union over a set that returns the usual union of K_i for all i in the set if the set is non-empty, or \perp if the set is empty. Similarly, $\bigwedge_i K_i$ denotes the intersection over a set that returns the usual intersection of K_i for all i in the set if the set is non-empty, or \top if the set is empty.

Example 16. Consider the IMDP in Fig. 3b, which is the zone graph of the PIPTA from Fig. 2b. Recall that the description of the symbolic states of the IMDP from Fig. 3b is given in Table 2. We illustrate the constructions for T_{out} and $ForbidD$ given above.

Clearly, there is only one outgoing transition from state \mathbf{s}_1 , which is labeled with e_2 . As a consequence, we have $T_{out}(\mathbf{s}_1) = \{(\mathbf{s}_1, e_2, I)\}$ with $I(\mathbf{s}_3) = [0, 0.2]$, $I(\mathbf{s}_4) = [0, 0.3]$, and $I(\mathbf{s}_i) = [0, 0]$ for $i \notin \{3, 4\}$.

Remark that, as explained above, all the states that are reachable through I have the same reachability condition (given in Table 2). As a consequence, we have $ForbidD(I) = \neg(\gamma \geq 2) \equiv \gamma < 2$.

We now propose a characterization of the set of parameter valuations that ensure consistency of a given PIPTA under the assumption that its parametric probabilistic zone graph is finite.

Let \mathcal{PIP} be a PIPTA, and let \mathcal{IM} be its parametric probabilistic zone graph. Assume \mathcal{IM} is finite with state space $\mathbf{S} = \{\mathbf{s}_0, \dots, \mathbf{s}_n\}$. Consider the formula $cons(c_{s_0}, \dots, c_{s_n})$ defined as:

$$(c_{s_0} = \top) \wedge \bigwedge_{\mathbf{s} \in \mathbf{S}} \bigwedge_{(\mathbf{s}, e, I) \in T_{out}(\mathbf{s})} \left(\neg c_{\mathbf{s}} \vee ForbidD(I) \vee \bigvee_{S' \in FS(I)} \bigwedge_{\mathbf{s}' \in S' \setminus \{\mathbf{s}\}} c_{\mathbf{s}'} \right).$$

Intuitively in this formula the variable $c_{\mathbf{s}}$ represents whether state \mathbf{s} can be reachable in an implementation. Recall that this can only be true if \mathbf{s} is consistent. As a consequence, the formula can only be true when the valuation of the parameters is coherent with the consistent states. Indeed, this formula ensures that the initial state is reachable and that, for any state \mathbf{s} and any outgoing transition of this state, either:

- the source state \mathbf{s} is not reachable ($\neg c_{\mathbf{s}}$), or
- the transition is disabled due to the valuation of the parameters ($ForbidD(I)$), or
- the transition is enabled and thus there must exist a feasible support for which all reachable states are also consistent.

The set of all solutions for the consistency synthesis problem is thus given as the set of solutions (in terms of parameter valuations) of the equation:

$$\bigvee_{(c_{s_0}, \dots, c_{s_n}) \in \{\top, \perp\}^{n+1}} cons(c_{s_0}, \dots, c_{s_n}) \quad (2)$$

In the following, this procedure (i.e., solving equation (2)) is called ConstSynth. The intuition behind procedure ConstSynth is that we “guess” the states that will be present in the implementation through the first disjunction (states for which $c_{\mathbf{s}} = \top$), and then verify using $cons(c_{s_0}, \dots, c_{s_n})$ that the resulting implementation is well-defined.

Obviously, the empty set of parameter valuations is always a solution to equation (2). Indeed, in this case, one can set $c_{s_0} = \top$ and $c_{\mathbf{s}} = \perp$ for all $\mathbf{s} \neq \mathbf{s}_0$ and then $ForbidD(I)$ is true for all outgoing transitions of \mathbf{s}_0 . If this is the only solution, then the PIPTA \mathcal{PIP} is inconsistent. Otherwise, \mathcal{PIP} is consistent.

