Quantitative Model-Checking of One-Clock Timed Automata under Probabilistic Semantics

Nathalie Bertrand[1] Patricia Bouyer[2] Thomas Brihaye[3] Nicolas Markey[2]

[1] IRISA, INRIA Rennes, France
nathalie.bertrand@irisa.fr

[2] LSV, ENS Cachan & CNRS, France
{bouyer,markey}@lsv.ens-cachan.fr

[3] Université de Mons-Hainaut, Belgium
thomas.brihaye@umh.ac.be

Abstract

In [3] a probabilistic semantics for timed automata has been defined in order to rule out unlikely (sequences of) events. The qualitative model-checking problem for LTL *properties has been investigated, where the aim is to check whether a given* LTL *property holds with probability* 1 *in a timed automaton, and solved for the class of single-clock timed automata.*

In this paper, we consider the quantitative model-checking problem for ω-regular properties: we aim at computing the exact probability that a given timed automaton satisfies an ω-regular property. We develop a framework in which we can compute a closed-form expression for this probability; we furthermore give an approximation algorithm, and finally prove that we can decide the threshold problem in that framework.

1 Introduction

Timed automata [1] are a well-established formalism for the modelling and analysis of timed systems. A timed automaton is roughly a finite-state automaton enriched with clocks and clock constraints. This model has been extensively studied, and several verification tools have been developed. However, like many models used in model-checking, timed automata are an idealized mathematical model, in which many hypotheses are implicitly made. For instance, a timed automaton can check the values of clocks with an infinite precision, events are instantaneous, *etc*. Recently, a new direction of research has consisted in proposing alternative semantics for timed automata that provide more realistic operational models for real-time systems. We can for instance mention the Almost ASAP semantics (AASAP for short) introduced in [13] and further investigated in [12, 2, 7, 8], which somewhat relaxes constraints on clocks, hence most of the idealization side-effects for timed automata. However, it induces a very strong notion of robustness, suitable for really critical systems, but maybe too strong for less critical systems. Another ‘robust semantics’, based on the notion of tube acceptance, has been proposed in [15, 16]: a metric is put on the set of traces of the timed automaton, and roughly, a trace is robustly accepted if and only if a tube around that trace is classically accepted. This language-focused notion of acceptance is not completely satisfactory because it does not take into account the structure of the automaton, and hence is not related to the most-likely behaviours of the automaton.

In [4, 3], a natural probabilistic semantics has been given to timed automata, which randomizes both delays and choices of transitions, and provides a way of measuring the ‘size’ of sets of behaviours in the timed automaton. That way, we can measure, for instance, how likely a timed automaton satisfies a given LTL property. In those two papers, the *almost-sure* model-checking problem for LTL is investigated,[1] where the probability of satisfying the property is compared to 1. A topological characterization of the almost-sure satisfaction is given, which helps understanding when a timed automaton almost-surely satisfies an LTL property. In [3], the almost-sure model-checking problem is shown decidable for *single-clock* timed automata, and an algorithm based on the construction of a (qualitatively equivalent) finite Markov chain is described. An intriguing two-clock example is presented, for which the above finite Markov chain abstraction is not correct.

In this paper, we investigate the *quantitative* probabilistic model-checking problem, which aims at computing

[1] Note that the work developed in [3] can straightforwardly be extended to the whole class of ω-regular properties.

the probability of a given ω-regular property in a timed automaton. The finite Markov chain abstraction that has been proposed in [3] is no more correct, and new techniques need to be developed. For a subclass of single-clock timed automata[2], we define a new abstraction of the timed automaton, which helps solving the quantitative model-checking problem. Given a timed automaton $\mathcal{A}$ and an ω-regular property φ, this abstraction is a finite Markov chain $\mathcal{M}_{\mathcal{A}}$, and there is a computable reachability property φ' such that the probability that $\mathcal{A}$ satisfies φ coincides with the probability that $\mathcal{M}_{\mathcal{A}}$ satisfies φ'. However this probability can in general not be expressed by a simple closed-form expression, and we provide a concrete framework (where the probability distributions over delays are given by exponential functions), in which we will be able to *(i)* compute a closed-form expression for the probability that $\mathcal{A}$ satisfies φ, *(ii)* approximate this probability, and *(iii)* decide the threshold problem (compare the probability to a given threshold).

The paper is organized as follows: in Section 2, we recall the classical definitions related to timed automata, and introduce the probabilistic semantics we are considering and the associated model-checking problem. In Section 3, we present an abstraction, in the form of a finite Markov chain, which allows to compute abstract expressions for the probabilities of ω-regular properties in single-clock timed automata. In Section 4, we present a restricted framework in which closed-form expressions can be computed for the probabilities of ω-regular properties; we then develop an approximation scheme, and finally prove the decidability of the threshold problem.

2 Definitions

2.1 The timed automaton model

Let X be a finite set of variables, called *clocks*. A *clock valuation* over X is a mapping $\nu\colon X \to \mathbb{R}_+$, where $\mathbb{R}_+$ is the set of nonnegative reals. We write $\mathbb{R}_+^X$ for the set of clock valuations over X. If $\nu \in \mathbb{R}_+^X$ and $\tau \in \mathbb{R}_+$, we write $\nu + \tau$ for the clock valuation defined by $(\nu + \tau)(x) = \nu(x) + \tau$ if $x \in X$. If $Y \subseteq X$, the valuation $[Y \leftarrow 0]\nu$ is the valuation assigning 0 to $x \in Y$ and $\nu(x)$ to $x \notin Y$. A *guard* (or *clock constraint*) over X is a finite conjunction of expressions of the form $x \sim c$ where $x \in X$, $c \in \mathbb{N}$, and $\sim \in \{<, \leq, =, \geq, >\}$. We denote by $\mathcal{G}(X)$ the set of guards over X. The satisfaction relation for guards over clock valuations is defined in a natural way, and we write $\nu \models g$, if ν satisfies g.

Definition 1 *A* timed automaton *is a tuple* $\mathcal{A} = (L, X, E, \mathcal{I})$ *such that:* (i) *L is a finite set of locations,* (ii) *X is a finite set of clocks,* (iii) *$E \subseteq L \times \mathcal{G}(X) \times 2^X \times L$ is a finite set of edges, and* (iv) *$\mathcal{I} : L \to \mathcal{G}(X)$ assigns an invariant to each location.*

The semantics of a timed automaton $\mathcal{A}$ is a timed transition system whose states are pairs $(\ell, \nu) \in L \times \mathbb{R}_+^{|X|}$ with $\nu \models \mathcal{I}(\ell)$, and whose transitions are of the form $(\ell, \nu) \xrightarrow{\tau, e} (\ell', \nu')$ if there exists an edge $e = (\ell, g, Y, \ell')$ such that for every $0 \leq \tau' \leq \tau$, $\nu + \tau' \models \mathcal{I}(\ell)$, $\nu + \tau \models g$, $\nu' = [Y \leftarrow 0](\nu + \tau)$, and $\nu' \models \mathcal{I}(\ell')$. A finite (resp. infinite) *run* ϱ of $\mathcal{A}$ is a finite (resp. infinite) sequence of transitions, *i.e.*, $\varrho = s_0 \xrightarrow{\tau_1, e_1} s_1 \xrightarrow{\tau_2, e_2} s_2 \ldots$ where for each $i \geq 0$, $s_i = (\ell_i, \nu_i)$ is a state. We write $\mathsf{Runs}_f(\mathcal{A}, s_0)$ (resp. $\mathsf{Runs}(\mathcal{A}, s_0)$) for the set of finite runs (resp. infinite runs) of $\mathcal{A}$ from state s_0.

