ENS CACHAN

Christel Baier, Nathalie Bertrand, Patricia Bouyer,

Thomas Brihaye, and Marcus Größer

# Almost-Sure Model Checking of Infinite Paths in One-Clock Timed Automata

Research Report LSV-07-29

19 September 2007

Laboratoire
Spécification
et
Vérification

CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

Ecole Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

# Almost-Sure Model Checking of Infinite Paths in One-Clock Timed Automata

Christel Baier[1], Nathalie Bertrand[1,⋆], Patricia Bouyer[2,3,⋆⋆],
Thomas Brihaye[2], and Marcus Größer[1]

[1] Technische Universität Dresden, Germany
[2] LSV - CNRS & ENS de Cachan, France
[3] Oxford University Computing Laboratory, UK

**Abstract.** In this paper, we define two relaxed semantics (one based on probabilities and the other one based on the topological notion of largeness) for LTL over infinite runs of timed automata which rule out unlikely sequences of events. We prove that these two semantics match in the framework of single-clock timed automata (and only in that framework), and prove that the corresponding relaxed model-checking problems are PSPACE-Complete. Moreover, we prove that the probabilistic non-Zenoness can be decided for single-clock timed automata in NLOGSPACE.

## 1 Introduction

Nowadays *timed automata* [AD94] are a well-established formalism for the modelling and analysis of timed systems. Roughly speaking timed automata are finite state automata enriched with clocks and clock constraints. This model has been intensively studied, and several verification tools have been developed. However, like most models used in model checking, timed automata are an idealized mathematical model. In particular it has infinite precision, instantaneous events, *etc.* Recently, more and more research is devoted to propose alternative semantics for timed automata that provide more realistic operational models for real-time systems. Let us first mention the *Almost ASAP semantics* introduced in [DDR04] and further studied in [DDMR04,ALM05,BMR06]. This AASAP semantics somewhat relaxes the constraints and precision of clocks. 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 [GHJ97,HR00]. In this framework, a metric is put on the set of traces of the timed automaton, and 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 for implementability issues, 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.

[⋆] Partly supported by a Lavoisier fellowship (program of the French Ministry of Foreign Affairs).
[⋆⋆] Partly supported by a Marie Curie fellowship (program of the European Commission).

Varacca and Völzer recently proposed in [VV06] a *probabilistic framework for finite-state (time-abstract) systems* to overcome side-effects of modelling. They use probabilities to define the notion of being fairly correct as having probability zero to fail, when every non-deterministic choice has been transformed into a "reasonable" probabilistic choice. Moreover, in their framework, a system is fairly correct with respect to some property if and only if the set of traces satisfying that property in the system is topologically large, which somehow attests the relevance of this notion of fair correctness.

In the recent paper [BBB$^+$07], we used similar concepts as in [VV06] and proposed two alternative semantics for reasoning about the *finite runs* of timed automata: *(i)* a *probabilistic semantics* which assigns probabilities both on delays and on discrete choices, and *(ii)* a *topological semantics*, following ideas of [GHJ97,HR00] but rather based on the structure of the automaton than on its accepted language. For both semantics, we naturally addressed a model-checking problem for LTL interpreted over finite paths. We proved, by means of Banach-Mazur games, that both semantics coincide and that both model-checking problems for LTL specifications on finite words are PSPACE-Complete.

The purpose of this paper is to adapt the techniques proposed in [BBB$^+$07] for analyzing the *infinite behaviours* of timed automata by means of a probabilistic *almost-sure* interpretation of LTL over infinite runs (which requires that the given LTL formula $\varphi$ holds with probability 1) and a topological interpretation (which requires topological largeness of the set of infinite runs where $\varphi$ holds). The formal definitions of the almost-sure and topological semantics of LTL interpreted over the infinite runs in a timed automata are rather straightforward adaptions of the corresponding definitions in the case of finite runs [BBB$^+$07]. However, to establish a link between the two semantics and to show that the topological semantics of LTL is reasonable in the sense that it matches the standard meaning of negation, the proof techniques used in [BBB$^+$07] are no longer appropriate. Instead, methods are required that are specific for infinite runs. To confirm that the topological semantics yields a reasonable interpretation for LTL, we prove that the underlying topology constitutes a *Baire space*. For the case of one-clock timed automata, we will show that some kind of strong *fairness* is inherent in the almost-sure semantics. This observation will be used to prove that the almost-sure and topological semantics for infinite paths in one-clock timed automata agree. As the topological semantics only relies on the graph-structure of the given automaton (but not on any quantitative assumption on the resolution of the nondeterministic choices as it is the case for the probabilistic setting), this result yields the key to establish a polynomially space-bounded model checking algorithm for LTL over infinite words with respect to our non-standard semantics. In addition, we introduce a notion of probabilistic non-Zenoness, which requires that the set of Zeno runs have measure 0, and show that it has a simple topological characterization which can serve as a basis for a nondeterministic logarithmic space-bounded algorithm to checking probabilistic non-Zenoness. We also show that analogous results cannot be established for timed automata with two or more clocks, as then the probabilistic and topological semantics for $\mathsf{LTL}$ over infinite words do not agree.

***Organisation of the paper.*** Section 2 summarizes our notations for timed automata, $\mathsf{LTL}$ and the relevant topological concepts. The probabilistic space and the topological space associated with a timed automaton together with the almost-sure and topological $\mathsf{LTL}$ semantics are defined in Section 3. The relation between the two semantics and the induced model checking problems are studied in Section 4. Probabilistic Zenoness is considered in Section 5.

## 2 Preliminaries

### 2.1 The timed automaton model

Let us first recall the notions of clocks, clock valuations, and guards needed to define timed automata [AD94]. We denote by $X = \{x_1, \dots, x_k\}$ a finite set of *clocks*. A *clock valuation* over $X$ is a mapping $\nu : X \to \mathbb{R}_+$, where $\mathbb{R}_+$ denotes the set of nonnegative reals. We write $\mathbb{R}_+^X$ for the set of clock valuations over $X$. Given a clock valuation $\nu$ and $\tau \in \mathbb{R}_+$, $\nu + \tau$ is the clock valuation defined by $(\nu + \tau)(x) = \nu(x) + \tau$ for every $x \in X$. If $Y \subseteq X$, the valuation $[Y \leftarrow 0]\nu$ is the valuation $\nu'$ such that $\nu'(x) = 0$ if $x \in Y$, and $\nu'(x) = \nu(x)$ otherwise. A *guard* over $X$ is a finite conjunction of expressions of the form $x \sim c$ where $x \in X$ is a clock, $c \in \mathbb{N}$ is an integer, and $\sim$ is one of the symbols $\{<, \leq, =, \geq, >\}$. 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 the clock valuation $\nu$ satisfies the guard $g$. We denote by $\mathsf{AP}$ a finite set of atomic propositions.

**Definition 1.** *A* timed automaton *is a tuple $\mathcal{A} = (L, X, E, \mathcal{I}, \mathcal{L})$ 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,* $(iv)$ *$\mathcal{I} : L \to \mathcal{G}(X)$ assigns an invariant to each location, and* $(v)$ *$\mathcal{L} : L \to 2^{\mathsf{AP}}$ is a labelling function.*

The semantics of a timed automaton $\mathcal{A}$ is given by a timed transition system $T_{\mathcal{A}} = (S, E, \mathbb{R}_+, \to)$ where the set $S$ of states is $\{s = (\ell, \nu) \in L \times \mathbb{R}_+^X \mid \nu \models \mathcal{I}(\ell)\}$, and the transition relation $\to \subseteq (S \times (E \cup \mathbb{R}_+) \times S)$ is composed of delay and discrete transitions as follows:

- *(delay transition)* $(\ell, \nu) \xrightarrow{\tau}_{\mathcal{A}} (\ell, \nu + \tau)$ if $\tau \in \mathbb{R}_+$ and if for all $0 \leq \tau' \leq \tau$, $\nu + \tau' \models \mathcal{I}(\ell)$,
- *(discrete transition)* $(\ell, \nu) \xrightarrow{e}_{\mathcal{A}} (\ell', \nu')$ if $e = (\ell, g, Y, \ell') \in E$ is such that $\nu \models \mathcal{I}(\ell) \wedge g$, $\nu' = [Y \leftarrow 0]\nu$, and $\nu' \models \mathcal{I}(\ell')$.

A finite (resp. infinite) *run* $\varrho$ of $\mathcal{A}$ is a finite (resp. infinite) sequence of states obtained by alternating delay and discrete transitions, *i.e.*, $\varrho = s_0 \xrightarrow{\tau_1} s'_1 \xrightarrow{e_1} s_1 \xrightarrow{\tau_2} s'_2 \xrightarrow{e_2} s_2 \dots$ or more compactly $s_0 \xrightarrow{\tau_1,e_1} s_1 \xrightarrow{\tau_2,e_2} s_2 \dots$ We write $\mathsf{Runs}_f(\mathcal{A}, s_0)$ (resp. $\mathsf{Runs}(\mathcal{A}, s_0)$) for the set of runs (resp. infinite runs) of $\mathcal{A}$ from state $s_0$.

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

$$\pi_{\mathcal{C}}(s, e_1 \dots e_n) \;=\; \{\varrho = s \xrightarrow{\tau_1,e_1} s_1 \dots \xrightarrow{\tau_n,e_n} s_n \mid \varrho \in \mathsf{Runs}_f(\mathcal{A}, s) \text{ and } (\tau_i)_{1\le i\le n} \models \mathcal{C}\}\,.$$

If $\mathcal{C}$ is equivalent to 'true', we simply write $\pi(s, e_1 \dots e_n)$. Let $\pi_{\mathcal{C}} = \pi_{\mathcal{C}}(s, e_1 \dots e_n)$ be a finite symbolic path, we define the *cylinder* generated by $\pi_{\mathcal{C}}$ as

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

In the following, we will also use infinite symbolic paths defined, given $s$ a state of $\mathcal{A}$ and $(e_i)_{i\ge1}$ an infinite sequence of edges, as:

$$\pi(s, e_1e_2 \dots) = \{\varrho = s \xrightarrow{\tau_1,e_1} s_1 \xrightarrow{\tau_2,e_2} s_2 \dots \mid \varrho \in \mathsf{Runs}(\mathcal{A}, s)\}\,.$$

If $\varrho \in \mathsf{Runs}(\mathcal{A}, s)$, we write $\pi_\varrho$ for the unique symbolic path containing $\varrho$. Given $s$ a state of $\mathcal{A}$ and $e$ an edge, we define $I(s, e) = \{\tau \in \mathbb{R}_+ \mid s \xrightarrow{\tau,e} s'\}$ and $I(s) = \bigcup_e I(s, e)$. Note that $I(s, e)$ is an interval, whereas $I(s)$ is a finite union of intervals. The timed automaton $\mathcal{A}$ is said *non-blocking* if, for every state $s$, $I(s) \neq \emptyset$.

### 2.2 The region automaton abstraction

The well-known region automaton construction [AD94] is an abstraction of timed automata which can be used for verifying many properties, for instance regular untimed properties.

Let $\mathcal{A}$ be a timed automaton. Define $M$ as the largest constant to which clocks are compared in guards or invariants of $\mathcal{A}$. Two clock valuations $\nu$ and $\nu'$ are said *region-equivalent* (written $\nu \approx \nu'$) whenever the following conditions hold:

- $\lfloor \nu(x) \rfloor = \lfloor \nu'(x) \rfloor$ or $\nu(x), \nu'(x) > M$, for all $x \in X$;
- $\{\nu(x)\} = 0$ iff $\{\nu'(x)\} = 0$, for all $x \in X$ with $\nu(x) \le M$;
- $\{\nu(x)\} \le \{\nu(y)\}$ iff $\{\nu'(x)\} \le \{\nu'(y)\}$, for all $x, y \in X$ with $\nu(x), \nu(y) \le M$.

where, $\lfloor\cdot\rfloor$ denotes the integral part, and $\{\cdot\}$ denotes the fractional part.