We first illustrate our construction on an example and then show that it is sound and complete when the parametric probabilistic zone graph of \mathcal{PIP} is finite.

Example 17. We now apply our construction to the IMDP from Fig. 3b. Recall that parameters are non-negative, therefore $\gamma < 0 \equiv \perp$. First observe that either \mathbf{s}_1 has to be non-reachable ($c_{s_1} = \perp$) or its outgoing transition needs to be forbidden, because there is no set $S' \in FS(I)$. As a consequence, we obtain the following constraint: $\neg c_{s_1} \vee (\gamma < 2)$. There are no constraints for states $\mathbf{s}_5, \mathbf{s}_7$ and \mathbf{s}_8 as they have no outgoing transitions. For state \mathbf{s}_6 , we have the following constraint:

$$\neg c_{s_6} \vee (\gamma < 0) \vee c_{s_7} \vee (c_{s_7} \wedge c_{s_8}) \equiv \neg c_{s_6} \vee (\gamma < 0) \vee c_{s_7} \equiv \neg c_{s_6} \vee c_{s_7}$$

Similarly, for state s_2 , we obtain:

$$\neg c_{s_2} \vee (\gamma < 0) \vee c_{s_5} \equiv \neg c_{s_2} \vee c_{s_5}$$

Finally, state s_0 yields the following:

$$\neg c_{s_0} \vee (\gamma < 0) \vee c_{s_1} \equiv \neg c_{s_0} \vee c_{s_1}$$

Clearly, when put together in equation (2), we obtain the following (after simplifications):

$$\bigvee_{(c_{s_0}, \dots, c_{s_n}) \in \{\top, \perp\}^{n+1}} (c_{s_0}) \wedge (c_{s_1}) \wedge (\gamma < 2) \wedge (\neg c_{s_2} \vee c_{s_5}) \wedge (\neg c_{s_6} \vee c_{s_7})$$

The solutions are therefore all parameter valuations such that $\gamma < 2$, and can be obtained for all assignments of c_s such that $c_{s_0} = c_{s_1} = \top$, $c_{s_2} \Rightarrow c_{s_5}$ and $c_{s_6} \Rightarrow c_{s_7}$.

We now prove that our construction is indeed correct whenever the parametric probabilistic zone graph of the given PIPTA \mathcal{PIP} is finite.

Proposition 4 (Correctness). *Let \mathcal{PIP} be a PIPTA, and let \mathcal{IM} be its parametric probabilistic zone graph. Assume \mathcal{IM} is finite. Assume that the set of parameter valuations satisfying equation (2) is not empty and let v be such a parameter valuation.*

Then $v(\mathcal{PIP})$ is consistent.

Proof. Let v be a solution of equation (2). As a consequence, there must exist an assignment $\kappa_{s_0}, \dots, \kappa_{s_n}$ of the variables c_s such that $v(\text{cons}(\kappa_{s_0}, \dots, \kappa_{s_n}))$. Moreover, for each state s such that $\kappa_s = \top$, the following equation is satisfied:

$$\bigwedge_{(s, e, I) \in T_{\text{out}}(s)} \left(v(\text{ForbidD}(I)) \vee \bigvee_{S' \in FS(I)} \bigwedge_{s' \in S' \setminus \{s\}} \kappa_{s'} \right)$$

Therefore, for each $(s, e, I) \in T_{\text{out}}(s)$, either $v(\text{ForbidD}(I))$ is true (in this case the transition cannot be taken due to timing parameters and is therefore absent from $v(\mathcal{IM})$), or there exists a distribution ι matching I such that all states s' such that $\iota(s') > 0$ are such that $\kappa_{s'} = \top$.

We can therefore construct an MDP whose states are exactly the states s such that $\kappa_s = \top$, and whose transitions are given the distributions ι defined above, that clearly satisfies the IMDP $v(\mathcal{IM})$.

By Proposition 2, we can therefore conclude that \mathcal{PIP} is consistent. \square

We now show that our construction is complete whenever the parametric probabilistic zone graph of the given PIPTA \mathcal{PIP} is finite.