If s is a state of $\mathcal{A}$ and $(e_i)_{1 \leq i \leq n}$ is a finite sequence of edges of $\mathcal{A}$, the *(symbolic) path* starting from s and determined by $(e_i)_{1 \leq i \leq n}$ is the following set of runs:

$$\pi(s, e_1 \ldots e_n) = \{\varrho = s \xrightarrow{\tau_1, e_1} s_1 \ldots \xrightarrow{\tau_n, e_n} s_n \mid \varrho \in \mathsf{Runs}_f(\mathcal{A}, s)\}.$$

Given an n-variable constraint $\mathcal{C}$, the *constrained* symbolic path $\pi_{\mathcal{C}}(s, e_1 \ldots e_n)$ is the subset of $\pi(s, e_1 \ldots e_n)$ where the delays τ_1 to τ_n satisfy the constraint $\mathcal{C}$. Let $\pi = \pi(s, e_1 \ldots e_n)$ be a finite symbolic path, we define the *cylinder* generated by π as

$$\mathsf{Cyl}(\pi) = \{\varrho \in \mathsf{Runs}(\mathcal{A}, s) \mid \exists \varrho' \in \mathsf{Runs}_f(\mathcal{A}, s), \text{ finite prefix of } \varrho, \text{ s.t. } \varrho' \in \pi\}$$

Also, given a state s of $\mathcal{A}$ and an infinite sequence of edges $(e_i)_{i \geq 1}$, we will need the notion of infinite symbolic paths defined as:

$$\pi(s, e_1 \ldots) = \{\varrho = s \xrightarrow{\tau_1, e_1} s_1 \ldots \mid \varrho \in \mathsf{Runs}(\mathcal{A}, s)\}.$$

Given a state s of $\mathcal{A}$ and an edge e, we define $I(s, e) = \{\tau \in \mathbb{R}_+ \mid s \xrightarrow{\tau, e} s'\}$ and $I(s) = \bigcup_e I(s, e)$. The automaton $\mathcal{A}$ is *non-blocking* if, for every state s, $I(s) \neq \emptyset$.

2.2 The region automaton abstraction

The well-known region automaton construction is a finite abstraction of timed automata which can be used for verifying many properties like ω-regular untimed properties [1]. For lack of space, we do not redefine the region equivalence relation, and we write $R_{\mathcal{A}}$ for the set of (clock) regions of automaton $\mathcal{A}$. Here we use a slight modification of the original construction, which is still a timed automaton.

If $\mathcal{A} = (L, X, E, \mathcal{I})$ is a timed automaton then the *region automaton* of $\mathcal{A}$ is the timed automaton $\mathsf{R}(\mathcal{A}) = (Q, X, T, \kappa)$ such that $Q = L \times R_{\mathcal{A}}$ and:

[2] Due to the results of [3], the restriction to timed automata with one clock seems necessary.

- $\kappa((\ell, r)) = \mathcal{I}(\ell)$ for all $(\ell, r) \in Q$;
- $T \subseteq (Q \times \mathsf{cell}(R_{\mathcal{A}}) \times E \times 2^X \times Q)$, and $(\ell, r) \xrightarrow{\mathsf{cell}(r''),e,Y} (\ell', r')$ is in T iff there exists $e = \ell \xrightarrow{g,Y} \ell'$ in E s.t. there exist $\nu \in r$, $\tau \in \mathbb{R}_+$ with $(\ell, \nu) \xrightarrow{\tau,e} (\ell', \nu')$, $\nu + \tau \in r''$ and $\nu' \in r'$ ($\mathsf{cell}(r'')$ is the tightest guard containing r'').

We recover the usual region automaton of [1] by labelling the transitions with 'e' instead of '$\mathsf{cell}(r''), e, Y$', and by interpreting $\mathsf{R}(\mathcal{A})$ as a finite automaton. The above timed interpretation satisfies strong timed bisimulation properties that we do not detail here. To every finite path $\pi((\ell, \nu), e_1 \dots e_n)$ in $\mathcal{A}$ corresponds a finite set of paths $\pi(((\ell, [\nu]), \nu), f_1 \dots f_n)$ in $\mathsf{R}(\mathcal{A})$,[3] each one corresponding to a choice in the regions that are visited. If ϱ is a run in $\mathcal{A}$, we denote $\iota(\varrho)$ its unique image in $\mathsf{R}(\mathcal{A})$. Note that if $\mathcal{A}$ is non-blocking, then so is $\mathsf{R}(\mathcal{A})$.

In the rest of the paper we assume, following [3], that timed automata are non-blocking.

2.3 The probabilistic semantics

Following [3], we define a probability measure over sets of infinite runs of timed automata, which measures in some sense their *likelihood*. Let $\mathcal{A}$ be a timed automaton. We assume probability distributions are given from every state s of $\mathcal{A}$ both over delays and over enabled moves. For every state s of $\mathcal{A}$, the probability measure μ_s over delays in $\mathbb{R}_+$ (equipped with the standard Borel σ-algebra) must satisfy several requirements:

- $\mu_s(I(s)) = \mu_s(\mathbb{R}_+) = 1$,[4]
- Denoting λ the Lebesgue measure, if $\lambda(I(s)) > 0$, μ_s is equivalent[5] to λ on $I(s)$; Otherwise, μ_s is equivalent on $I(s)$ to the uniform distribution over points of $I(s)$. We recall that the Lebesgue measure is the unique measure that assigns $b - a$ to the interval (a, b).
- We also assume technical hypotheses which we do not detail here (see [3] for details) but that are natural and satisfied in all our further developments.

The second condition is a fairness condition w.r.t. enabled transitions, in that we cannot disallow one transition by assigning probability 0 to delays enabling that transition.

Example 2 *Examples of possible distributions are uniform (resp. exponential) distributions over bounded (resp. unbounded) intervals.*

For every state s of $\mathcal{A}$, we also assume a probability distribution p_s over edges, such that for every edge e, we have $p_s(e) > 0$ iff e is enabled in s. Moreover, for the sake of simplicity, we assume that p_s is given by weights on transitions, as it is classically done for resolving non-determinism: we associate with each edge e a weight $w(e) > 0$, and for every state s and every edge e, $p_s(e) = 0$ if e is not enabled in s, and $p_s(e) = w(e)/(\sum_{e' \text{ enabled in } s} w(e'))$ otherwise. As a consequence, if s and s' are region equivalent, then for every edge e, $p_s(e) = p_{s'}(e)$. We then inductively define a measure over finite symbolic paths from state s as

$$\mathbb{P}_{\mathcal{A}}(\pi(s, e_1 \dots e_n)) = \int_{t \in I(s,e_1)} p_{s+t}(e_1)\, \mathbb{P}_{\mathcal{A}}(\pi(s_t, e_2 \dots e_n))\, \mathrm{d}\mu_s(t)$$

where $s \xrightarrow{t} (s + t) \xrightarrow{e_1} s_t$, and we initialize with $\mathbb{P}_{\mathcal{A}}(\pi(s)) = 1$. The formula for $\mathbb{P}_{\mathcal{A}}$ relies on the fact that the probability of taking transition e_1 at time t coincides with the probability of waiting t time units and then choosing e_1 among the enabled transitions, *i.e.*, $p_{s+t}(e_1)\mathrm{d}\mu_s(t)$. Note that time passage and actions are independent events.