This equivalence relation on clock valuations has a finite (exponential) index, and extends to the states of $\mathcal{A}$, saying that $(\ell, \nu) \approx (\ell', \nu')$ iff $\ell = \ell'$ and $\nu \approx \nu'$. We use $[\nu]$ (resp. $[(\ell, \nu)]$) to denote the equivalence class to which $\nu$ (resp. $(\ell, \nu)$) belongs. A *region* is an equivalence class of valuations. The set of all the regions is denoted by $R$. If $r$ is a region, we denote by $\mathsf{cell}(r)$ the smallest guard defined with constants smaller than $M$, and which contains $r$. We denote by $\mathsf{cell}(R)$ the set of all the $\mathsf{cell}(r)$.

The original region automaton [AD94] is a finite automaton which is the quotient of the timed transition system $T_{\mathcal{A}}$ by the equivalence relation $\approx$. Here, we use a slight modification of the original construction, which is still a timed automaton, but which satisfies very strong properties.

**Definition 2.** *Let $\mathcal{A} = (L, X, E, \mathcal{I}, \mathcal{L})$ be a timed automaton. The* region automaton *of $\mathcal{A}$ is the timed automaton $R(\mathcal{A}) = (Q, X, T, \kappa, \lambda)$ such that:*

- *$Q = L \times R$;*
- *$\kappa((\ell, r)) = \mathcal{I}(\ell)$, and $\lambda((\ell, r)) = \mathcal{L}(\ell)$ for every $(\ell, r) \in L \times R$;*
- *$T \subseteq (Q \times \mathit{cell}(R) \times 2^X \times Q)$, and $(\ell, r) \xrightarrow{\mathit{cell}(r''),e,Y} (\ell', r')$ is in $T$ iff $e = \ell \xrightarrow{g,Y} \ell'$ is in $E$, and there exists $\nu \in r$, $\tau \in \mathbb{R}_+$ with $(\ell, \nu) \xrightarrow{\tau,e} (\ell', \nu')$, $\nu + \tau \in r''$, and $\nu' \in r'$.*

We recover the usual region automaton of [AD94] by labelling the transitions "$e$" instead of "$\mathsf{cell}(r''), e, Y$", and by interpreting $\mathsf{R}(\mathcal{A})$ as a finite automaton. However, the above timed interpretation satisfies strong timed bisimulation properties that we do not detail here (we assume the reader is familiar with this construction). To every finite path $\pi((\ell, \nu), e_1 \dots e_n)$ in $\mathcal{A}$ corresponds a finite set of paths in $\pi(((\ell, [\nu]), \nu), f_1 \dots f_n)$ in $\mathsf{R}(\mathcal{A})$, each one corresponding to a choice in the regions that are crossed. If $\varrho$ is a run in $\mathcal{A}$, then we write $\iota(\varrho)$ 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 that timed automata are non-blocking, even though general timed automata could also be handled (but at a technical extra cost). In all examples, if a state has no outgoing transition, we implicitly add a self-loop on that state with no constraints, so that the automaton is non-blocking.

### 2.3 The logic LTL

We consider the linear-time temporal logic $\mathsf{LTL}$ [Pnu77] defined inductively as:

$$\mathsf{LTL} \ni \varphi ::= p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \neg\varphi \mid \varphi \,\mathbf{U}\, \varphi$$

where $p \in \mathsf{AP}$ is an atomic proposition. We use classical shorthands like $\mathtt{tt} \overset{\text{def}}{=} p \vee \neg p$, $\mathtt{ff} \overset{\text{def}}{=} p \wedge \neg p$, $\varphi_1 \Rightarrow \varphi_2 \overset{\text{def}}{=} \neg\varphi_1 \vee \varphi_2$, $\mathbf{F}\,\varphi \overset{\text{def}}{=} \mathtt{tt}\,\mathbf{U}\,\varphi$, and $\mathbf{G}\,\varphi \overset{\text{def}}{=} \neg\mathbf{F}\,(\neg\varphi)$. We assume the reader is familiar with the semantics of $\mathsf{LTL}$, that we interpret here on infinite runs of a timed automaton.

### 2.4 Largeness, meagerness and the Banach-Mazur game

We assume the reader is familiar with basic notions of topology (see *e.g.* [Mun00]). However, we recall the more elaborate notions of *meagerness* and *largeness*. If $(A, \mathcal{T})$ is a topological space, a set $B \subseteq A$ is *nowhere dense* if the interior of the closure of $B$ is empty. A set is *meager* if it is a countable union of nowhere dense sets, and a set is *large* if its complement is meager. These notions have very nice characterizations in terms of Banach-Mazur games.

**Definition 3 (Banach-Mazur game).** *Let $(A, \mathcal{T})$ be a topological space and $\mathcal{B}$ be a family of subsets of $A$ satisfying the two following properties:*

- *for all $B \in \mathcal{B}$, $\mathring{B} \neq \emptyset$, and*
- *given $O$ a non-empty open set of $A$, there exists $B \in \mathcal{B}$ such that $B \subseteq O$.*

*Fix $C$ a subset of $A$. Two players alternate their moves: Player 1 starts and chooses an element $B_1$ of $\mathcal{B}$; Player 2 then responds by choosing an element $B_2$ of $\mathcal{B}$ such that $B_1 \supseteq B_2$; Then Player 1 responds by choosing $B_3$ in $\mathcal{B}$ such that $B_2 \supseteq B_3$, and so on. This way, they build a decreasing sequence of sets $B_i$:*

$$A \supseteq B_1 \supseteq B_2 \supseteq B_3 \cdots$$

*where the $B_{2i+1}$'s (resp. $B_{2i}$'s) are the choices of Player 1 (resp. Player 2), for $i \in \mathbb{N}$. Player 1 wins the game if the intersection of all $B_i$'s intersects $C$, i.e.,*

$$\bigcap_{i=1}^{\infty} B_i \cap C \neq \emptyset \,.$$

*Otherwise, Player 2 wins the game.*

Banach-Mazur games are not always determined, even for simple topological spaces (see [Oxt57, Remark 1]). Still a natural question is to know when the players have winning strategies. The following result gives a partial answer and relates Banach-Mazur games and meagerness:

**Theorem 4 (Banach-Mazur [Oxt57]).** *Player 2 has a winning strategy in the Banach-Mazur game with target set $C$ if and only if $C$ is meager.*

## 3 Probabilistic and Topological Semantics for Timed Automata

In [BBB$^+$07], we defined two relaxed semantics for LTL over finite runs of timed automata: the *almost-sure* semantics, based on probabilities, and the *large* semantics, based on the topological notion of largeness. In this section, we extend, in a natural way these semantics and the notion of dimension to infinite runs of timed automata.

### 3.1 A probabilistic semantics for LTL

Let $\mathcal{A}$ be a timed automaton. As in [BBB$^+$07], 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 $\mu_s$ over delays in $\mathbb{R}_+$ (equipped with the standard Borel $\sigma$-algebra) must satisfy several requirements. A first series, is denoted $(\star)$ in the sequel:

- $\mu_s(I(s)) = \mu_s(\mathbb{R}_+) = 1$,[4]
- Denoting $\lambda$ the Lebesgue measure, if $\lambda(I(s)) > 0$, $\mu_s$ is equivalent[5] to $\lambda$ on $I(s)$; Otherwise, $\mu_s$ is equivalent on $I(s)$ to the uniform distribution over points of $I(s)$.

This last condition denotes some kind of fairness w.r.t. enabled transitions, in that we cannot disallow one transition by putting a probability 0 to delays enabling that transition.

For technical reasons, we also ask for additional requirements (denoted $(\dagger)$):

- $(s,a,b) \mapsto \mu_s\big(\{d \mid s+d \in [a,b]\}\big)$ is continuous on $\{(s,a,b) \mid \exists e \text{ s.t. } [a,b] \subseteq I(s,e)\}$;
- If $s' = s+t$ for some $t \geq 0$, and $0 \notin I(s+t',e)$ for every $0 \leq t' \leq t$, then $\mu_s(I(s,e)) \leq \mu_{s'}(I(s',e))$;
- There is $0 < \lambda_0 < 1$ s.t. for every state $s$ with $I(s)$ unbounded, $\mu_s([0,1/2]) \leq \lambda_0$.

*Remark 5.* The three last requirements are technical and needed to deal with infinite behaviours, but they are natural and easily satisfiable. For instance, a timed automaton equipped with uniform (resp. exponential[6]) distributions on bounded (resp. unbounded) intervals satisfy these conditions. If we assume exponential distributions on unbounded intervals, the very last requirement corresponds to the bounded transition rate condition in [DP03], required to have reasonable and realistic behaviours.

For every state $s$ of $\mathcal{A}$, we also assume a probability distribution $p_s$ over edges, such that for every edge $e$, $p_s(e) > 0$ iff $e$ is enabled in $s$ (*i.e.*, $s \xrightarrow{e} s'$ for some $s'$). Moreover, to simplify, 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 $s$, for 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 define a measure over finite symbolic paths from state $s$ as

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

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

[6] With bounded transition rates, see [DP03].

where $s \xrightarrow{t} (s+t) \xrightarrow{e_1} s_t$, and we initialize with $\mathbb{P}_{\mathcal{A}}(\pi(s)) = 1$.[7] 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 is needed to deal with Zeno behaviours (see Section 5). 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 $\pi$ is a finite (constrained) symbolic path. Finally we extend $\mathbb{P}_{\mathcal{A}}$ in a standard and unique way to the $\sigma$-algebra generated by these cylinders, that we note $\Omega^s_{\mathcal{A}}$.

**Proposition 6.** *Let $\mathcal{A}$ be a timed automaton. For every state $s$, $\mathbb{P}_{\mathcal{A}}$ is a probability measure over $(\mathsf{Runs}(\mathcal{A}, s), \Omega^s_{\mathcal{A}})$.*

The proof of the proposition also justifies the construction for the probability measure $\mathbb{P}_{\mathcal{A}}$ and is given in the appendix (see page i).

*Example 7.* Consider the timed automaton $\mathcal{A}$ depicted on Fig. 1, and assume for all states both uniform distributions over delays and discrete moves. If $s_0 = (\ell_0, 0)$ is the initial state, then

$$\mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi(s_0, e_1 e_1))) = \mathbb{P}_{\mathcal{A}}(\pi(s_0, e_1 e_1)) = \frac{1}{4} \quad \text{and} \quad \mathbb{P}_{\mathcal{A}}(\pi(s_0, {e_1}^{\omega})) = 0\,.$$

- Indeed,

$$\begin{aligned} \mathbb{P}_{\mathcal{A}}(\pi(s_0, e_1 e_1)) &= \int_{t \in I(s, e_1)} p_{s+t}(e_1)\, \mathbb{P}_{\mathcal{A}}(\pi((\ell_0, t), e_1))\, \mathrm{d}\mu_{s_0}(t) \\ &= \int_0^1 \frac{1}{2}\, \mathbb{P}_{\mathcal{A}}(\pi((\ell_0, t), e_1))\, \mathrm{d}\lambda(t) \\ &= \frac{1}{2} \int_0^1 \Big( \int_{t \in I(s_t, e_2)} p_{s_t + u}(e_2)\, \mathbb{P}_{\mathcal{A}}(\pi((\ell_1, u)))\, \mathrm{d}\mu_{s_t}(u) \Big)\, \mathrm{d}\lambda(t) \\ &= \frac{1}{2} \int_0^1 \Big( \int_t^1 \frac{1}{2} \frac{1}{1-t}\, \mathrm{d}\lambda(u) \Big)\, \mathrm{d}\lambda(t) = \frac{1}{4} \end{aligned}$$

because $\mu_{s_0} = \lambda$ (resp. $\mu_{s_t} = \frac{\lambda}{1-t}$) is the uniform distribution over $[0,1]$ (resp. over $[t,1]$).