Proposition 5 (Completeness). *Let \mathcal{PIP} be a PIPTA, and let \mathcal{IM} be its parametric probabilistic zone graph. Assume \mathcal{IM} is finite. Let v be such that $v(\mathcal{PIP})$ is consistent. Then v is a solution of equation (2).*

Proof. Since $v(\mathcal{PIP})$ is consistent, there must exist, by Proposition 2 and Lemma 1, an MDP \mathcal{M} with the same structure as $v(\mathcal{IM})$ that satisfies $v(\mathcal{IM})$.

We now propose a valuation of the variables c_s and show that, for this valuation, equation (2) is true.

For all state s in \mathcal{IM} , let $c_s = \top$ if s is reachable (and present) in \mathcal{M} , and \perp otherwise. We now show that the equation $v(\text{cons}(c_{s_0}, \dots, c_{s_n}))$ is true. Clearly, we have $c_{s_0} = \top$, so we just have to show that for all state s in \mathcal{IM} ,

$$\bigwedge_{(s, e, I) \in T_{\text{out}}(s)} \left(\neg c_s \vee v(\text{ForbidD}(I)) \vee \bigvee_{S' \in FS(I)} \bigwedge_{s' \in S' \setminus \{s\}} c_{s'} \right) = \text{true}$$

If s is such that $c_s = \perp$, then this is trivial. Otherwise, let $(s, e, I) \in T_{\text{out}}(s)$. Clearly, since the transition is present in $v(\mathcal{IM})$, we have $v(\text{ForbidD}(I)) = \text{false}$. Moreover, since s is present (and reachable) in \mathcal{M} , there is a transition (s, e, ι) in \mathcal{M} such that $\iota \preceq_{\mathcal{R}_M} I$ for the witnessing relation \mathcal{R}_M . As a consequence, the set $S' = \{s' \mid \iota(s') > 0\}$ is such that $S' \in FS(I)$. Moreover, all states in S' are reachable by construction, thus $s' \in S' \Rightarrow c_{s'} = \top$. Therefore, $v(\text{cons}(c_{s_0}, \dots, c_{s_n}))$ is true. \square

Remark 2. Our construction is based on the parametric probabilistic zone graph. It is sound and complete when this zone graph is finite. However, the resulting equation contains infinite conjunctions and disjunctions when the parametric probabilistic zone graph is infinite, rendering it useless in practice in this case.

However, in practice, one could truncate the parametric probabilistic zone graph up to a certain depth, which would allow computing an approximation of the set of parameter valuations ensuring consistency.




Fig. 8. A PIPTA for which no valuation allows for consistent reachability of l_1 .

4.4. Parametric consistent reachability

A model that is inconsistent is a model that can be considered as ill-formed; therefore, synthesizing valuations for a model to be consistent is an important problem. However, it may not be seen as the final problem a system designer aims at solving. More common problems are reachability, safety, unavoidability, or more complex properties expressed, e.g., using logic formulas.

In this section, we illustrate how consistency synthesis can be combined with existing synthesis algorithms. As a proof of concept, we consider the following parametric consistent reachability synthesis problem:

parametric consistent reachability synthesis problem:

INPUT: A PIPTA \mathcal{PIP} , a set of goal locations G

PROBLEM: find all parameter valuations v for which $v(\mathcal{PIP})$ is consistent and at least one location in G is reachable in $v(\mathcal{PIP})$.

The corresponding emptiness problem, i.e., the emptiness of the valuation set for which a PIPTA is consistent and at least one goal location is reachable, is clearly undecidable: it suffices to consider a PIPTA with no probabilities. This gives a PTA, for which reachability emptiness is undecidable: so, clearly, a PTA is always consistent and therefore consistent reachability emptiness reduces to reachability emptiness, which is undecidable, as shown in Alur et al. [2].

Still, we will propose a method to perform parametric consistent reachability synthesis for PIPTA; again, this method only works when the parametric probabilistic zone graph is finite.