The value $\mathbb{P}_{\mathcal{A}}(\pi(s, e_1 \dots e_n))$ is the result of n successive one-dimensional integrals, but it can also be viewed as the result of an n-dimensional integral. Hence, we can easily extend the above definition to finite constrained paths $\pi_{\mathcal{C}}(s, e_1 \dots e_n)$ when $\mathcal{C}$ is Borel-measurable. This extension to constrained paths will allow to express (and later, measure) various and rather complex sets of paths, for instance Zeno runs.[6] The measure $\mathbb{P}_{\mathcal{A}}$ can then be defined on cylinders, letting $\mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi)) = \mathbb{P}_{\mathcal{A}}(\pi)$ if π is a finite (constrained) symbolic path. Finally we extend $\mathbb{P}_{\mathcal{A}}$ in a standard and unique way to the σ-algebra generated by these cylinders, which we note $\Omega^s_{\mathcal{A}}$ (see [3] for details).

Proposition 3 ([3]) *Let $\mathcal{A}$ be a timed automaton. For every state s, the function $\mathbb{P}_{\mathcal{A}}$ is a probability measure over $(\mathit{Runs}(\mathcal{A}, s), \Omega^s_{\mathcal{A}})$.*

For instance, the set $\mathsf{Zeno}(s)$ of all the Zeno runs starting from s belongs to $\Omega^s_{\mathcal{A}}$. Indeed, it can be defined as:

$$\bigcup_{M \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} \bigcup_{(e_1, \dots, e_n) \in E^n} \mathsf{Cyl}(\pi_{\sum_{i \leq n} \tau_i \leq M}(s, e_1 \dots e_n))$$

Example 4 *Consider the timed automaton $\mathcal{A}$ depicted on Fig. 1, and assume that we assign the uniform distribution over delays to all locations except ℓ_1 and over discrete moves, and that we put the distribution with density function $t \mapsto e^{-t}$ over $\mathbb{R}_+$ in ℓ_1. If $s_0 = (\ell_0, 0)$ is the initial*

[3] Where $[\nu]$ is the unique region in $R_{\mathcal{A}}$ containing valuation ν.

[4] Note that this is possible, as we assume $\mathcal{A}$ is non-blocking, hence $I(s) \neq \emptyset$ for every state s of $\mathcal{A}$.

[5] Two measures ν and ν' are *equivalent* whenever for each measurable set A, $\nu(A) = 0 \Leftrightarrow \nu'(A) = 0$.

[6] An infinite run ϱ: $s_0 \xrightarrow{\tau_1,e_1} s_1 \xrightarrow{\tau_2,e_2} \cdots$ is said *Zeno* whenever $\sum_{i=1}^{\infty} \tau_i$ is bounded.

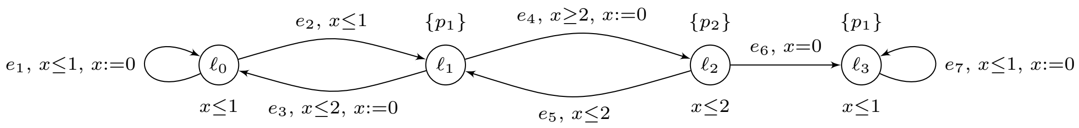

Fig. 1. A running example

state, then

$$\mathbb{P}_{\mathcal{A}}(\mathit{Cyl}(\pi(s_0, e_2 e_3))) = \frac{1}{2} \cdot (1 - e^{-1} + e^{-2})$$

In [3], it is explained how to transfer probabilities from $\mathcal{A}$ to $\mathsf{R}(\mathcal{A})$, thus allowing to prove results on $\mathsf{R}(\mathcal{A})$ and to recover them on the original automaton $\mathcal{A}$. We assume that, for every state s in $\mathcal{A}$, $\mu_s^{\mathcal{A}} = \mu_{\iota(s)}^{\mathsf{R}(\mathcal{A})}$, and for every $t \in \mathbb{R}_+$, $p_{s+t}^{\mathcal{A}} = p_{\iota(s)+t}^{\mathsf{R}(\mathcal{A})}$. Under those assumptions, we have the following correctness result.

Lemma 5 ([3]) *Assume measures in $\mathcal{A}$ and in $\mathsf{R}(\mathcal{A})$ are related as above. Then, for every set S of runs in $\mathcal{A}$ we have: $S \in \Omega_{\mathcal{A}}^s$ iff $\iota(S) \in \Omega_{\mathsf{R}(\mathcal{A})}^{\iota(s)}$, and in this case $\mathbb{P}_{\mathcal{A}}(S) = \mathbb{P}_{\mathsf{R}(\mathcal{A})}(\iota(S))$.*

Therefore, in the sequel, we assume w.l.o.g. that timed automata are given as region automata, *i.e.*, $\mathcal{A} = \mathsf{R}(\mathcal{A})$.

2.4 The quantitative model-checking problem

In this paper we consider ω-regular properties. We assume that an ω-regular property is given by a deterministic finite automaton $\mathcal{B}$ with a Streett acceptance condition of the form $\psi_{\mathcal{B}} = \bigwedge_{i=1}^{n}(\Box\Diamond Q_i \Rightarrow \Box\Diamond Q'_i)$, where $(Q_1, Q'_1), \cdots, (Q_n, Q'_n)$ are pairs of subsets of states in $\mathcal{B}$. The linear-time temporal logic LTL [18] defines a subclass of ω-regular properties.

In [3], the *qualitative* LTL *model-checking problem* is investigated: given a timed automaton $\mathcal{A}$ and an LTL formula φ, this problem consists in deciding whether $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) = 1$, *i.e.*, whether the automaton $\mathcal{A}$ almost-surely satisfies the property φ. It has been proved that for single-clock timed automata, under some technical and reasonable conditions on the various distributions, the almost-sure model-checking problem for LTL is decidable, that it does not depend on the distributions that are used in the automaton, and that the introduction of probabilities does not increase the theoretical complexity of the problem (which is PSPACE-complete). Though the results are stated for LTL properties, the decidability result carries over to ω-regular properties. It relies on the construction of a finite Markov chain abstraction, based on the region automaton $\mathsf{R}(\mathcal{A})$, which preserves the qualitative properties of $\mathcal{A}$.

In this paper we consider the *quantitative model-checking problem*: given a single-clock timed automaton $\mathcal{A}$ with initial state s_0 and an ω-regular property φ, we want to compute $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$. Unfortunately, the abstraction developed in [3] for solving the qualitative model-checking problem is of no interest here, as it does not preserve any precise information about the values of the probabilities.

We first notice that, contrary to the qualitative model-checking problem, the answer to the quantitative model-checking problem does depend on the choice of the distributions that are assigned to delays and edges. This is no surprise since it is already the case for finite discrete-time Markov chains. Furthermore, the probabilities that we compute (when we manage to) are not always satisfactory, as they crucially depend on the possible representation and evaluation of non-rational numbers. As a consequence, we also investigate the *approximate model-checking problem*, where, given a positive real ε, we will aim at computing two rationals P_ε^+ and P_ε^- such that:

$$\begin{cases} P_\varepsilon^- \leq \mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) \leq P_\varepsilon^+ \\ P_\varepsilon^+ - P_\varepsilon^- < \varepsilon \end{cases}$$

It is quite natural to consider this approximate variant in our framework since we show that even for reachability properties, the probability isn't rational (and even not algebraic) in general, and hence cannot be represented easily.