[7] In [BBB$^+$07] the definition was slightly different since we wanted the measure of all finite paths to be 1. We therefore used a normalisation factor 1/2 so that the measure of all paths of length $i$ were $1/2^{i+1}$.

- In a similar way we can show that $\mathbb{P}_{\mathcal{A}}(\pi(s_0, {e_1}^n)) = \frac{1}{2^n}$, for $n \in \mathbb{N}$; and thus conclude that $\mathbb{P}_{\mathcal{A}}(\pi(s_0, {e_1}^\omega)) = 0$. △

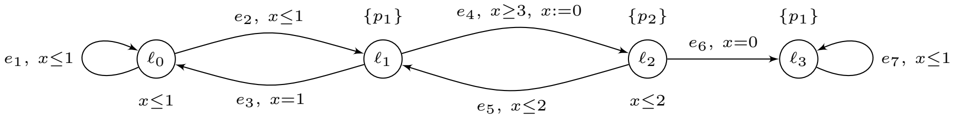

**Fig. 1.** A running example

We have seen in [BBB$^+$07] how to transfer probabilities from $\mathcal{A}$ to $\mathsf{R}(\mathcal{A})$, and proved the correctness of the transformation. Under the same hypotheses (for every state $s$ in $\mathcal{A}$, $\mu_s^{\mathcal{A}} = \mu_{\iota(s)}^{\mathsf{R}(\mathcal{A})}$, and for every $t \in \mathbb{R}_+$ $p_{s+t}^{\mathcal{A}} = p_{\iota(s)+t}^{\mathsf{R}(\mathcal{A})}$ ) this correctness still holds in our case by definition of the probability measure:

**Lemma 8.** *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 $\mathbb{P}_{\mathcal{A}}(S) = \mathbb{P}_{\mathsf{R}(\mathcal{A})}(\iota(S))$.*

*Proof.* The two measures coincide on finite constraints paths (see [BBB$^+$07, Lemma 7]), thus on cylinders, and finally on any measurable set of infinite runs. □

We can therefore lift results proved on $\mathsf{R}(\mathcal{A})$ to $\mathcal{A}$. In the sequel, we write $\mathcal{A} = \mathsf{R}(\mathcal{A})$ when we consider a region automaton rather than a general timed automaton.

Given an infinite symbolic path $\pi$ and an $\mathsf{LTL}$ formula $\varphi$, either all concretizations of $\pi$ (*i.e.*, concrete runs $\varrho \in \pi$) satisfy $\varphi$, or they all do not satisfy $\varphi$. Hence, the set $\{\varrho \in \mathsf{Runs}(\mathcal{A}, s_0) \mid \varrho \models \varphi\}$ is measurable (in $\Omega_{\mathcal{A}}^{s_0}$), as it is an $\omega$-regular property [Var85]. In the sequel, we write $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ for $\mathbb{P}_{\mathcal{A}}\{\varrho \in \mathsf{Runs}(\mathcal{A}, s_0) \mid \varrho \models \varphi\}$.

**Definition 9.** *Let $\varphi$ be an $\mathsf{LTL}$ formula and $\mathcal{A}$ a timed automaton. We say that $\mathcal{A}$* almost-surely satisfies *$\varphi$ from $s_0$, and we then write $\mathcal{A}, s_0 \mathrel{|}\!\!\approx_{\mathbb{P}} \varphi$, whenever $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi) = 1$. The* almost-sure model-checking problem *asks, given $\mathcal{A}$, $\varphi$ and $s_0$, whether $\mathcal{A}, s_0 \mathrel{|}\!\!\approx_{\mathbb{P}} \varphi$.*

*Example 10.* Consider the timed automaton $\mathcal{A}$ of Fig. 1 again with both uniform distributions over delays and discrete moves in all states and initial state $s_0 = (\ell_0, 0)$. Then, $\mathcal{A}, s_0 \mathrel{|}\!\!\approx_{\mathbb{P}} \mathbf{F}\,(p_1 \wedge \mathbf{G}\,(p_1 \Rightarrow \mathbf{F}\,p_2))$. Indeed, in state $(\ell_0, \nu)$ with $0 \le \nu \le 1$, the probability of firing $e_2$ (after some delay) is always $1/2$ (guards of $e_1$ and $e_2$ are the same, there is thus a uniform distribution over both edges), thus the location $\ell_1$ is reached with probability 1. In $\ell_1$, the transition $e_3$ will unlikely happen, because its guard $x = 1$ is much too "small" compared to the guard $x \ge 3$ of the transition $e_4$. The same phenomenon arises in location $\ell_2$ between the transitions $e_5$ and $e_6$. In conclusion, the runs of the timed automaton $\mathcal{A}$ (from $s_0$) are almost surely following sequences of transitions of the form ${e_1}^* e_2 (e_4 e_5)^\omega$. Hence, with probability 1, the formula $\mathbf{F}\,(p_1 \wedge \mathbf{G}\,(p_1 \Rightarrow \mathbf{F}\,p_2))$ is satisfied. Note that the previous formula is not satisfied with the classical $\mathsf{LTL}$ semantics. Indeed several counter-examples to the satisfaction of the formula can be found: '*staying in $\ell_0$ forever*', '*reaching $\ell_3$*', etc... All these counter-examples are unlikely and vanish thanks to our probabilistic semantics. △

Although the values $\mathbb{P}_{\mathcal{A}}(s_0 \models \varphi)$ depend on the chosen weights $p_s(e)$ and measures $\mu_s$, we will see that for one-clock timed automata the almost-sure satisfaction relation is not affected by the choice of the weights and distributions. This will be crucial for the decidability of the almost-sure model checking problem. The way to establish this result is to prove the equivalence of the almost-sure semantics and a topological semantics, which is defined on the basis of the so-called dimension of symbolic paths.

### 3.2 The dimension, a tool to almost-surely analyze timed automata

In [BBB$^+$07], we introduced a notion of dimension for finite constrained symbolic paths. Intuitively, a path is of *defined dimension* if it corresponds to a polyhedron of maximal dimension (in the space induced by the automaton). Formally, let $\pi_{\mathcal{C}} = \pi_{\mathcal{C}}(s, e_1 \dots e_n)$ be a constrained path of a timed automaton $\mathcal{A}$. We define its associated polyhedron as follows:

$$\mathsf{Pol}(\pi_{\mathcal{C}}) = \left\{ (\tau_i)_{1 \le i \le n} \in (\mathbb{R}_+)^n \mid s \xrightarrow{\tau_1, e_1} s_1 \cdots \xrightarrow{\tau_n, e_n} s_n \in \pi_{\mathcal{C}}(s, e_1 \dots e_n) \right\}.$$

For each $0 \le i \le n$, we write $\mathcal{C}_i$ for the constraint induced by the projection of $\mathsf{Pol}(\pi_{\mathcal{C}})$ over the $i$ first coordinates, with the convention that $\mathcal{C}_0$ is true. We say that the dimension of $\pi_{\mathcal{C}}$ is *undefined*, denoted $\dim_{\mathcal{A}}(\pi_{\mathcal{C}}) = \bot$, whenever there exists some index $1 \le i \le n$ with

$$\dim\big(\mathsf{Pol}\big(\pi_{\mathcal{C}_i}(s, e_1 \dots e_i)\big)\big) < \dim\big(\cup_e \, \mathsf{Pol}\big(\pi_{\mathcal{C}_{i-1}}(s, e_1 \dots e_{i-1} e)\big)\big).$$

Otherwise we say that the dimension of $\pi_{\mathcal{C}}$ is *defined*, denoted $\dim_{\mathcal{A}}(\pi_{\mathcal{C}}) = \top$.

The notion of dimension naturally extends to infinite symbolic paths: If $\pi = \pi(s, e_1 e_2 \dots)$ is an infinite symbolic path, its *dimension* is

$$\dim_{\mathcal{A}}(\pi) \;=\; \lim_{n \to \infty} \dim_{\mathcal{A}}(\pi(s, e_1 \dots e_n)).$$

*Remark 11.* If for some index $n$, $\dim_{\mathcal{A}}(\pi(s, e_1 e_2 \dots e_n)) = \bot$, then for every index $m \ge n$, $\dim_{\mathcal{A}}(\pi(s, e_1 e_2 \dots e_m)) = \bot$. This is a consequence of [BBB$^+$07, Lemma 22].

*Example 12.* On the automaton $\mathcal{A}$ of Fig. 1 with $s_0 = (\ell_0, 0)$, $\dim_{\mathcal{A}}(\pi(s_0, {e_1}^\omega)) = \top$ and $\dim_{\mathcal{A}}(\pi(s_0, e_1 (e_2 e_3)^\omega)) = \bot$.

- Let us first consider the infinite path $\pi(s_0, {e_1}^\omega)$, and show that all its finite prefixes $\pi(s_0, e_1^n)$ have defined dimension.

  $$\mathsf{Pol}(\pi(s_0, {e_1}^n)) = \{(\tau_1, \ldots, \tau_n) \in (\mathbb{R}_+)^n \mid (0 \le \tau_1 \le 1) \wedge \cdots \wedge (0 \le \tau_1 + \cdots + \tau_n \le 1)\}\,,$$

  thus clearly enough $\dim(\mathsf{Pol}(\pi(s_0, {e_1}^n))) = n$. Moreover we have that $\mathsf{Pol}(\pi(s_0, {e_1}^{n-1}))$ is equal to the projection of $\mathsf{Pol}(\pi(s_0, {e_1}^n))$ on the $n-1$ first coordinates; we denote by $\mathcal{C}_{n-1}$ the constraint induced by this projection. In particular, we have that: $\pi_{\mathcal{C}_{n-1}}(s_0, {e_1}^{n-1}, e) = \pi(s_0, {e_1}^{n-1}, e)$. We can now conclude that, for $1 \le i \le n$:

  $$\begin{aligned}\dim\Big(\bigcup_e \mathsf{Pol}\big(\pi_{\mathcal{C}_{i-1}}(s_0, {e_1}^{i-1}e)\big)\Big) &= \dim\Big(\bigcup_e \mathsf{Pol}\big(\pi(s_0, {e_1}^{i-1}e)\big)\Big) \\ &= \dim\big(\mathsf{Pol}\big(\pi(s_0, {e_1}^{i})\big)\big) = i,\end{aligned}$$

  proving that $\dim(\pi(s_0, {e_1}^n)) = \top$, for $n \in \mathbb{N}$, and thus $\dim(\pi(s_0, {e_1}^\omega)) = \top$.
- Let us now consider the infinite path $\pi(s_0, e_1(e_2e_3)^\omega)$. In order to show that its dimension is undefined, we exhibit a finite prefix of undefined dimension. First notice that:

  $$\begin{aligned}\mathsf{Pol}(\pi(s_0, e_1e_2e_3)) &= \{(\tau_1, \tau_2, \tau_3) \mid (0 \le \tau_1 \le 1) \wedge (0 \le \tau_1 + \tau_2 \le 1) \wedge (\tau_1 + \tau_2 + \tau_3 = 1)\}\,, \\ \mathsf{Pol}(\pi(s_0, e_1e_2e_4)) &= \{(\tau_1, \tau_2, \tau_3) \mid (0 \le \tau_1 \le 1) \wedge (0 \le \tau_1 + \tau_2 \le 1) \wedge (\tau_1 + \tau_2 + \tau_3 \ge 3)\}\,,\end{aligned}$$

  thus $\dim(\mathsf{Pol}(\pi(s_0, e_1e_2e_3))) = 2 < \dim(\mathsf{Pol}(\pi(s_0, e_1e_2e_4))) = 3$ which implies that $\dim(\pi(s_0, e_1e_2e_3) = \perp$. The definiion of the dimension clearly implies that all extensions of paths of undefined dimension have undefined dimension too. Hence we conclude that $\dim(\pi(s_0, e_1(e_2e_3)^\omega) = \perp$. △