First, let us rule out the following naive method. We could have considered the PTA obtained from a PIPTA by removing all probabilities, then we could have synthesized valuations for which reachability of location G is ensured, which can be obtained using the algorithm described in, e.g., Jovanović et al. [24], and that we will call EFsynth . This gives a constraint K_{reach} . Then, we could synthesize the constraint K_{cons} obtained from ConstSynth . Finally, we could have considered the intersection $K_{\text{reach}} \wedge K_{\text{cons}}$. However this is not satisfactory (and wrong), as shown in the example below.

Example 18. Consider the PIPTA in Fig. 8. Assume $G = \{l_1\}$. Clearly, l_1 is inconsistent, as its successor has an interval distribution that admits no implementation. l_1 can easily be discarded by assigning it a 0-probability from l_0 while keeping the interval consistent.

On this PIPTA without probabilities, EFsynth will output \top as any parameter valuation may reach l_1 . ConstSynth will also output \top . The intersection gives \top , while the set of valuations for which l_1 is reachable and the system is consistent is empty.

We propose the following construction, which we adapt from our construction ConstSynth . Recall that in this construction, we use for each state \mathbf{s} a variable $c_{\mathbf{s}}$ that encodes the potential presence of state \mathbf{s} in an implementation (and therefore imposes that this state is consistent). Unfortunately, the presence of such a state in an implementation is not sufficient to ensure that this state is reachable from the initial state. In order to guarantee reachability, we therefore have to add variables and constraints to ConstSynth .

We therefore add new variables $r_{\mathbf{s}}$ for all states \mathbf{s} in the parametric probabilistic zone graph. These variables will be assigned values in $[0, N] \cup \{\infty\}$, where N is the total number of states of the parametric probabilistic zone graph.

Then, in order to ensure reachability of the goal locations, we add the following constraints:

- $r_{(l,C)} = 0 \iff l \in G$,
- for all states $\mathbf{s} = (l, C)$ in the parametric probabilistic zone graph such that $l \notin G$, we impose that either $(r_{\mathbf{s}} = \infty)$ or

$$\bigvee_{(\mathbf{s}, e, l) \in T_{\text{out}}(\mathbf{s})} \left(\neg \text{ForbidD}(l) \wedge \bigvee_{\mathbf{s}' \in FS(l)} \left(\bigwedge_{\mathbf{s}' \in \mathbf{s}' \setminus \{\mathbf{s}\}} c_{\mathbf{s}'} \wedge \bigvee_{\mathbf{s}' \in \mathbf{s}'} (r_{\mathbf{s}'} = r_{\mathbf{s}} + 1) \right) \right)$$

Now, solving the conjunction of equation (2) from ConstSynth and the constraints presented above, while imposing that $r_{s_0} < \infty$ will yield exactly the set of parameter valuations ensuring the consistent reachability of goal locations from G . We call this new procedure ConstEFSynth.

Proposition 6. *Let v be a parameter valuation satisfying the result of ConstEFSynth. Then, $v(PIP)$ is consistent and at least one location in G is reachable in $v(PIP)$.*

Proof. First observe that any parameter valuation v obtained through ConstEFSynth needs to satisfy ConstSynth. As a consequence, $v(PIP)$ is consistent. Moreover, the additional constraints provided in ConstEFSynth ensure that whenever $r_s < \infty$, there is an execution of length at most r_s from s to a state (l, C) such that $l \in G$ in the parametric probabilistic zone graph of $v(PIP)$. Since we impose that $r_{s_0} < \infty$, G is indeed reachable in $v(PIP)$. \square

Example 19. Let us come back to Fig. 8. ConstSynth yields the entire set of parameter valuations. By construction, the parametric probabilistic zone graph of this PIPTA is almost identical to the PIPTA itself (the states will be $s_i = (l_i, \gamma \geq 0)$ for all i). However, in ConstEFSynth, it will be impossible to set r_{s_1} to a finite value as the feasible support of its outgoing transition is empty. As a consequence, ConstEFSynth will yield the empty set of parameter valuations, as expected.