Remark 6 *Algorithms for the approximate quantitative model-checking of probabilistic systems have for instance been proposed for infinite-state systems represented as infinite-state discrete Markov chains (e.g. probabilistic lossy channel systems [19] or probabilistic pushdown automata [14]).*

Finally, we also focus on the *threshold problem*, which asks, given a timed automaton $\mathcal{A}$ with its initial state s_0, an ω-regular property φ, and a threshold $\sim c$ with $\sim \in \{<, \leq, =, \geq, >\}$ and $c \in \mathbb{Q}$, whether $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) \sim c$.

For all the problems we consider (quantitative model-checking, approximate quantitative model-checking and threshold problem), following [3], we restrict our study to single-clock timed automata (because the decidability of the *qualitative* model-checking is already an open problem for multi-clock timed automata).

$$\begin{cases} \mathfrak{p}_\ell(x) = 1 & \text{if } \ell \in B \\ \mathfrak{p}_\ell(x) = \displaystyle\sum_{e \text{ resetting edge}} \int_{t\in I((\ell,x),e)} p_{s+t}(e)\cdot \mathfrak{p}_{\ell'}(0)\,\mathrm{d}\mu_s(t) & \text{otherwise} \\ \qquad\qquad + \displaystyle\sum_{e \text{ non resetting edge}} \int_{t\in I((\ell,x),e)} p_{s+t}(e)\cdot \mathfrak{p}_{\ell'}(x+t)\,\mathrm{d}\mu_s(t) & \end{cases}$$

Table 1. Integral equations for reachability properties

2.5 Methodology

We first solve the quantitative model-checking problem for prefix-independent location-based properties.[7] Given a single-clock timed automaton $\mathcal{A}$ and a prefix-independent location-based property φ, the method follows the two steps below:

- we first abstract the timed automaton $\mathcal{A}$ into a finite Markov chain $\mathcal{M}_\mathcal{A}$;
- we then compute in $\mathcal{M}_\mathcal{A}$ the probability of property φ.

Following techniques of Courcoubetis and Yannakakis [11], computing the probability of a prefix-independent property φ in $\mathcal{M}_\mathcal{A}$ amounts to computing the probability of reaching the BSCCs of $\mathcal{M}_\mathcal{A}$ that are 'good' w.r.t. φ.

The result for general ω-regular properties will then be derived, applying a classical product approach, which we briefly describe now. We assume that φ is an ω-regular property, and we build $\mathcal{B}_\varphi$ a deterministic Streett automaton for φ. Now, given the timed automaton $\mathcal{A}$, we consider the product automaton $\mathcal{A}_\varphi = \mathcal{A} \times \mathcal{B}_\varphi$. Under the assumption that the distributions over delays and actions are naturally transferred from $\mathcal{A}$ to $\mathcal{A}_\varphi$ (*i.e.*, for all state q of $\mathcal{B}_\varphi$, for all state s of $\mathcal{A}$, we set $\mu^{\mathcal{A}_\varphi}_{(s,q)} = \mu^{\mathcal{A}}_s$ and for all edges e, $p^{\mathcal{A}_\varphi}_{(s,q)} = p^{\mathcal{A}}_s$), we have that

$$\mathbb{P}_\mathcal{A}(s_0 \models \varphi) = \mathbb{P}_{\mathcal{A}_\varphi}((s_0, q_0) \models \psi_{\mathcal{B}_\varphi})$$

where q_0 is the initial state of $\mathcal{B}_\varphi$, and $\psi_{\mathcal{B}_\varphi}$ is the acceptance condition induced by automaton $\mathcal{B}_\varphi$.

Hence this product construction allows to lift computability (and decidability) results for prefix-independent location-based properties to general ω-regular properties.

In the sequel, we only consider prefix-independent location-based properties, but all results hold for general ω-regular properties.

Remark 7 *Consider that we want to compute the probability of reaching a set B of locations in $\mathcal{A}$. One is easily convinced that it can be defined by the integral equations of Table 1, where $\mathfrak{p}_\ell(x)$ is the probability of reaching B from (ℓ, x). However, these integral equations can a priori not be solved, in the sense that in general the $\mathfrak{p}_\ell$'s cannot be expressed as functions in closed-form. It is true that most of the time, we will be able to solve numerically this kind of equations, but what we aim at is to obtain closed-form expressions, in order to approximate the values and this way decide the threshold problem, which cannot be done by only applying numerical methods.*

3 Abstraction into a Finite Markov Chain

In this section, we present an abstraction of a timed automaton into a finite Markov chain which we prove is sound and complete for the quantitative model-checking problem for a slight restriction of single-clock timed automata. Let $\mathcal{A} = (L, \{x\}, E, \mathcal{I})$ be a single-clock timed automaton with initial state $s_0 = (\ell_0, 0)$ and assume that M is the maximal constant that appears in a guard of $\mathcal{A}$. W.l.o.g. (thanks to Lemma 5), we assume that $\mathcal{A} = \mathsf{R}(\mathcal{A})$, and we assume moreover that *(i)* if $s = (\ell, \alpha)$ and $s' = (\ell, \alpha')$ are two states s.t. $\alpha, \alpha' > M$, then $\mu_s = \mu_{s'}$, and *(ii)* any bounded cycle of $\mathsf{R}(\mathcal{A})$ contains at least one resetting edge. We write (†) for these restrictions.

Remark 8 *The first restriction is such that it will not be possible to distinguish between region-equivalent states which are in the unbounded region.*

The second restriction (no bounded cycle without reset) is a common assumption when one wants to get rid of some Zeno behaviours. Indeed Alur and Dill introduced in [1] the progress condition*, which ensures the existence of accepted non-Zeno behaviours. This condition is the existence of a reachable SCC in* $\mathsf{R}(\mathcal{A})$ *which is unbounded or which resets a clock.*

From $\mathcal{A}$, we will derive a finite Markov chain $\mathcal{M}_\mathcal{A}$ such that for every location-based prefix-independent property φ, there is a set F_φ of states in $\mathcal{M}_\mathcal{A}$, s.t. $\mathbb{P}_\mathcal{A}(s_0 \models \varphi) = \mathbb{P}_{\mathcal{M}_\mathcal{A}}(q_0 \models \Diamond F_\varphi)$, where q_0 is a distinguished state of $\mathcal{M}_\mathcal{A}$ and $\mathbb{P}_{\mathcal{M}_\mathcal{A}}$ is the classical probability measure on sets of runs

[7] Formally a property $L \subseteq \Sigma^\omega$ is *location-based* if its truth value along a run only depends on the locations that are visited, and not on the clock valuations. It is *prefix-independent* if for all $w \in \Sigma^\omega$ and $u \in \Sigma^*$, $uw \in L$ iff $w \in L$. In other words, the satisfaction of a prefix-independent property by a word only depends on the set of atomic proposition true in infinitely many positions of that word. Note that this kind of property is commonly used for games objectives (see *e.g.*[6, 17] or [10] where they are referred to as "tail objectives".)