In the context of finite paths, a symbolic path has probability 0 iff it has an undefined dimension. In the context of infinite paths, this is not quite true as infinite paths with defined dimension can have probability 0, like $\pi(s_0, e_1^\omega)$ in the automaton of Fig. 1. However, writing $\mathbb{P}_\mathcal{A}(s \models \mathsf{dim\_undef})$ for $\mathbb{P}_\mathcal{A}\{\varrho \in \mathsf{Runs}(\mathcal{A}, s) \mid \dim_\mathcal{A}(\varrho) = \perp\}$, the following holds:

**Lemma 13.** *If $\mathcal{A}$ is a timed automaton, for every state $s$ in $\mathcal{A}$,*

$$\mathbb{P}_\mathcal{A}(s \models \mathsf{dim\_undef}) = 0\,.$$

*Proof.* Let $\pi$ be an infinite path in $\mathcal{A}$ with undefined dimension. By definition of the dimension for infinite paths, $\pi$ admits a finite prefix[8] $\pi_1$ that has undefined dimension too. Moreover, any continuation of $\pi_1$ also has an undefined dimension. Therefore, the whole cylinder set generated by $\pi_1$ is composed of infinite paths of undefined dimension. By definition, $\mathbb{P}_\mathcal{A}(\mathsf{Cyl}(\pi_1))$ is the probability of the finite

[8] As $\mathcal{A} = \mathsf{R}(\mathcal{A})$, projections and prefixes match.

symbolic path $\pi_1$, which is equal to 0 thanks to the equivalence for finite paths between zero-probability and undefined dimension (see [BBB$^+$07]).

The set of infinite paths which have undefined dimension can be written has the denumerable union of cylinders generated by finite prefixes with undefined dimension:

$$\{\varrho \in \mathsf{Runs}(\mathcal{A}, s) \mid \dim(\pi_\varrho) = \bot\} = \bigcup_{\substack{\pi=\pi(s,e_1\ldots e_n)\\ \text{s.t. } \dim(\pi)=\bot}} \mathsf{Cyl}(\pi).$$

Hence

$$\mathbb{P}_\mathcal{A}(s \models \mathsf{dim\_undef}) \leq \sum_{\substack{\pi=\pi(s,e_1\ldots e_n)\\ \text{s.t. } \dim(\pi)=\bot}} \mathbb{P}_\mathcal{A}(\mathsf{Cyl}(\pi)) \quad = 0\,.$$

□

### 3.3 A topological semantics for LTL

Let $\mathcal{A}$ be a timed automaton, and $s$ be a state of $\mathcal{A}$. Let $\mathcal{T}^s_\mathcal{A}$ be the topology over the set of runs of $\mathcal{A}$ starting in $s$ defined with the following basic opens sets: either the set $\mathsf{Runs}(\mathcal{A}, s)$, or the cylinders $\mathsf{Cyl}(\pi_\mathcal{C})$ where $\pi_\mathcal{C} = \pi_\mathcal{C}(s, e_1e_2 \ldots e_n)$ is a finite constrained symbolic path of $\mathcal{A}$ such that: (*i*) $\dim(\pi_\mathcal{C}) = \top$, (*ii*) $\mathcal{C}$ is convex (and Borel-measurable), and (*iii*) $\mathsf{Pol}(\pi_\mathcal{C})$ is open in $\mathsf{Pol}(\pi)$ for the classical topology on $\mathbb{R}^n$.

We first prove that our topological space is a *Baire space*:[9] indeed, in non Baire spaces, the notions of largeness and meagerness do not always make sense. For instance, in $\mathbb{Q}$ with the classical topology, every set is both meager and large. Hence negation would have little meaning in our topological satisfaction. In Baire spaces, however, if a set is large its complement is not.

**Proposition 14.** *Let $\mathcal{A}$ be a timed automaton. For every state $s$ of $\mathcal{A}$, the topological space $(\mathsf{Runs}(\mathcal{A}, s), \mathcal{T}^s_\mathcal{A})$ is a Baire space.*

The proof of Proposition 14 heavily relies on the Banach-Mazur game but is not a consequence of the same result for finite runs [BBB$^+$07].

*Proof.* To prove that $(\mathsf{Runs}(\mathcal{A}, s), \mathcal{T}^s_\mathcal{A})$ is a Baire space, we prove that every non-empty basic open set in $\mathcal{T}^s_\mathcal{A}$ is not meager. Let $\mathsf{Cyl}(\pi_\mathcal{C}(s, e_1 \ldots e_n))$ be a basic open set, where $\pi_\mathcal{C}(s, e_1 \ldots e_n)$ is a finite constrained symbolic path. Using Banach-Mazur games (see page 6 or [Oxt57]), we prove that $\mathsf{Cyl}(\pi_\mathcal{C}(s, e_1 \ldots e_n))$ is not meager by proving that Player 2 does not have a winning strategy for the Banach-Mazur game playing with basic open sets and where the goal set is $C = \mathsf{Cyl}(\pi_\mathcal{C}(s, e_1 \ldots e_n))$.

Player 1 starts by choosing a set $B_1 = \mathsf{Cyl}(\pi_\mathcal{C}(s, e_1 \ldots e_n))$. Then Player 2 picks some basic open set $B_2 = \mathsf{Cyl}(\pi_{\mathcal{C}^2}(s, e_1 \ldots e_n \ldots e_{n_1}))$ such that $B_1 \supseteq B_2$.

[9] Recall that a topological space $(A, \mathcal{T})$ is a *Baire space* if every non-empty open set in $\mathcal{T}$ is not meager.

Let us now explain how Player 1 can build her move in order to avoid to reach the empty set. Since $B_2$ is an open set, we have that $(i)$ $\dim(\pi_{\mathcal{C}^2}) = \top$ and $(ii)$ $\mathsf{Pol}(\pi_{\mathcal{C}^2}(s, e_1 \dots e_{n_1}))$ is open in $\mathsf{Pol}(\pi(s, e_1 \dots e_{n_1})) \subseteq \mathbb{R}_+^{n_1}$. Since the topology on $\mathsf{Pol}(\pi(e_1 \dots e_{n_1}))$ is induced from a distance, we know that there exists a closed, bounded and convex set denoted $K_1$ such that $\overset{\circ}{K_1} \neq \emptyset$ and $K_1 \subseteq \mathsf{Pol}(\pi_{\mathcal{C}^2}(s, e_1 \dots e_{n_1}))$. Let $\mathcal{D}^1$ be the set of constraints associated with $K_1$, we clearly have that the cylinder $\mathsf{Cyl}(\pi_{\mathcal{D}^1}(s, e_1 \dots e_{n_1}))$ is included in $B_2$. Let $O$ be a convex open set included in $K_1$ and $\mathcal{C}^3$ be the set of constraints associated with $O$. Applying Corollary D of the research report correponding to [BBB$^+$07], we know that $\dim_{\mathcal{A}}(\pi_{\mathcal{C}^3}(s, e_1 \dots e_{n_1})) = \top$. Hence clearly enough, we have that $\mathsf{Cyl}(\pi_{\mathcal{C}^3}(s, e_1 \dots e_{n_1}))$ is an open set. Player 1's move will be to take $B_3 = \mathsf{Cyl}(\pi_{\mathcal{C}^3}(s, e_1 \dots e_{n_1}))$. By iterating the same process for the strategy of Player 1, we obtain the following sequence:

$$B_1 \supseteq B_2 \supseteq \mathsf{Cyl}(\pi_{\mathcal{D}^1}) \supseteq B_3 \supseteq B_4 \supseteq \mathsf{Cyl}(\pi_{\mathcal{D}^2}) \supseteq \dots \supseteq B_{2i-1} \supseteq B_{2i} \supseteq \mathsf{Cyl}(\pi_{\mathcal{D}^i}) \supseteq \cdots$$

where for each $i$, $K_i = \mathsf{Pol}(\pi_{\mathcal{D}^i})$ is a closed and bounded subset of $\mathsf{Pol}(\pi(e_1, \dots, e_{n_i})) \subseteq \mathbb{R}_+^{n_i}$ (where the $n_i$'s form a non-decreasing sequence of $\mathbb{N}$). We then have that:

$$\bigcap_{i=1}^{\infty} B_i \;=\; \bigcap_{i=1}^{\infty} \mathsf{Cyl}(\pi_{\mathcal{D}^i})\,.$$

We would like to guarantee that the above intersection in non-empty This is not completely straightforward since the polyhedra $K_i = \mathsf{Pol}(\pi_{\mathcal{D}^i})$ belong to different powers of $\mathbb{R}_+$. We distinguish between two cases:

- either the sequence $(n_i)_{i\geq 1}$ diverges to $+\infty$. In that case, we will embed $\bigcap_{i=1}^{\infty} K_i$ into a compact set of $\mathbb{R}_+^{\mathbb{N}}$. We first define

  $$\widetilde{K_j} = \mathsf{Proj}_{\{n_{j-1}+1,\dots,n_j\}} K_j \quad \text{and} \quad \widetilde{K} = \prod_{j \geq 1} \widetilde{K_j}\,.$$

  Note that $\widetilde{K_j}$ is a compact set, since it is the projection of a compact set. Each $K_i$ can naturally be embedded in $\widetilde{K}$ by considering the sets $K'_i$ defined by

  $$K'_i = K_i \times \prod_{j > i} \widetilde{K_j}\,.$$

  The decomposition is illustrated on Figure 2. The $K'_i$'s form a nested chain of closed sets of $\widetilde{K}$. By Tychonoff's theorem, $\widetilde{K}$ is compact. Hence we can ensure that $\bigcap_{i=1}^{\infty} K'_i$ is non-empty (Heine-Borel Theorem). Take a sequence $(\tau_j)_{j\geq 1}$ in $\bigcap_{i=1}^{\infty} K'_i$. Each subsequence $(\tau_j)_{1\leq j\leq n_i}$ straightforwardly belongs to $K_i$. Hence, the run $s \xrightarrow{\tau_1, e_1} s_1 \xrightarrow{\tau_2, e_2} s_2 \dots$ is in $\bigcap_{i=1}^{\infty} B_i$, which completes the proof in this case.
- either the sequence $(n_i)_{i\geq 1}$ is upper bounded. In that case, we embed $\bigcap_{i=1}^{\infty} K_i$ into a compact set of $\mathbb{R}_+^N$ where $N = \lim_{i\to+\infty}$. We let the details to the reader, as they are very similar to (and easier than) the previous case. □

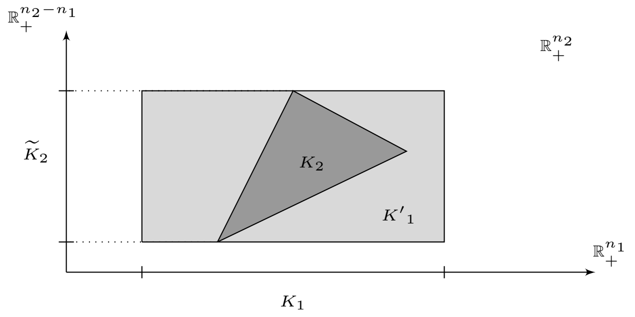

**Fig. 2.** The decomposition of the $K_i$'s

We can now define a topological semantics for LTL based on the (topological) notion of largeness.