Example 20. Assume now that we would like to synthesize the parameter values ensuring the consistent reachability of l_5 in the PIPTA given in Fig. 2b. Recall that the probabilistic zone graph is given in Fig. 3b and the solutions of ConstSynth are given in Example 17. In the process of solving ConstEFSynth on this example, we are required to set $r_{s_i} = \infty$ for $i \in \{1, 3, 4, 8\}$. We also have to set $r_{s_5} = r_{s_7} = 0$. Finally, in addition to the above constraints and those obtained in ConstSynth, ConstEFSynth yields the following:

- For state s_6 : either $(r_{s_6} = \infty)$, or

$$(\gamma \geq 0) \wedge ((c_{s_7} \wedge (r_{s_6} = r_{s_7} + 1)) \vee (c_{s_7} \wedge c_{s_8} \wedge ((r_{s_6} = r_{s_7} + 1) \vee (r_{s_6} = r_{s_8} + 1))))$$

- For state s_2 : either $(r_{s_2} = \infty)$, or

$$(\gamma \geq 0) \wedge ((c_{s_5} \wedge (r_{s_2} = r_{s_5} + 1)) \vee (c_{s_5} \wedge c_{s_6} \wedge ((r_{s_2} = r_{s_5} + 1) \vee (r_{s_2} = r_{s_6} + 1))))$$

- For state s_0 : either $(r_{s_0} = \infty)$, or

$$(\gamma \geq 0) \wedge ((c_{s_1} \wedge (r_{s_0} = r_{s_1} + 1)) \vee (c_{s_1} \wedge c_{s_2} \wedge ((r_{s_0} = r_{s_1} + 1) \vee (r_{s_0} = r_{s_2} + 1))))$$

In the end, we can set $(r_{s_6} = 1)$, $(r_{s_2} = 1)$, and $(r_{s_0} = 2)$ for instance. In this case, we still obtain the same set of parameter valuations as in ConstSynth: all those satisfying $(\gamma < 2)$.

Remark 3. Observe that, for acyclic PIPTAs (i.e., the underlying graph of which contains no cycle), the answer to the parametric consistent reachability synthesis problem can be effectively computed. Indeed, the procedure presented above consists in a procedure to be solved on a set of states. If that set is finite, the procedure can be effectively solved with an exact result.

This result can also be extended to PIPTAs whose symbolic semantics is acyclic (i.e., the underlying IMDP contains no cycle). However, it may not be possible to decide whether an arbitrary PIPTA has a finite symbolic semantics.

5. Conclusion

In this work, we provided abstractions to reason on systems involving real-time constraints and probabilities: first, by allowing probabilities to range in some intervals, and, second, by allowing timing constants to be abstracted in the form of parameters. Without parameters, we proposed an approach to decide whether an interval probabilistic timed automaton is consistent, i.e., admits an implementation based on a simulation relation. When adding parameters, the mere existence of a parameter valuation yielding consistency is undecidable. However, when the set of parameters is partitioned between lower-bound parameters and upper-bound parameters, decidability is ensured. We also proposed a procedure to synthesize valuations ensuring consistency for PIPTAs whose parametric probabilistic zone graph is finite, as well as to ensure consistent reachability.

Discussion. We believe our definition of consistency allows for incremental design: one can first define range for probabilities and range for timing parameters. Then, depending on refined design choices, one will assign interval probabilities with punctual values, and value timing parameters. Clearly, inconsistent probabilistic distributions can be seen as ill-formed

models—just as deadlocks, for examples. One could argue that, contrarily to deadlocks, one could statically detect such situations, or even forbid them statically. However, we see two reasons not to do so. First, we believe that allowing these situations could be used as an additional freedom, that can be then detected and corrected using the methods described in this manuscript. That is, inconsistent intervals do not need to be removed statically if there is another way to remove them (using other probabilities or timing parameters). Second, our work builds on top on works where *parametric* probabilistic bounds can be used (e.g., Delahaye [15], Delahaye et al. [18]). In this latter case, the static detection does not work. As our ultimate goal is to reintroduce parametric intervals in the future (see below), we believe our definition of consistency is worth exploring.