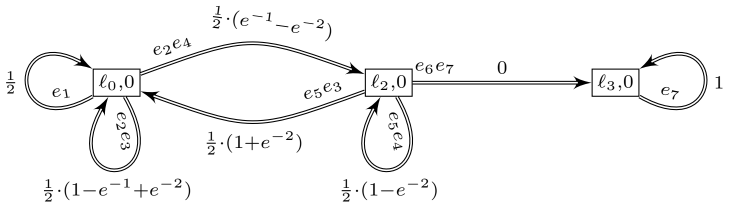
$\frac{1}{2}\cdot(e^{-1}-e^{-2})$
e_2e_4
$\frac{1}{2}$
e_1
ℓ_0,0
e_5e_3
ℓ_2,0
e_6e_7
0
ℓ_3,0
e_7
1
e_2e_3
e_5e_4
$\frac{1}{2}\cdot(1+e^{-2})$
$\frac{1}{2}\cdot(1-e^{-1}+e^{-2})$
$\frac{1}{2}\cdot(1-e^{-2})$

Fig. 2. The finite Markov chain $\mathcal{M}_{\mathcal{A}}$ for the running example

in $\mathcal{M}_{\mathcal{A}}$. The value $\mathbb{P}_{\mathcal{M}_{\mathcal{A}}}(q_0 \models \Diamond F_\varphi)$, which is the probability of a reachability property in a finite Markov chain, can be computed as the solution of a system of linear equations [11]. Hence the probability of the set of runs satisfying a location-based property in $\mathcal{A}$ will be expressible using the probability distributions put on edges of the Markov chain $\mathcal{M}_{\mathcal{A}}$.

We now detail the construction of the finite Markov chain $\mathcal{M}_{\mathcal{A}}$. The idea is that, thanks to hypotheses (†), a run in $\mathsf{R}(\mathcal{A})$ will either often visit an unbounded state of $\mathsf{R}(\mathcal{A})$, or often reset clock x. The set of states of $\mathcal{M}_{\mathcal{A}}$ is then $\{(\ell,0) \mid (\ell, x=0) \text{ state of } \mathsf{R}(\mathcal{A})\} \cup \{(\ell,\infty) \mid (\ell, x>M) \text{ state of } \mathsf{R}(\mathcal{A})\}$. We note $E_{:=0}$ the set of edges of $\mathcal{A}$ which reset clock x, and $E_{>M}$ the set of edges of $\mathcal{A}$ guarded by the constraint $x > M$. The set of transitions of $\mathcal{M}_{\mathcal{A}}$ is defined as follows.

1. Let $\pi((\ell,0), e_1 \dots e_p)$ be a non-empty loop-free (*i.e.*, the e_i's are all distinct) symbolic path such that for every $1 \le i < p$, $e_i \notin E_{:=0} \cup E_{>M}$, and $e_p \in E_{:=0} \cup E_{>M}$. If $e_p \in E_{:=0}$, we add a transition $(\ell,0) \xrightarrow{e_1 \dots e_p} (\ell',0)$ in $\mathcal{M}_{\mathcal{A}}$. If $e_p \in E_{>M} \setminus E_{:=0}$, we add a transition $(\ell,0) \xrightarrow{e_1 \dots e_p} (\ell',\infty)$ in $\mathcal{M}_{\mathcal{A}}$. In both cases, we label the transition with $\mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi))$;

2. For each edge $e \in E_{>M}$ leaving a state $(\ell, x > M)$ of $\mathsf{R}(\mathcal{A})$, we add a transition $(\ell,\infty) \xrightarrow{e} (\ell',0)$ if $e \in E_{:=0}$ and $(\ell', x=0)$ is the target state of e in $\mathsf{R}(\mathcal{A})$, and we add a transition $(\ell,\infty) \xrightarrow{e} (\ell',\infty)$ if $e \notin E_{:=0}$ and $(\ell', x > M)$ is the target state of e in $\mathsf{R}(\mathcal{A})$. In both cases, we label the edge with $w(e)/\Big(\sum_{e' \text{ enabled from } (\ell,x>M)} w(e')\Big)$.

Example 9 *We illustrate the construction on the automaton depicted in Figure 1. To locations ℓ_0, ℓ_2 and ℓ_3, we assign the uniform distribution over delays, whereas the density of the distribution over delays in location ℓ_1 is supposed to be $t \mapsto e^{-t}$ over $\mathbb{R}_+$. We assume that the weight of each edge is* 1*, so that the discrete choices are uniform. In that case,*[8] *we have $E_{:=0} = \{e_1, e_3, e_4, e_6, e_7\}$, and $E_{>2} = \{e_4\}$. Note that as $E_{>2} \subseteq E_{:=0}$, there won't be any transition of the second type, and no state of the form (ℓ,∞) will be reachable. The construction is depicted on Figure 2, where each transition is labelled with the corresponding sequence of edges together with its probability. An edge of $\mathcal{M}_{\mathcal{A}}$ corresponds to a (finite) sequence of edges in $\mathcal{A}$, hence is somehow a* macro-edge*, explaining the use of double arrows in the figure.*

The first property we have to check is that $\mathcal{M}_{\mathcal{A}}$ is indeed a finite Markov chain, which is not obvious from the above construction. This result is however true as stated in the following lemma.

Lemma 10 *$\mathcal{M}_{\mathcal{A}}$ is a finite Markov chain.*

We can now state the correctness of our abstraction into a finite Markov chain, under assumption (†):

Theorem 11 *Let φ be a location-based prefix-independent property on $\mathcal{A}$. We can compute a set F_φ of states of $\mathcal{M}_{\mathcal{A}}$ that is SCC-closed*[9]*and s.t.*

$$\mathbb{P}_{\mathcal{A}}\Big((\ell_0,0) \models \varphi\Big) = \mathbb{P}_{\mathcal{M}_{\mathcal{A}}}\Big((\ell_0,0) \models \Diamond F_\varphi\Big).$$

We have reduced the quantitative model-checking problem for location-based prefix-independent properties in $\mathcal{A}$ to a quantitative reachability question in a finite Markov chain. However we are still not done, because the values labelling the edges of $\mathcal{M}_{\mathcal{A}}$ may not have closed-form representations in general, even for very simple distributions over the time (see Example 12 below). In the next section, we will further restrict our model, and provide a framework in which the probabilities can effectively be computed.

Example 12 *Consider the automaton $\mathcal{A}$ of Fig. 3, on which we assume uniform distributions over delays, and assign weight* 1 *to every edge. We can easily compute the following*

[8] Note that formally, $\mathcal{A} \neq \mathsf{R}(\mathcal{A})$, in that case, but the construction can still be done.

[9] Which means that for any $q \in F_\varphi$ and any q' in the same SCC of $\mathsf{R}(\mathcal{A})$ as q, we have $q' \in F_\varphi$.

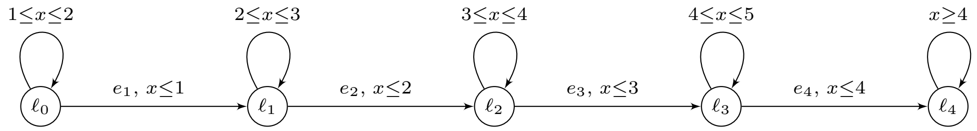

Fig. 3. Automaton $\mathcal{A}$

probability:

$$\mathbb{P}_{\mathcal{A}}\Big(\mathit{Cyl}(\pi((\ell_0,0),e_1e_2e_3))\Big) \;=\; \frac{1}{2}\Big(1+(\ln(2)-\ln(3))(1-\ln(2))+ \mathrm{dilog}(4)-\mathrm{dilog}(3)\Big) \;\approx\; 0.19$$

where $x \mapsto \mathrm{dilog}(x)$ *is the primitive of* $x \mapsto -\frac{\ln(x)}{1-x}$.