**Definition 15.** *Let $\varphi$ be an* LTL *formula and $\mathcal{A}$ a timed automaton. We say that $\mathcal{A}$* largely satisfies *$\varphi$ from $s_0$, and we write $\mathcal{A}, s_0 \mathrel{|\!\!\approx}_{\mathcal{T}} \varphi$, if $\{\varrho \in \textsf{Runs}(\mathcal{A}, s_0) \mid \varrho \models \varphi\}$ is topologically large. The* large model-checking problem *asks, given $\mathcal{A}$, $\varphi$ and $s_0$, whether $\mathcal{A}, s_0 \mathrel{|\!\!\approx}_{\mathcal{T}} \varphi$.*

*Example 16.* On the timed automaton $\mathcal{A}$ of Fig. 1 with initial state $s_0 = (\ell_0, 0)$, $\mathcal{A}, s_0 \mathrel{|\!\!\approx}_{\mathcal{T}} \mathbf{F}\,(p_1 \wedge \mathbf{G}\,(p_1 \Rightarrow \mathbf{F}\,p_2))$.

In order to prove that formula $\varphi \equiv \mathbf{F}\,(p_1 \wedge \mathbf{G}\,(p_1 \Rightarrow \mathbf{F}\,p_2))$ is largely satisfied, we show that the set $C \overset{\text{def}}{=} \{\pi(s_0, {e_1}^i e_2 (e_4 e_5)^\omega) | i \in \mathbb{N}\}$ is large. Indeed, each run of $C$ satisfies $\varphi$, and thus, if $C$ is large then $\mathcal{A}, s_0 \mathrel{|\!\!\approx}_{\mathcal{T}} \varphi$, since largeness is closed under subsumption. To prove that $C$ is large (or equivalently that its complement is meager) we use a Banach-Mazur game [Oxt57], and show that Player 2 has a strategy to avoid the complement of $C$, hence to reach $C$. The game is played with the basic open sets of $(\textsf{Runs}(\mathcal{A}, s_0))$. The strategy of Player 2 is as follows:

- We assume Player 1 has chosen a cylinder $\textsf{Cyl}(\pi(s_0, {e_1}^{n_1}))$, for some $n_1 \in \mathbb{N}_0$ (if Player 1 leaves $\ell_0$ at her first move, we skip the first move of Player 2)
- Player 2 chooses $\textsf{Cyl}(\pi(s_0, {e_1}^{n_1} e_2))$,
- Notice that Player 1 is not allowed to extend the symbolic path $\pi(s_0, {e_1}^{n_1} e_2)$ with sequences of transitions including $e_3$ or $e_6$, since both symbolic paths $\pi(s_0, {e_1}^{n_1} e_2 e_3)$ and $\pi(s_0, {e_1}^{n_1} e_2 e_4 e_6)$ have undefined dimension. Thus she can only play moves of the form $\textsf{Cyl}(\pi(s_0, {e_1}^{n_1} (e_2 e_3)^{n_2}))$ or $\textsf{Cyl}(\pi(s_0, {e_1}^{n_1} e_2 (e_3 e_2)^{n_2}))$.
- Player 2 takes $\textsf{Cyl}(\pi(s_0, {e_1}^{n_1} (e_2 e_3)^{n_3}))$, with $n_3 > n_2$.

One can easily be convinced that by repeating infinitely often the two last moves, we will obtain a run of $C$, proving that Player 2 won the game and thus that $C$ is large.

Notice that both players could also play with constrained paths. This would not be interesting for Player 1, since it could only cause the intersection to be empty (in which case Player 2 wins as well). △

Although the topological spaces given by $\mathcal{A}$ and $\mathsf{R}(\mathcal{A})$ are not homeomorphic, the topologies in $\mathcal{A}$ and in $\mathsf{R}(\mathcal{A})$ somehow match, as stated by the next proposition. This allows to lift result from $\mathsf{R}(\mathcal{A})$ to $\mathcal{A}$.

**Proposition 17.** *Let $\mathcal{A}$ be a timed automaton, and $s$ a state of $\mathcal{A}$. Let $S \subseteq \mathsf{Runs}(\mathcal{A}, s)$. Then, $S$ is large in $(\mathsf{Runs}(\mathcal{A}, s), \mathcal{T}^s_{\mathcal{A}})$ iff $\iota(S)$ is large in $(\mathsf{Runs}(\mathsf{R}(\mathcal{A}), \iota(s)), \mathcal{T}^{\iota(s)}_{\mathsf{R}(\mathcal{A})})$.*

The proof of this proposition relies on the following technical lemma. Indeed, it is now sufficient to simulate a Banach-Mazur game from $\mathcal{A}$ to $\mathsf{R}(\mathcal{A})$ and *vice-versa* to get the expected result.

**Lemma 18.** *Let $\iota : \mathsf{Runs}_f(\mathcal{A}, s) \to \mathsf{Runs}_f(\mathsf{R}(\mathcal{A}), \iota(s))$ be the projection of finite runs $\varrho$ in $\mathcal{A}$ onto the region automaton (see page 5). Then $\iota$ is continuous, and for every non-empty open set $O \in \mathcal{T}^s_{\mathcal{A}}$, $\overset{\circ}{\widehat{\iota(O)}} \neq \emptyset$.*

*Proof.* Let us first proof that $\iota$ is continuous. Let $\mathsf{Cyl}(\pi_{\mathcal{C}})$ be a basic open set of $\mathcal{T}^{\iota(s)}_{\mathsf{R}(\mathcal{A})}$, we need to prove that $\iota^{-1}(\mathsf{Cyl}(\pi_{\mathcal{C}}))$ is an open set of $\mathcal{T}^s_{\mathcal{A}}$. One can easily be convinced that $\iota^{-1}(\mathsf{Cyl}(\pi_{\mathcal{C}})) = \mathsf{Cyl}(\iota^{-1}(\pi_{\mathcal{C}}))$. By [BBB+07, Lemma 16], we have that $\iota^{-1}(\pi_{\mathcal{C}})$ is a finite symbolic path with defined dimension whose polyhedron is open in its ambient space. Hence $\mathsf{Cyl}(\iota^{-1}(\pi_{\mathcal{C}}))$ is open, and $\iota$ is thus continuous.

Let us now prove that for every non-empty open set $\mathcal{O} \in \mathcal{T}^s_{\mathcal{A}}$, $\overset{\circ}{\widehat{\iota(\mathcal{O})}} \neq \emptyset$. Let $\mathsf{Cyl}(\pi_{\mathcal{C}})$ be a basic open set of $\mathcal{T}_{\mathcal{A}}$. Again using [BBB+07, Lemma 16], we obtain that $\iota(\pi_{\mathcal{C}})$ contains a symbolic path $\pi'$ with defined dimension whose polyhedron is open in its ambient space. Hence $\mathsf{Cyl}(\pi')$ is open and since $\mathsf{Cyl}(\pi') \subseteq \mathsf{Cyl}(\iota(\pi_{\mathcal{C}}))$, we obtain the desired result. □

*Remark 19.* Note that $\iota$ is not an homeomorphism from $\mathsf{Runs}_f(\mathcal{A}, s)$ to $\mathsf{Runs}_f(\mathsf{R}(\mathcal{A}), \iota(s))$ since $\iota^{-1} : \mathsf{Runs}_f(\mathsf{R}(\mathcal{A}), s) \to \mathsf{Runs}_f(\mathcal{A}, \iota^{-1}(s))$ is not continuous. Indeed, let us consider the automaton $\mathcal{A}$ of Fig. 3, with $s_0 = (\ell_0, 0)$. The set of runs $\mathcal{O} = \mathsf{Cyl}(\pi(s_0, e_1))$ is open in $\mathcal{T}^{s_0}_{\mathcal{A}}$ since $\pi(s_0, e_1)$ is a symbolic unconstrained path of defined dimension. However, $\iota(\pi(s_0, e_1)) = \mathsf{Cyl}(\pi(s_0, f_1)) \cup \mathsf{Cyl}(\pi(s_0, f_2))$ is not open in $\mathcal{T}^{\iota(s_0)}_{\mathsf{R}(\mathcal{A})}$ as $\dim_{\mathsf{R}(\mathcal{A})}(\pi(s_0, f_1)) = \perp$ and hence $\mathsf{Cyl}(\pi(s_0, f_1))$ is not a basic open. Thus $\iota(\mathcal{O})$ is not open and $\iota^{-1}$ is not continuous.

## 4 The Two Semantics Match

In the previous section, we defined two relaxed semantics for $\mathsf{LTL}$ over infinite runs in timed automata: the almost-sure satisfaction based on probabilistic interpretations

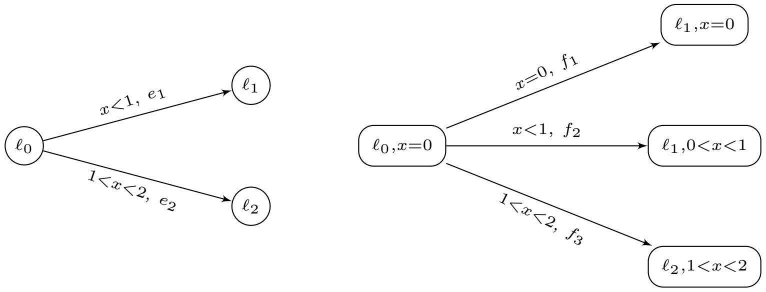

**Fig. 3.** An automaton and its region automaton

of delays and discrete choices, and the large satisfaction based on a topology defined on runs of the automaton. In this section, we prove that the two semantics match in the case of one-clock timed automata, and provide a decidability algorithm for the almost-sure (or equivalently large) $\mathsf{LTL}$ model-checking problem. It is however not a straightforward consequence of the same result for finite runs [BBB$^+$07]. It is indeed rather involved and requires the development of techniques mixing classical probabilistic techniques and strong properties of one-clock timed automata. Note that these techniques only apply in the one-clock framework!

We first recall a construction made in [BBB$^+$07] to decide the almost-sure model checking of $\mathsf{LTL}$ interpreted over finite paths. Any edge $e$ in $\mathsf{R}(\mathcal{A})$ is colored in red if $\mu_s(I(s,e)) = 0$, and in blue otherwise. Then, a finite path in $\mathsf{R}(\mathcal{A})$ has an undefined dimension iff it crosses a red edge. Hence, having a defined (or undefined) dimension for a path can be specified locally in $\mathsf{R}(\mathcal{A})$. We say that a blue (resp. red) edge has a defined (resp. undefined) dimension. We call $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ the restriction of $\mathsf{R}(\mathcal{A})$ to edges with defined dimension.

### 4.1 A notion of fairness

In the case of finite paths, if $\mathbb{P}_{\mathcal{A}}(s \models \varphi) = 1$, then only paths of undefined dimension may not satisfy $\varphi$. Unfortunately, this is in general wrong for infinite paths. Indeed, on the timed automaton $\mathcal{A}$ of Fig. 1, when starting from $s = (\ell_0, 0)$, location $\ell_1$ is clearly reached with probability 1. However the infinite path $\pi(s, {e_0}^{\omega})$ has defined dimension although it never reaches $\ell_1$. This kind of behaviours forces us to restrict our study to *fair* infinite paths, which is rather natural since probabilities and strong fairness are closely related [Pnu83].