Future works. We envision several future works. First, exhibiting subclasses of PIP-TAs for which exact synthesis can be achieved is on our agenda. As the use of timing parameters seems critical in our undecidability results, relying on recent works exhibiting decidable subclasses of parametric timed automata, such as bounded integer parameters (see Jovanović et al. [24]) or reset-parametric timed automata (see André et al. [8,7]), can serve as a first basis for a probabilistic extension.

Finally, we are interested in considering higher-level abstractions of probabilities; notably, using parameters instead of intervals with constant bounds (as in Delahaye et al. [18] for parametric interval Markov chains) is of high interest, and makes the notion of consistency even more delicate, as tuning the parametric bounds in an interval may impact the consistency of other probabilistic distributions.

Declaration of Competing Interest

We hereby declare that there is no competing interest.

Acknowledgements

We would like to thank the anonymous reviewers for useful comments that helped us to improve the manuscript.

References

- [1] R. Alur, D.L. Dill, A theory of timed automata, *Theor. Comput. Sci.* 126 (2) (Apr. 1994) 183–235.
- [2] R. Alur, T.A. Henzinger, M.Y. Vardi, Parametric real-time reasoning, in: S.R. Kosaraju, D.S. Johnson, A. Aggarwal (Eds.), *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, STOC'93*, ACM, New York, NY, USA, 1993, pp. 592–601.
- [3] É. André, What's decidable about parametric timed automata?, *Int. J. Softw. Tools Technol. Transf.* 21 (2) (2019) 203–219.
- [4] É. André, Th. Chatain, E. Encrenaz, L. Fribourg, An inverse method for parametric timed automata, *Int. J. Found. Comput. Sci.* 20 (5) (2009) 819–836.
- [5] É. André, L. Fribourg, J. Sproston, An extension of the inverse method to probabilistic timed automata, *Form. Methods Syst. Des.* (2) (2013) 119–145.
- [6] É. André, D. Lime, Liveness in L/U-parametric timed automata, in: A. Legay, K. Schneider (Eds.), *ACSD, IEEE*, 2017, pp. 9–18.
- [7] É. André, D. Lime, M. Ramparison, Timed automata with parametric updates, in: *Proceedings of the 18th International Conference on Application of Concurrency to System Design (ACSD'18)*, IEEE, 2018, pp. 21–29.
- [8] É. André, D. Lime, O.H. Roux, Decision problems for parametric timed automata, in: K. Ogata, M. Lawford, S. Liu (Eds.), *Proceedings of the 18th International Conference on Formal Engineering Methods (ICFEM'16)*, in: *Lecture Notes in Computer Science*, vol. 10009, Springer, 2016, pp. 400–416.
- [9] É. André, N. Markey, Language preservation problems in parametric timed automata, in: S. Sankaranarayanan, E. Vicario (Eds.), *FORMATS*, in: *Lecture Notes in Computer Science*, vol. 9268, Springer, Sep. 2015, pp. 27–43.
- [10] G. Behrmann, P. Bouyer, K.G. Larsen, R. Pelánek, Lower and upper bounds in zone-based abstractions of timed automata, *Int. J. Softw. Tools Technol. Transf.* 8 (3) (2006) 204–215.
- [11] N. Beneš, P. Bezděk, K.G. Larsen, J. Srba, Language emptiness of continuous-time parametric timed automata, in: M.M. Halldórsson, K. Iwama, N. Kobayashi, B. Speckmann (Eds.), *ICALP, Part II*, in: *Lecture Notes in Computer Science*, vol. 9135, Springer, Jul. 2015, pp. 69–81.