Note that there does not seem to exist a closed-form solution for $\mathbb{P}_{\mathcal{A}}\Big(\mathit{Cyl}(\pi((\ell_0,0),e_1e_2e_3e_4))\Big)$.

4 Quantitative Model-Checking Made Decidable

We have seen in the previous section a correct abstraction for computing the probabilities of prefix-independent location-based properties. However we have also seen that we could not always obtain a closed-form expression for those probabilities, hence we cannot really compute values. In this section, on top of hypotheses (†) made in the previous section, we assume that for every state s of $\mathcal{A}$, it holds $I(s)=\mathbb{R}_+$, and that we have exponential distributions over delays which are "uniform by location": for every location ℓ of $\mathcal{A}$, there is a positive constant $\lambda_\ell \in \mathbb{Q}_{>0}$ (called the rate of ℓ) such that for every state $s=(\ell,u)$, the measure μ_s has density $t \mapsto \lambda_\ell \cdot \exp(-\lambda_\ell \cdot t)$. We write (‡) for these additional restrictions.

Remark 13 *Note that single-clock timed automata, even under restrictions (†) and (‡), are still a generalization of continuous-time Markov chains [5]. Indeed continuous-time Markov chains can be seen as single-clock timed automata without guards, that reset the clock after each transition, and for which the probabilistic distribution over delays is a decreasing exponential.*

4.1 Expressing the Probability

We assume that $\mathcal{A}=\mathsf{R}(\mathcal{A})$ is a single-clock timed automaton satisfying hypotheses (†) and (‡). We let $s_0=(\ell_0,0)$ be the initial state of $\mathcal{A}$. For every location ℓ of $\mathcal{A}$, we write λ_ℓ for the rate of ℓ.

Proposition 14 *Let* φ *be a prefix-independent location-based property. Then,* $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ *can be expressed as* $f\left(e^{-\frac{1}{q}}\right)$ *for some positive integer q, where* $f \in \mathbb{Q}(X)$ *is a rational function.*

To prove this result, we first show that any value labelling a transition of $\mathcal{M}_{\mathcal{A}}$, the finite Markov chain constructed in the previous section, is the evaluation of some polynomial at $e^{-\frac{1}{q}}$ (for some $q \in \mathbb{N}_{>0}$).

Lemma 15 *Let* $e_1,\dots,e_n$ *be edges of* $\mathcal{A}$ *and let* (ℓ,r) *be a state of* $\mathsf{R}(\mathcal{A})$. *Then the function*

$$\begin{array}{rcl} r & \to & [0,1] \\ t & \mapsto & \mathbb{P}_{\mathcal{A}}\Big(\mathit{Cyl}\big(\pi((\ell,t),e_1\dots e_n)\big)\Big) \end{array}$$

can be written as a function of the form:

$$t \in r \mapsto \sum_{\ell \in L} \exp(\lambda_\ell t)\cdot P_\ell\Big((e^{\lambda_{\ell'}})_{\ell' \in L},(e^{-\lambda_{\ell'}})_{\ell' \in L}\Big) + P\Big((e^{\lambda_{\ell'}})_{\ell' \in L},(e^{-\lambda_{\ell'}})_{\ell' \in L}\Big)$$

where $(P_\ell)_{\ell \in L}$ *and* P *are multivariate polynomials in* $\mathbb{Q}[(X_\ell)_{\ell \in L},(Y_\ell)_{\ell \in L}]$.

The proof of this lemma is by induction on the length of unconstrained symbolic paths. It consists in a simple but tedious case inspection.

We now come to the proof of Proposition 14.

Proof. By Theorem 11 we know that computing the probability of satisfying φ in $\mathcal{A}$ can be converted into the computation of the probability of a reachability property in $\mathcal{M}_{\mathcal{A}}$. We then use the following two facts:

- Computing the probability to reach a set of states in a finite Markov chain amounts to solving a system of linear equations, whose coefficients are probability values labelling the transitions of the Markov chain [9].
- By construction, values labelling transitions leaving a state of the form (ℓ,∞) are rational. According to Lemma 15, the value labelling a transition leaving a state $(\ell,0)$ is of the form $P\Big((e^{\lambda_\ell})_\ell,(e^{-\lambda_\ell})_\ell\Big)$ for some polynomial $P \in \mathbb{Q}[(X_\ell)_{\ell \in L},(Y_\ell)_{\ell \in L}]$. Hence the transition probabilities in $\mathcal{M}_{\mathcal{A}}$ can all be written in the previous form.

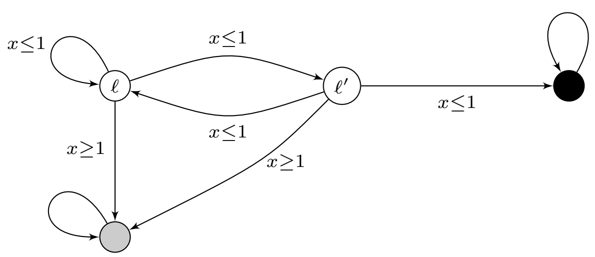

Fig. 4. An example with a non-resetting bounded cycle

We now prove that solving the linear equation system yields a solution of the desired form. Since the λ_ℓ's are all assumed to be positive rational numbers, there exists $q \in \mathbb{N}_{>0}$ and integers $(p_\ell)_\ell$ such that for every ℓ, $\lambda_\ell = \frac{p_\ell}{q}$. As a consequence, and using the property of the exponential function that $e^{-a/b} = (e^{-1/b})^a$, each transition probability $P((e^{\lambda_\ell})_\ell, (e^{-\lambda_\ell})_\ell)$ can be rewritten as $(e^{1/q})^k \cdot Q(e^{-1/q})$ where $k \in \mathbb{N}$, and $Q \in \mathbb{Q}[X]$. With such coefficients, the solution of the linear equations system has the desired form $f(e^{-1/q})$, with $f \in \mathbb{Q}(X)$ a rational function. □

Example 16 *We illustrate on an example why Proposition 14 really relies on hypothesis* $(\dagger)$*, and more precisely on the hypothesis that any bounded cycle of* $\mathsf{R}(\mathcal{A})$ *contains at least one resetting edge. Consider the automaton in Figure 4, in which we assume a weight* 1 *per edge, and an exponential distribution of density* $t \mapsto \lambda \cdot e^{-\lambda t}$ *in locations* ℓ *and* ℓ'*. The probability of reaching the black location is* 1 *from the black location, and* 0 *from the grey location. Now we write* $\mathfrak{p}_\ell$ *(resp.* $\mathfrak{p}_{\ell'}$*) the function which associate to every* $x \in \mathbb{R}_+$ *the probability of reaching the black location from* (ℓ, x) *(resp.* (ℓ', x)*). It is not hard to be convinced that for every* $x \geq 1$*,* $\mathfrak{p}_\ell(x) = \mathfrak{p}_{\ell'}(x) = 0$*, and that for every* $x \leq 1$*,*

$$\begin{cases} \mathfrak{p}_\ell(x) = \displaystyle\int_{t=0}^{1-x} \frac{\lambda}{2} \cdot e^{-\lambda t} \cdot \mathfrak{p}_{\ell'}(x+t)\,\mathrm{d}t \\ \qquad\qquad + \displaystyle\int_{t=0}^{1-x} \frac{\lambda}{2} \cdot e^{-\lambda t} \cdot \mathfrak{p}_\ell(x+t)\,\mathrm{d}t \\ \mathfrak{p}_{\ell'}(x) = \displaystyle\int_{t=0}^{1-x} \frac{\lambda}{2} \cdot e^{-\lambda t} \mathfrak{p}_\ell(x+t)\,\mathrm{d}t \\ \qquad\qquad + \displaystyle\int_{t=0}^{1-x} \frac{\lambda}{2} \cdot e^{-\lambda t}\,\mathrm{d}t \end{cases}$$