Let $\mathcal{A} = \mathsf{R}(\mathcal{A})$ be a timed automaton. An infinite region path $q_0 \xrightarrow{e_1} q_1 \xrightarrow{e_2} q_2 \ldots$ in $\mathcal{A}$ is *fair* iff for every edge $e$ with defined dimension, if $e$ is enabled in infinitely many $q_i$ with $i \in \mathbb{N}$, then $e_i = e$ for infinitely many $i \in \mathbb{N}$. Note that region paths and symbolic paths are closely related, as we assume $\mathcal{A} = \mathsf{R}(\mathcal{A})$: to any non-empty symbolic path $\pi(s, e_1 e_2 \ldots)$, we associate a unique region path $q_0 \xrightarrow{e_1} q_1 \xrightarrow{e_2} q_2 \ldots$ with $s \in q_0$. Hence, we say that a symbolic path $\pi(s, e_1 e_2 \ldots)$ is fair whenever its corresponding region path is fair. Finally, we say that an infinite run $\varrho$ is fair whenever $\pi_\varrho$ is fair. Obviously, the set of fair infinite runs from $s$ is $\Omega^s_{\mathcal{A}}$-measurable, as fairness is an $\omega$-regular property over infinite paths. Writing $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair})$ for $\mathbb{P}_{\mathcal{A}}\{\varrho \in \mathsf{Runs}(\mathcal{A}, s) \mid \varrho \text{ is fair}\}$, we get the following property:

**Lemma 20.** *If $\mathcal{A}$ is a one-clock timed automaton, for every state $s$ in $\mathcal{A}$,*

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair}) = 1\,.$$

The proof of this lemma is fairly involved, we first briefly sketch the main steps of the proof, and will then give the complete proof.

($i$) We first prove that any edge with defined dimension is almost-surely taken infinitely often within a compact (for the value of the unique clock), provided it is enabled infinitely often within that compact.

($ii$) Then, restricting to runs with infinitely many resets, those paths will pass infinitely often in a given configuration (because we only have one clock, hence resetting the clock and going to location $q$ means entering the configuration $(q, 0)$). We can then apply the previous lemma, and get that any sequence of edges with defined dimension will be taken infinitely often with probability 1.

($iii$) Concerning the runs ending up in the unbounded region (with no more resets of the clock), we prove that the distributions over edges correspond ultimately to a finite Markov chain, and hence that these runs are fair with probability 1.

($iv$) Finally, restricting to runs ending up in a bounded region (with no more resets of the clock), only edges labelled with that precise region as a constraint can be enabled, and it will ultimately behave like a finite Markov chain, hence leading to the fairness property with probability 1.

***Proof of Lemma 20.*** The first step ($i$) relies on Lemma 21 below. A *subregion* of a region $q$ is a pair $(q, J)$ such that $J \subseteq q$ is an interval. If $s \in J$, we may write $s \in (q, J)$ as well. If $(q, J)$ and $(q', I)$ are subregions, we write $(q, J) \xrightarrow{e} (q', I)$ to express that $(q, v) \xrightarrow{e,t} (q', v')$ for some $v \in J$, $v' \in I$ and $t \in \mathbb{R}_+$. In the sequel to ease the reading, we will use $\mathsf{LTL}$-like notations, like $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond(q, J) \xrightarrow{e} (q', I) \mid \Box\Diamond(q, J))$, which denotes the conditional probability of the set of real runs $s_0 \xrightarrow{t_1,e_1} s_1 \xrightarrow{t_2,e_2} s_2 \cdots$ such that $s_0 = s$ and $\{s_i \xrightarrow{e_{i+1}} s_{i+1} \mid s_i \in J,\ e_{i+1} = e, \text{ and } s_{i+1} \in I\}$ is infinite, assuming that the set $\{s_i \mid s_i \in J\}$ is infinite. We will use other such notations, that we expect are sufficiently explicit to be understandable.

**Lemma 21.** *1. For every subregion $(q, J)$ of $q$ such that ($i$) $J$ is non-empty and open in $q$ (for the induced topology), and ($ii$) $\overline{J} \subseteq q$ is compact,*

*2. for every edge $e$ enabled in $q$ such that* $\dim_{\mathcal{A}}(e) \neq \bot$*,*

3. *for every subregion $(q', I)$ of $q'$ such that for every $s \in (q, J)$, $e(s) \cap I$ is non-empty and open in $q'$ (for the induced topology), where $e(s) = \{s' \mid \exists t \in \mathbb{R}_+ \ s.t.\ s \xrightarrow{t,e} s'\}$,*
4. *for every state $s$ of $\mathcal{A}$ such that $\mathbb{P}_\mathcal{A}(s, \Box\Diamond(q, J)) > 0$,*[10]

$$\mathbb{P}_\mathcal{A}(s, \Box\Diamond(q, J) \xrightarrow{e} (q', I) \mid \Box\Diamond(q, J)) = 1\,.$$

*Proof.* We write $\mathbb{P}_\mathcal{A}(s \xrightarrow{e} (q', I))$ for the probability of the set of runs starting from $s$ with a move $s \xrightarrow{t,e} s'$ with $s' \in (q', I)$ and for some $t \in \mathbb{R}_+$.[11]

Let $\lambda \overset{\text{def}}{=} \inf_{s \in (q,J)} \mathbb{P}_\mathcal{A}(s \xrightarrow{e} (q', I))$. Since $\overline{J} \subseteq q$ is compact and $\forall s \in q$, $\mathbb{P}_\mathcal{A}(s \xrightarrow{e} (q', I)) > 0$ (because $\dim_\mathcal{A}(e) \neq \bot$ and $e(s) \cap I$ is non-empty and open), $\lambda > 0$. Indeed we have supposed that $\mu_s\big(\{d \mid s + d \in [a, b]\}\big)$ is continuous on $\{(s, a, b) \mid [a, b] \subseteq I(s)\}$, see the first hypothesis in (†), hence $s \mapsto \mathbb{P}_\mathcal{A}(s \xrightarrow{e} (q', I))$ is continuous.

Denote $E_k$ the set of paths in $\mathcal{A}$ that visit $(q, J)$ infinitely often, but from the $k$-th passage in $(q, J)$ on never fire $(q, J) \xrightarrow{e} (q', I)$ anymore. Note that the set $E_k$ is $\mathbb{P}_\mathcal{A}$-measurable, and that $\mathbb{P}_\mathcal{A}(E_k) \leq \prod_k^\infty (1 - \lambda) = 0$. Then note that the set $\bigcup_{k \geq 1} E_k$ can be equivalently defined by $B \wedge \neg A$ where $B$ is '$\Box\Diamond(q, J)$' and $A$ is '$\Box\Diamond(q, J) \xrightarrow{e} (q', I)$'. Hence, we get that $\mathbb{P}_\mathcal{A}(s, B \wedge \neg A) \leq \lim_{k \to +\infty} \mathbb{P}_\mathcal{A}(E_k) = 0$, and thus

$$\begin{aligned}
\mathbb{P}_\mathcal{A}(s, A \mid B) &= \tfrac{\mathbb{P}_\mathcal{A}(s, A \wedge B)}{\mathbb{P}_\mathcal{A}(s, B)} && \text{by definition}\\
&= \tfrac{\mathbb{P}_\mathcal{A}(s, A \wedge B)}{\mathbb{P}_\mathcal{A}(s, A \wedge B) + \mathbb{P}_\mathcal{A}(s, \neg A \wedge B)} && \text{by Bayes formulas}\\
&= 1 && \text{because } \mathbb{P}_\mathcal{A}(s, B \wedge \neg A) = 0
\end{aligned}$$

which is exactly $\mathbb{P}_\mathcal{A}(s, \Box\Diamond(q, J) \xrightarrow{e} (q', I) \mid \Box\Diamond(q, J)) = 1$. □

*Remark 22.* This lemma holds for all timed automata, not only one-clock timed automata.

We have done the proof for a single transition, but this lemma can be extended straightforwardly to finite sequences of edges as follows:

**Lemma 23.** *1. For all regions $(q_i)_{0 \leq i \leq p}$,*

2. *for all edges $(e_i)_{1 \leq i \leq p}$ such that $e_i$ is enabled in $q_{i-1}$ and $\dim_\mathcal{A}(e_i) \neq \bot$*
3. *for all subregions $((q_i, J_i))_{0 \leq i \leq p}$ such that (i) $J_i$ is non-empty and open in $q_i$ (for the induced topology), (ii) $\overline{J_i} \subseteq q_i$ is compact, and (iii) for every $s \in J_i$, $e_i(s) \cap J_{i+1}$ is non-empty and open, where $e_i(s) = \{s' \mid \exists t \in \mathbb{R}_+ \ s.t.\ s \xrightarrow{t,e_i} s'\}$,*
4. *for every edge $e$ enabled in $q$ such that $\dim_\mathcal{A}(e) \neq \bot$,*
5. *for every subregion $(q', I)$ of $q'$ such that for every $s \in (q_p, J_p)$, $e(s) \cap I$ is non-empty and open, where $e(s) = \{s' \mid \exists t \in \mathbb{R}_+ \ s.t.\ s \xrightarrow{t,e} s'\}$,*

[10] This is for the next conditional probability to be defined.

[11] Note that this set is $\mathbb{P}_\mathcal{A}$-measurable because it can be seen as $\mathsf{Cyl}(\pi_{\mathcal{C}_I}(s, e))$ for some constraint $\mathcal{C}_I$ enforcing the first move to lead to $I$.

*6. for every state $s$ of $\mathcal{A}$ such that $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond(q_0, J_0)) > 0$*

$$\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond\sigma \xrightarrow{e} (q', I) \mid \Box\Diamond\sigma) = 1$$

*where $\sigma = (q_0, J_0) \xrightarrow{e_1} (q_1, J_1) \ldots \xrightarrow{e_p} (q_p, J_p)$.*

Now, we can turn back to the proof of Lemma 20.

*Proof (of Lemma 20).* Let $s$ be a state. We decompose the set of infinite runs into:

($F_1$) the set of runs with infinitely many resets,

($F_2$) the set of runs with finitely many resets, and which are ultimately in the unbounded region $(M, +\infty)$,

($F_3$) the set of runs with finitely many resets, and which ultimately stay forever in a bounded region, either $\{c\}$ with $0 \leq c \leq M$, or $(c, c+1)$ with $0 \leq c < M$. We write $(F_3^{(c,c+1)})$ (resp. $(F_3^c)$) for condition $F_3$ restricted to $(c, c+1)$ (resp. $\{c\}$).

We write $\mathbb{P}_{\mathcal{A}}(s, F_i)$ for the probability of the runs starting in $s$ and satisfying condition $F_i$. The three sets of runs above are disjoint, cover the set of all runs, and are measurable. Hence $\sum_{i=1,2,3} \mathbb{P}_{\mathcal{A}}(s, F_i) = 1$, and $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair}) = \sum_{i=1,2,3} \mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_i) \cdot \mathbb{P}_{\mathcal{A}}(s, F_i)$ (application of the Bayes formula).[12] We now distinguish between the three cases.

**Case** $F_1$ We consider the set of runs with infinitely many resets. Let $\pi = s_0 \xrightarrow{e_1} s_1 \xrightarrow{e_2} \ldots$ be such a run. There exists $q$ such that for infinitely many $i$ with $i \in \mathbb{N}$, $s_i = (q, 0)$. Now, fix a state $(q, 0)$ and assume that $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond(q, 0)) > 0$ (otherwise the set of runs visiting infinitely often $(q, 0)$ will be negligible). For every sequence $\sigma$ of edges and compact sets (as in the statement of Lemma 23), we get that

$$\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond\sigma \mid \Box\Diamond(q, 0)) = 1\,.$$

Hence, for sequences of edges $(e_i)_{1 \leq i \leq p}$ such that such a $\sigma$ exists, we get that

$$\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond(q, 0) \xrightarrow{e_1} q_1 \ldots \xrightarrow{e_p} q_p \mid \Box\Diamond(q, 0)) = 1\,. \tag{1}$$

Now notice that such a $\sigma$ always exists whenever these edges have defined dimension.