- [12] J. Bengtsson, W. Yi, Timed automata: semantics, algorithms and tools, in: J. Desel, W. Reisig, G. Rozenberg (Eds.), *Lectures on Concurrency and Petri Nets*, *Advances in Petri Nets*, in: *Lecture Notes in Computer Science*, vol. 3098, Springer, 2003, pp. 87–124.
- [13] L. Bozzelli, S. La Torre, Decision problems for lower/upper bound parametric timed automata, *Form. Methods Syst. Des.* 35 (2) (2009) 121–151.
- [14] A. Collomb-Annichini, M. Sighireanu, Parameterized reachability analysis of the IEEE 1394 root contention protocol using TReX, in: *RT-TOOLS*, 2001.
- [15] B. Delahaye, Consistency for parametric interval Markov chains, in: *2nd International Workshop on Synthesis of Complex Parameters, SynCoP 2015*, April 11, 2015, London, United Kingdom, in: *OASICS*, vol. 44, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2015, pp. 17–32.
- [16] B. Delahaye, J. Katoen, K.G. Larsen, A. Legay, M.L. Pedersen, F. Sher, A. Wasowski, Abstract probabilistic automata, *Inf. Comput.* 232 (2013) 66–116.
- [17] B. Delahaye, K.G. Larsen, A. Legay, M.L. Pedersen, A. Wasowski, Consistency and refinement for interval Markov chains, *J. Log. Algebraic Program.* 81 (3) (2012) 209–226.
- [18] B. Delahaye, D. Lime, L. Petrucci, Parameter synthesis for parametric interval Markov chains, in: *Verification, Model Checking, and Abstract Interpretation - Proceedings of the 17th International Conference, VMCAI 2016*, St. Petersburg, FL, USA, January 17–19, 2016, in: *Lecture Notes in Computer Science*, vol. 9583, Springer, 2016, pp. 372–390.
- [19] L. Doyen, Robust parametric reachability for timed automata, *Inf. Process. Lett.* 102 (5) (2007) 208–213.
- [20] H. Gregersen, H.E. Jensen, Formal Design of Reliable Real Time Systems, Master's thesis, Department of Mathematics and Computer Science, Aalborg University, 1995.
- [21] T. Hune, J. Romijn, M. Stoelinga, F.W. Vaandrager, Linear parametric model checking of timed automata, *J. Log. Algebraic Program.* 52–53 (2002) 183–220.
- [22] B. Jonsson, K. Larsen, Specification and refinement of probabilistic processes, in: *LICS, IEEE*, 1991, pp. 266–277.
- [23] A. Jovanović, M.Z. Kwiatkowska, Parameter synthesis for probabilistic timed automata using stochastic game abstractions, in: J. Ouaknine, I. Potapov, J. Worrell (Eds.), *Proceedings of the 8th International Workshop on Reachability Problems (RP 2014)*, in: *Lecture Notes in Computer Science*, vol. 8762, Springer, 2014, pp. 176–189.
- [24] A. Jovanović, D. Lime, O.H. Roux, Integer parameter synthesis for timed automata, *IEEE Trans. Softw. Eng.* 41 (5) (2015) 445–461.
- [25] M.Z. Kwiatkowska, G. Norman, D. Parker, J. Sproston, Performance analysis of probabilistic timed automata using digital clocks, *Form. Methods Syst. Des.* 29 (1) (2006) 33–78.

- [26] M.Z. Kwiatkowska, G. Norman, R. Segala, J. Sproston, Automatic verification of real-time systems with discrete probability distributions, *Theor. Comput. Sci.* 282 (2002) 101–150.
- [27] M.Z. Kwiatkowska, G. Norman, J. Sproston, Probabilistic model checking of deadline properties in the IEEE 1394 FireWire root contention protocol, *Form. Asp. Comput.* 14 (3) (2003) 295–318.
- [28] J.S. Miller, Decidability and complexity results for timed automata and semi-linear hybrid automata, in: N.A. Lynch, B.H. Krogh (Eds.), HSCC, in: *Lecture Notes in Computer Science*, vol. 1790, Springer, 2000, pp. 296–309.