Deriving the above integral equations, we get differential equations that we can solve, and we get the following solutions for all $x \leq 1$*:*

$$\begin{cases} \mathfrak{p}_\ell(x) = 1 - \frac{5+3\sqrt{5}}{10} \exp(\frac{\lambda}{4}(3-\sqrt{5})(x-1)) \\ \qquad\qquad + \frac{3\sqrt{5}-5}{10} \exp(\frac{\lambda}{4}(3+\sqrt{5})(x-1)) \\ \mathfrak{p}_{\ell'}(x) = 1 - \frac{5+\sqrt{5}}{10} \exp(\frac{\lambda}{4}(3-\sqrt{5})(x-1)) \\ \qquad\qquad - \frac{5-\sqrt{5}}{10} \exp(\frac{\lambda}{4}(3+\sqrt{5})(x-1)) \end{cases}$$

We immediately notice that these expressions do not match the general form described in Proposition 14.

4.2 Approximating the Probability

From the previous subsection, given a timed automaton $\mathcal{A}$ with initial state $s_0 = (\ell_0, 0)$, and a location-based prefix-independent property φ, we can effectively compute a rational function $f \in \mathbb{Q}(X)$ and a positive integer $q \in \mathbb{N}_{>0}$ such that $\mathbb{P}_\mathcal{A}((\ell_0, 0) \models \varphi) = f\left(e^{-1/q}\right)$.

We now explain how to approximate this quantity with a precision $\varepsilon > 0$.

First we notice that we can compute two approximating sequences $(a_i)_{i\in\mathbb{N}}$ and $(b_i)_{i\in\mathbb{N}}$ of rational numbers such that:

- $\forall i$, $a_i \leq a_{i+1} \leq e^{-\frac{1}{q}} \leq b_{i+1} \leq b_i$, and
- $\lim_{i\to\infty} a_i = \lim_{i\to\infty} b_i = e^{-\frac{1}{q}}$.

These two sequences $(a_i)_{i\in\mathbb{N}}$ and $(b_i)_{i\in\mathbb{N}}$ can be obtained using the Maclaurin series of the exponential function. Indeed, for all $x \in \mathbb{R}_{>0}$, $e^{-x} = \sum_{k=0}^{\infty} \frac{(-x)^k}{k!}$. Hence, in order to approximate $e^{-\frac{1}{q}}$, one can set $b_i = \sum_{k=0}^{2i} \frac{(-1/q)^k}{k!}$ and $a_i = \sum_{k=0}^{2i+1} \frac{(-1/q)^k}{k!}$.

Then, we remark that $e^{-1/q}$ is a transcendental number (because e is), and we prove that, on a sufficiently small (computable) interval (a, b) containing a transcendental real ζ, a rational function $f \in \mathbb{Q}(X)$ is monotonic.

Lemma 17 *Let* $f \in \mathbb{Q}(X)$ *be a rational function, and* $\zeta \in \mathbb{R}$ *be a transcendental number. There exist two rational numbers* $\alpha, \beta \in \mathbb{Q}$ *such that* $\zeta \in (\alpha, \beta)$*, and* f *is monotonic over the interval* (α, β)*. Moreover, if* ζ *has two approximating sequences as described above, then* α *and* β *can be effectively computed.*

Proof. Let $P, Q \in \mathbb{Q}[X]$ such that $f = P/Q$. Since $f' = P'Q - PQ'/Q^2$ it is sufficient to prove that the polynomial $R \overset{\text{def}}{=} P'Q - PQ'$ has a constant sign over some interval (α, β) containing ζ. The reason for that is that ζ is transcendental, hence $R(\zeta) \neq 0$ (provided $R \neq 0$) and by continuity, R has a constant sign over some neighbourhood of ζ.

To show the effectiveness of the construction of (α, β), provided that ζ can be approximated by two sequences (one increasing and one decreasing) $(a_i), (b_i) \in \mathbb{Q}^{\mathbb{N}}$, one first proves that given a polynomial $R \in \mathbb{Q}[X]$, there exist $\alpha, \beta \in \mathbb{Q}$ such that $\zeta \in (\alpha, \beta)$ and R has a constant sign over (α, β) (see Lemma 18 below). Applying this result to R yields an interval (α, β) that contains ζ and over which f is monotonic. □

Lemma 18 *Let $R \in \mathbb{Q}[X]$ be a non-zero polynomial and $\zeta \in \mathbb{R}$ be a transcendental number. Then, there exist $\alpha, \beta \in \mathbb{Q}$ such that $\zeta \in (\alpha, \beta)$ and R has constant sign over (α, β). Moreover, if there are approximating sequences $(a_i)_{i\in\mathbb{N}}$ and $(b_i)_{i\in\mathbb{N}}$ in $\mathbb{Q}^{\mathbb{N}}$ as described above, then α and β can be effectively computed.*

Proof. The existence of α, β is due both to the fact that $R(\zeta) \neq 0$ (since ζ is transcendental) and to the continuity of R.

The computability of some α, β requires assumptions on ζ. We assume that there are two approximating sequences $(a_i)_{i\in\mathbb{N}}$ and $(b_i)_{i\in\mathbb{N}}$ in $\mathbb{Q}^{\mathbb{N}}$ as described before. Under this assumption, we now prove that we can compute such values α and β by induction on the degree of polynomial R.

degree 0: Assume R has degree 0, or equivalently R is a constant function over $\mathbb{R}$. Letting, *e.g.*, $\alpha = a_1$ and $\beta = b_1$ works.

degree n+1: Assume now that the degree of R is $n + 1$ for $n \in \mathbb{N}$. The induction hypothesis applied to R' yields the existence and computability of $\alpha_n, \beta_n \in \mathbb{Q}$ such that $\zeta \in (\alpha_n, \beta_n)$ and R' is of constant sign over the interval (α_n, β_n). Hence R is monotonic over (α_n, β_n). Since $R(\zeta) \neq 0$ and R is continuous, the monotonicity of R over (α_n, β_n) implies the existence of an interval $I \subseteq (\alpha_n, \beta_n)$ containing ζ and over which R has a constant sign. Now, starting from a_i, b_i with i large enough to have $(a_i, b_i) \subseteq (\alpha_n, \beta_n)$, it suffices to find some index $j \geq i$ with $R(a_j) \cdot R(b_j) > 0$. Letting $(\alpha_{n+1}, \beta_{n+1}) = (a_j, b_j)$ yields the expected result.