Fix an edge $e$ with defined dimension, and assume that the set of paths passing through $(q, 0)$ infinitely often and enabling $e$ infinitely often, has a positive probability. We will then prove that

$$\mathbb{P}_{\mathcal{A}}(s, (\Box\Diamond e \text{ enabled}) \Rightarrow (\Box\Diamond \xrightarrow{e}) \mid \Box\Diamond(q, 0)) = 1\,,$$

which will imply that $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_1) = 1$.

[12] If $\mathbb{P}_{\mathcal{A}}(s, F_i) = 0$, we remove the $i$-th term from the sum, as the conditional probability $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_i)$ is then not defined, but the restricted sum is then still equal to $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair})$.

- Assume that $e$ is reachable from $(q,0)$ following edges of defined dimension, say $(e_i)_{1\leq i\leq p}$ with $e_p = e$. Then, applying $(\star)$, we get that $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond(q,0) \xrightarrow{e_1} q_1 \ldots \xrightarrow{e_p} q_p \mid \Box\Diamond(q,0)) = 1$, hence that $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond \xrightarrow{e} \mid \Box\Diamond(q,0)) = 1$.
- Assume on the contrary that $e$ is not reachable from $(q,0)$ following edges of defined dimension. If $e$ is not reachable from $(q,0)$, then $\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond e \text{ enabled} \mid \Box\Diamond(q,0)) = 0$. Let $W$ be the set of finite sequences of edges $(e_i)_{1\leq i\leq p}$ leading from $(q,0)$ to a state where $e$ is enabled. Then:

$$\begin{aligned}
\mathbb{P}_{\mathcal{A}}(\Box\Diamond e \text{ enabled} \mid \Box\Diamond(q,0)) &= \mathbb{P}_{\mathcal{A}}(\Box\Diamond \bigcup_{w\in W} w \mid \Box\Diamond(q,0)) \\
&\leq \mathbb{P}_{\mathcal{A}}(\Diamond \bigcup_{w\in W} w \mid \Box\Diamond(q,0)) \\
&= 0 \qquad \begin{array}{l}\text{because one of the edges in } w \\ \text{has undefined dimension}\end{array}
\end{aligned}$$

In both cases, we get the expected property.

**Case** $F_2$ We consider the set of runs with finitely many resets and which end up in the unbounded region $(M, +\infty)$. Let $\pi = s \xrightarrow{e_1} s_1 \xrightarrow{e_2} \ldots$ be such a run, and assume that from $s_n$ on, all states are in the unbounded region. From that state on, all edges which are enabled have defined dimension and have guard $x > M$ where $M$ is the maximal constant of the automaton. From region $q$, the probability of $\mathsf{Cyl}(\pi(s,e))$ for every $s \in q$ is independent of the choice of $s$. Hence, ultimately, after having reached the unbounded region (and never leave it anymore), it will behave like a finite Markov chain!

Assume now that a resetting edge $e$ is enabled infinitely often along $\pi$. Then, by a similar argument to the one in the proof of Lemma 21 with the $E_k$, as the probability distribution of taking an edge is lower-bounded (because we are now in a finite Markov chain), then any edge will be almost surely taken infinitely often. Hence,

$$\mathbb{P}_{\mathcal{A}}(s, \Box\Diamond \text{resetting edge enabled} \mid F_2) = 0\,,$$

and thus

$$\mathbb{P}_{\mathcal{A}}(s, \neg(\Box\Diamond \text{resetting edge enabled}) \mid F_2) = 1\,.$$

Once more, due to the distribution over edges (which is a finite Markov chain), when there is no more resetting edges, we get

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_2) = 1\,.$$

**Case** $F_3$ We consider the set of runs with finitely many resets and which end up in a bounded region (either $x = c$ with $c \leq M$ or $c < x < c+1$ with $c < M$). We assume the region $c < x < c+1$. Let $\pi = s \xrightarrow{e_1} s_1 \xrightarrow{e_2} \ldots$ be a witness run, and we assume that from $s_n$ on, we are in region $c < x < c+1$. If $s_i$ and $s_j$ with $n \leq i < j$ correspond to the same location, then the clock value of $s_i$ is less than (or equal to) that of $s_j$. Hence, if an edge $e$ with defined dimension and whose guard is included in $[c+1, +\infty)$ is enabled in $s_i$ (and thus also in $s_j$), the probability of taking $e$ from $s_j$ is greater than (or equal to) the probability of taking $e$ from $s_i$ (due to the second hypothesis in (†) on $\mu$'s and to the fact that the discrete probability over edges is constant by regions). Hence, there is a positive lower bound for the probability of taking $e$, and if $e$ is enabled infinitely often, it will be taken infinitely often. Such an enabled edge is thus only possible with probability 0 under the assumption made in this case. Hence, with probability 1, only edges with guard $c < x < c+1$ are enabled. For these edges, as previously, the system behaves like a finite Markov chain. We thus get that

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_3^{(c,c+1)}) = 1\,.$$

If we now assume the region $x = c$, the reasoning is very similar to the previous one. Given a location $\ell$ along the suffix of the path where $x = c$ always holds, the edges enabled in $(\ell, x = c)$ are equipped with a distribution defining a finite Markov chain. Hence any edge enabled infinitely often will be taken infinitely often almost surely, which implies that

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair} \mid F_3^{c}) = 1\,.$$

Gathering all cases, we get the desired property, *i.e.*, $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{fair}) = 1$. □

*Remark 24.* This proof really relies on the one-clock hypothesis (cases $F_1$ and $F_3$).

***A two-clock counter-example to Lemma 20.*** We shortly argue why Lemma 20 requires the restriction to single-clock timed automata. As pointed out in [CHR02], timed automata admit various *time converging* behaviours, and some of these behaviours, not occurring in one-clock timed automata, can lead to "big" sets of *unfair* executions. Inspired by an example of [CHR02], we design a two-clock timed automaton $\mathcal{A}$ (see Fig. 4) which does not satisfy Lemma 20. When $\mathcal{A}$ is equipped with uniform distributions, one can show that the probability to run forever through the cycle $\ell_0\,\ell_3\,\ell_4\,\ell_0$ is positive and therefore $\mathbb{P}_{\mathcal{A}}((\ell_0, 0, 0) \models \mathsf{fair}) < 1$.

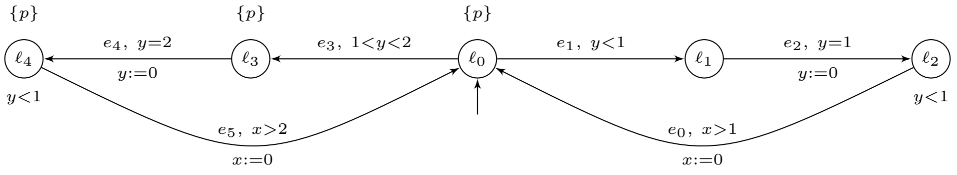

**Fig. 4.** A two-clock example with non negligible set of unfair runs

We prove that this automaton does not satisfy Lemma 20: Indeed, in $\ell_0$, both edges leading to the leftmost and the rightmost loops have defined dimension, but we show that with a positive probability, the rightmost loop will never be taken.

**Proposition 25.** *Let* $0 < t_0 < 1$*. We let* $S_{t_0}$ *be the set of runs starting in* $(\ell_0, (0, t_0))$*, which only take the leftmost loop of the automaton. Then,*

$$\mathbb{P}\big(S_{t_0}\big) > 0\,.$$

*Proof.* For every $N \geq 1$, we write $S_{t_0}^N$ for the set of runs starting in $s_0 = (\ell_0, (0, t_0))$, which only take the leftmost loop of the automaton for the $N$ first times. Then, obviously, $\mathbb{P}\big(S_{t_0}\big) = \lim_{N\to+\infty} \mathbb{P}\big(S_{t_0}^N\big)$.

We now would like to express $\mathbb{P}_{\mathcal{A}}(S_{t_0}^N)$ as a multiple integral. First notice that:

$$\mathbb{P}\big(S_{t_0}^N\big) = \mathbb{P}\Big(\pi(s_0, (e_3e_4e_5)^N)\Big)$$

In order to take the leftmost loop, we need to choose a first delay ensuring that the valuation of the clock $y$ satifies the guard $1 < y < 2$. The location $\ell_4$ is then reached with the clock valuation $(2 - t_0, 0)$ . From there a second positive time delay has to be chosen in order to reach location $\ell_0$. We thus have that:

$$\begin{aligned}\mathbb{P}\big(\pi(s_0, (e_3e_4e_5)^N)\big) &= \frac{1}{2-t_0}\int_{\tau=1}^{2} \frac{1}{1-t_0}\int_{t_1=t_0}^{1} \mathbb{P}\big(\pi(s_1, (e_3e_4e_5)^{N-1})\big)\mathrm{d}t_1\mathrm{d}\tau \\ &= \frac{1}{2-t_0}\cdot\frac{1}{1-t_0}\int_{t_1=t_0}^{1} \mathbb{P}\big(\pi(s_1, (e_3e_4e_5)^{N-1})\big)\mathrm{d}t_1\end{aligned}$$

where $s_1 = (\ell_0, (0, t_1))$. By iterating this process, we obtain that:

$$\mathbb{P}\big(S_{t_0}^N\big) = \frac{1}{2-t_0}.\frac{1}{1-t_0}\int_{t_1=t_0}^{1}\frac{1}{2-t_1}\cdot\frac{1}{1-t_1}\int_{t_2=t_1}^{1}\cdots\frac{1}{2-t_{N-1}}\cdot\frac{1}{1-t_{N-1}}\int_{t_N=t_{N-1}}^{1}\mathrm{d}t_N\ldots\mathrm{d}t_1$$

We write

$$\gamma_i^N = \frac{1}{1-t_{i-1}}\int_{t_i=t_{i-1}}^{1}\frac{1}{2-t_i}\cdot\frac{1}{1-t_i}\int_{t_{i+1}=t_i}^{1}\cdots\frac{1}{2-t_{N-1}}\cdot\frac{1}{1-t_{N-1}}\int_{t_N=t_{N-1}}^{1}\mathrm{d}t_N\ldots\mathrm{d}t_i$$

and we can prove by a descending induction on $i$ (see Lemma E in the appendix for a detailed proof) that

$$\gamma_i^N \;\geq\; \frac{2^{N+1-i}-1}{2^{N-i}} - \frac{2^{N-i}-1}{2^{N-i}}\cdot(2-t_{i-1}).$$

Thus, we deduce that

$$\begin{aligned}\mathbb{P}\big(S_{t_0}^N\big) &= \frac{1}{2-t_0}\cdot\gamma_1^N \\ &\geq \frac{1}{2-t_0}\cdot\Big(\frac{2^N-1}{2^{N-1}} - \frac{2^{N-1}-1}{2^{N-1}}\cdot(2-t_0)\Big)\end{aligned}$$

Hence, computing the limit, we get that

$$\mathbb{P}\big(S_{t_0}\big) \geq \frac{t_0}{2-t_0} > 0$$

This concludes the proof of the proposition. □

*Remark 26.* In the previous proof, we don't get that $\mathbb{P}(S_0)$ (*i.e.* $\mathbb{P}(S_{t_0})$ for $t_0 = 0$) is positive. However, roughly, after one loop, we will have $t_1 > 0$, hence we can apply the above result from the second loop on. Hence, we can write:

$$\begin{aligned}
\mathbb{P}(S_0) &= \frac{1}{2-0} \cdot \frac{1}{1-0} . \int_0^1 \mathbb{P}(S_{t_1})\, \mathrm{d}t_1 \\
&\geq \frac{1}{2} . \int_0^1 \frac{t_1}{2-t_1}\, \mathrm{d}t_1 \\
&\geq \frac{1}{2} . \int_0^1 \left(-1 + \frac{2}{2-t_1}\right) \mathrm{d}t_1 \\
&\geq \frac{1}{2} . \left[-t_1 - 2\log(2-t_1)\right]_{t_1=0}^1 \\
&\geq \log(2) - \frac{1}{2} \\
&> 0
\end{aligned}$$

This allows to extend Proposition 25 to the case $t_0 = 0$.