This ends the proof of Lemma 18. □

Approximation scheme. Let $\varepsilon > 0$ be an approximant. To approximate $f\left(e^{-1/q}\right)$ ε-closely, the idea is to evaluate f at $(a_i)_{i\geq N}$ and $(b_i)_{i\geq N}$ for some $N \in \mathbb{N}$ large enough so that f is monotonic over the interval (a_N, b_N). These evaluations lead to two sequences $(f(a_i))_{i\geq N}$ and $(f(b_i))_{i\geq N}$, one of which is increasing and the other decreasing, both converging towards $f(e^{-1/q})$ (because f is continuous). The difference $(|f(a_i) - f(b_i)|)_{i\geq N}$ decreases to 0, hence eventually, for some index i, we will have that $|f(a_i) - f(b_i)| < \varepsilon$. Hence one of $f(a_i)$ or $f(b_i)$ will be an over-approximation for $f\left(e^{-1/q}\right)$, and the other will be an under-approximation of $f\left(e^{-1/q}\right)$. We thus get the following result:

Theorem 19 *Let $\mathcal{A}$ be a single-clock timed automaton satisfying the hypotheses (†) and (‡), let φ be a location-based prefix-independent property. Assume that s_0 is the initial state of $\mathcal{A}$. We can decide if $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ is a rational, compute it if it is rational, and if not, for every $\varepsilon > 0$, we can compute two rationals P^-_ε and P^+_ε such that:*

$$\begin{cases} P^-_\varepsilon \leq \mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) \leq P^+_\varepsilon \\ P^+_\varepsilon - P^-_\varepsilon < \varepsilon \end{cases}$$

4.3 Deciding the Threshold Problem

We recall that the *threshold problem* asks, given a timed automaton $\mathcal{A}$ with its initial state s_0, an ω-regular property φ, and a threshold $\sim c$ with $\sim \in \{<, \leq, =, \geq, >\}$ and $c \in \mathbb{Q}$, whether $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) \sim c$.

As a consequence of the previous subsection, we get the decidability of the threshold problem.

Theorem 20 *The threshold problem is decidable for single-clock timed automata satisfying hypotheses (†) and (‡).*

Proof. Thanks to Theorem 19, we can decide whether $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ is rational or not, and compute it (and answer the threshold problem) if it is rational.

Now assume that $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ is not rational. Then the answer to the threshold problem is negative when $\sim$ is $=$ (since c is rational), and the answer to the problem coincide when $\sim$ is $<$ and $\leq$ (similarly for $>$ and $\geq$). Hence we need only be able to solve the problem when $\sim$ is $<$ or $>$.

We have seen that we could compute ε-close upper and lower approximations of $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ for arbitrarily small $\varepsilon > 0$. Hence, it suffices to obtain ε-approximations for $\varepsilon \leq |c - \mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)|$. This is achieved as follows: for every $n \in \mathbb{N}$, compute $\frac{1}{2^n}$-approximations γ_1 and γ_2, and stop when both are on the same side of c. □

5 Conclusion

In this paper we have studied the probabilistic (and quantitative) model-checking problem for single-clock timed automata, in which choices for delays and discrete events are probabilized. We have defined an abstraction, which takes the form of a finite Markov chain, which is correct for a subclass of automata for computing the probability that an ω-regular property holds in the system. However, the probability that is computed might not be a closed-form expression. Hence we have described a more restricted framework, where distributions over delays are given as exponential distributions, and in which we can compute closed-form expressions for the probability of ω-regular properties, we can approximate these values, and decide the threshold problem.

Further work includes approximation schemes for more general frameworks than the one described here, for instance for bounded automata, when distributions over delays are given as uniform distributions, since this also constitutes a natural framework.

References

[1] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

[2] R. Alur, S. La Torre, and P. Madhusudan. Perturbed timed automata. In *Proc. 8th International Workshop on Hybrid Systems: Computation and Control (HSCC'05)*, Lecture Notes in Computer Science 3414, p. 70–85. Springer, 2005.

[3] C. Baier, N. Bertrand, P. Bouyer, T. Brihaye, and M. Größer. Almost-sure model checking of infinite paths in one-clock timed automata. In *Proc. 23rd Annual Symposium on Logic in Computer Science (LICS'08)*. IEEE Computer Society Press, 2008. To appear.

[4] C. Baier, N. Bertrand, P. Bouyer, Th. Brihaye, and M. Größer. Probabilistic and topological semantics for timed automata. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, Lecture Notes in Computer Science 4855, p. 179–191. Springer, 2007.

[5] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(7):524–541, 2003.

[6] D. Berwanger. Admissibility in infinite games. In *Proc. 24th Annual Symposium on Theoretical Aspects of Computer Science (STACS'07)*, Lecture Notes in Computer Science 4393, p. 188–199. Springer, 2007.

[7] P. Bouyer, N. Markey, and P.-A. Reynier. Robust model-checking of timed automata. In *Proc. 7th Latin American Symposium on Theoretical Informatics (LATIN'06)*, Lecture Notes in Computer Science 3887, p. 238–249. Springer, 2006.

[8] P. Bouyer, N. Markey, and P.-A. Reynier. Robust analysis of timed automata via channel machines. In *Proc. 11th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'08)*, Lecture Notes in Computer Science 4962, p. 157–171. Springer, 2008.

[9] P. Brémaud. *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*. Springer, 1999.

[10] K. Chatterjee. Concurrent games with tail objectives. *Theoretical Computer Science*, 388(1-3):181–198, 2007.

[11] C. Courcoubetis and M. Yannakakis. The complexity of probabilistic verification. *Journal of the ACM*, 42(4):857–907, 1995.

[12] M. De Wulf, L. Doyen, N. Markey, and J.-F. Raskin. Robustness and implementability of timed automata. In *Proc. Joint Conference on Formal Modelling and Analysis of Timed Systems and Formal Techniques in Real-Time and Fault Tolerant System (FORMATS+FTRTFT'04)*, Lecture Notes in Computer Science 3253, p. 118–133. Springer, 2004.

[13] M. De Wulf, L. Doyen, and J.-F. Raskin. Almost ASAP semantics: From timed models to timed implementations. In *Proc. 7th International Workshop on Hybrid Systems: Computation and Control (HSCC'04)*, Lecture Notes in Computer Science 2993, p. 296–310. Springer, 2004.

[14] J. Esparza, A. Kučera, and R. Mayr. Model checking probabilistic pushdown automata. *Logical Methods in Computer Science*, 2006.

[15] V. Gupta, Th. A. Henzinger, and R. Jagadeesan. Robust timed automata. In *Proc. International Workshop on Hybrid and Real-Time Systems (HART'97)*, Lecture Notes in Computer Science 1201, p. 331–345. Springer, 1997.

[16] Th. A. Henzinger and J.-F. Raskin. Robust undecidability of timed and hybrid systems. In *Proc. 3rd International Workshop on Hybrid Systems: Computation and Control (HSCC'00)*, Lecture Notes in Computer Science 1790, p. 145–159. Springer, 2000.

[17] E. Kopczynski. Omega-regular half-positional winning conditions. In *Proc. 21th International Workshop on Computer Science Logic (CSL'07)*, Lecture Notes in Computer Science 4646, p. 41–53. Springer, 2007.

[18] A. Pnueli. The temporal logic of programs. In *Proc. 18th Annual Symposium on Foundations of Computer Science (FoCS'77)*, p. 46–57. IEEE Computer Society Press, 1977.

[19] A. Rabinovich. Quantitative analysis of probabilistic lossy channel systems. In *Proc. 30th International Colloquium on Automata, Languages and Programming (ICALP'03)*, Lecture Notes in Computer Science 2719, p. 1008–1021. Springer, 2003.