**Corollary 27.** *Let $0 \leq t_0 < 1$, and $s_0 = (\ell_0, (0, t_0))$. Then $\mathbb{P}(s_0 \models \mathsf{fair}) < 1$.*

*Proof.* Assume $0 \leq t_0 < 1$, and consider the set of runs in $\mathcal{A}$ starting in state $(\ell_0, (0, t_0))$. For all these runs, edge $e_1$ (from $\ell_0$ to $\ell_1$) is infinitely often enabled. However, the subset of runs that always ignore edge $e_1$ is not negligible. Recall that $e_1$ has a defined dimension. As a consequence, the set of fair runs has probability strictly less than 1. □

### 4.2 Relating probabilities and fair symbolic paths

We now come to one of the main results of this paper: relating fair infinite paths of defined dimension and sets of runs of probability 0.

**Theorem 28.** *Let $\mathcal{A}$ be a one-clock (non-blocking) timed automaton such that $\mathcal{A} = \mathsf{R}(\mathcal{A})$, and $\varphi$ be an $\mathsf{LTL}$ formula. If $s$ is a state of $\mathcal{A}$, then $\mathbb{P}_{\mathcal{A}}(s \models \varphi) > 0$ iff there exists a fair infinite symbolic path $\pi = \pi(s, e_1 e_2 \ldots)$ such that $\dim_{\mathcal{A}}(\pi) = \top$, and $\pi \models \varphi$.*

*Proof.* We prove the two implications separately.

($\Longrightarrow$) Let us assume that $\mathbb{P}_{\mathcal{A}}(s \models \varphi) > 0$. Thanks to Lemmas 13 and 20, $\mathbb{P}_{\mathcal{A}}(s \models \varphi) = \mathbb{P}_{\mathcal{A}}(s \models \varphi \wedge \mathsf{fair} \wedge \neg \mathsf{dim\_undef})$. Hence,

$$\mathbb{P}_{\mathcal{A}}(s \models \varphi \wedge \mathsf{fair} \wedge \neg \mathsf{dim\_undef}) > 0$$

In particular, there must exist an infinite path $\pi = \pi(s, e_1 e_2 \cdots)$ satisfying the three following conditions: $\pi \models \varphi$, $\pi$ is fair, and $\dim_{\mathcal{A}}(\pi) = \top$.

($\Longleftarrow$) Let $\pi = \pi(s, e_1 e_2 \cdots)$ be a symbolic path in $\mathcal{A}$ such that $\pi$ is fair, $\dim_{\mathcal{A}}(\pi) = \top$, and $\pi \models \varphi$.

We first assume $\varphi$ is a location-based $\omega$-regular accepting condition[13]. We color the graph $\mathcal{A}$ as it is done on page 16 (remind that $\mathcal{A} = R(\mathcal{A})$). Since $\dim_{\mathcal{A}}(\pi) = \top$, all edges in $\pi$ are "blue" edges. Hence $\pi$ is also a path in the graph $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, restriction of $\mathcal{A}$ to blue edges. Let us consider $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ in more details, and particularly its strongly connected components. As $\pi$ is a fair path, it eventually reaches a BSCC in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ and from then on takes each edge of the BSCC infinitely often. Otherwise, this would mean that $\pi$ ignores a blue edge (hence an edge with dimension) forever, and contradict the fairness assumption. Let $B_\pi$ be the BSCC that $\pi$ eventually reaches and $\pi_{\mathsf{pref}}$ the shortest prefix of $\pi$ leading from $s$ to $B_\pi$ (note that it has a defined dimension). Consider the following set of paths in $\mathcal{A}$:

$$E \overset{\text{def}}{=} \{\pi' \in \mathsf{Cyl}(\pi_{\mathsf{pref}}) \mid \dim_{\mathcal{A}}(\pi') = \top \text{ and } \pi' \models \mathsf{fair}\}\,.$$

Thanks to Lemmas 13 and 20, $\mathbb{P}_{\mathcal{A}}(E) = \mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi_{\mathsf{pref}})) > 0$. It now suffices to show that all paths in $E$ satisfy $\varphi$. This is rather clear if we assume that $\varphi$ is a location-based regular property. Indeed, in such cases, the satisfiability of $\varphi$ only depends on the set of states that are visited infinitely often, and it is the case that such states for paths in $E$ are exactly the states in $B_\pi$.

We now assume that $\varphi$ is an $\mathsf{LTL}$ formula. We first need to build the product of a deterministic Muller automaton[14] for $\varphi$ and our timed automaton $\mathcal{A}$. We detail below this product construction and show that the probability distribution over paths in $\mathcal{A}$ is closely related to the one in the product. Given $\mathcal{A} = (L, X, E, \mathcal{I}, \mathcal{L})$ a timed automaton, and $\mathcal{B} = (\mathsf{S}, \mathsf{s}_0, 2^{\mathsf{AP}}, \rightarrow_{\mathcal{B}}, \mathcal{F})$ a complete and deterministic Muller automaton with alphabet $2^{\mathsf{AP}}$, initial state $\mathsf{s}_0$, and accepting condition $\mathcal{F} \subseteq 2^{\mathsf{S}}$, the product $\mathcal{A} \times \mathcal{B}$ is the timed automaton $\widetilde{\mathcal{A}} = (\widetilde{L}, X, \widetilde{E}, \widetilde{\mathcal{I}}, \widetilde{\mathcal{L}})$ where:

- $\widetilde{L} = L \times \mathsf{S}$,
- if $e = (\ell, g, Y, \ell') \in E$ then for all $\mathsf{s} \in \mathsf{S}$ s.t. $\mathsf{s} \xrightarrow{\mathcal{L}(\ell)}_{\mathcal{B}} \mathsf{s}'$ (this edge is unique by determinism), there is an edge $\widetilde{e}_{\mathsf{s}} = ((\ell, \mathsf{s}), g, Y, (\ell', \mathsf{s}'))$ in $\widetilde{E}$,
- for all $\mathsf{s} \in \mathsf{S}$, $\widetilde{\mathcal{I}}(\ell, \mathsf{s}) = \mathcal{I}(\ell)$, and $\widetilde{\mathcal{L}}(\ell, \mathsf{s}) = \emptyset$.

*Remark 29.* It should be clear enough that $\widetilde{\mathcal{A}} = \mathcal{A} \times \mathcal{B}$ is non-blocking as soon as $\mathcal{A}$ is. Moreover, for all states $s$ of $\mathcal{A}$, for all states $\mathsf{s}$ of $\mathcal{B}$, and for all edges $e \in E$, $I(s, e) = I((s, \mathsf{s}), \widetilde{e}_{\mathsf{s}})$.

Since the intervals $I(s, e)$ and $I((s, \mathsf{s}), \widetilde{e}_{\mathsf{s}})$ coincide (for all $\mathsf{s} \in \mathsf{S}$), it is legitimate to assign the same measure over delays in $(s, \mathsf{s})$ and in $s$.

[13] *I.e.*, a Büchi, Muller, parity, Rabin or Streett condition.

[14] Or Streett, Rabin, parity etc...

**Lemma 30.** *Let $\mathcal{A}$ be a timed automaton (such that $\mathcal{A} = R(\mathcal{A})$) and $\mathcal{B}_\varphi$ a complete deterministic Muller automaton for formula $\varphi$ with accepting condition $\mathcal{F}$. Assume furthermore that $\mu_{(s,\mathsf{s})}$ is set to $\mu_s$ for every state $\mathsf{s}$ of $\mathcal{B}_\varphi$ and every state $s$ of $\mathcal{A}$. Then:*

$$\mathbb{P}_{\mathcal{A}}(s \models \varphi) = \mathbb{P}_{\mathcal{A}\times\mathcal{B}_\varphi}((s,\mathsf{s}_0) \models \mathcal{F})\,.$$

*Proof.* To each run $\varrho = s_0 \xrightarrow{\tau_1,e_1} s_1 \xrightarrow{\tau_2,e_2} s_2 \cdots$ in $\mathcal{A}$ corresponds a unique run of the form $(s_0,\mathsf{s}_0) \xrightarrow{\tau_1,e_1} (s_1,\mathsf{s}_1) \xrightarrow{\tau_2,e_2} (s_2,\mathsf{s}_2) \cdots$ in $\mathcal{A}\times\mathcal{B}_\varphi$. This run is denoted $\varrho^{\mathcal{B}_\varphi}$ and its existence and unicity are consequences of $\mathcal{B}_\varphi$ being complete and deterministic. Moreover, $I_{\mathcal{A}}(s_i) = I_{\mathcal{A}\times\mathcal{B}_\varphi}(s_i,\mathsf{s}_i)$. Conversely, each run in $\mathcal{A}\times\mathcal{B}_\varphi$ has a unique preimage in $\mathcal{A}$ (obtained by removing the $\mathcal{B}_\varphi$ component in each state). Together with the assumption $\mu_{(s,\mathsf{s})} = \mu_s$, this yields that the measure of a set of runs in $\mathcal{A}$ is the same as the measure of the set of their images in $\mathcal{A}\times\mathcal{B}_\varphi$: $\mathbb{P}_{\mathcal{A}}(E) = \mathbb{P}_{\mathcal{A}\times\mathcal{B}_\varphi}(\{\varrho^{\mathcal{B}_\varphi} \mid \varrho \in E\})$. The lemma is then a consequence of the following observation: $\pi$ in $\mathcal{A}$ is accepted by $\mathcal{B}_\varphi$ if and only if $\pi^{\mathcal{B}_\varphi} \models \mathcal{F}$. □

This completes the proof of Theorem 28. □

*Remark 31.* Note that the latter theorem would not hold for general timed automata: For the two-clock example of Fig. 4, $\mathbb{P}((\ell_0,0,0) \models \mathbf{G}\,p) > 0$ but there is no fair path satisfying $\mathbf{G}\,p$.

Indeed, from Proposition 25 (and Remark 26), we know that the set of runs starting in $s_0$ and which only take the left circuit has positive probability. Hence $\mathbb{P}(s_0 \models \mathbf{G}\,p) > 0$. However, any fair run issued from $s_0$ cannot always avoid taking edge $e_1$ since this edge has defined dimension. Hence, there is no fair run satisfying $\mathbf{G}\,p$.

### 4.3 Equivalence of the almost-sure and large semantics

We can now state the second main result of this paper, relating the almost-sure and the large semantics for $\mathsf{LTL}$. In particular, this result shows that the almost-sure semantics does not depend on the concrete choice of the weights $p_s(e)$ and the measures $\mu_s$.

**Theorem 32.** *Let $\mathcal{A}$ be a one-clock (non-blocking) timed automaton, and $\varphi$ an $\mathsf{LTL}$ formula. Let $s$ be a state of $\mathcal{A}$. Then,*

$$\mathcal{A}, s \approx\!\!\!\models_{\mathbb{P}} \varphi \quad\Leftrightarrow\quad \mathcal{A}, s \approx\!\!\!\models_{\mathcal{T}} \varphi\,.$$

*Proof.* We prove the theorem for $\omega$-regular location-based acceptance conditions. The extension to $\mathsf{LTL}$ formulae can be done as in the proof of Theorem 28.

Applying Lemma 8, we get that

$$\mathbb{P}_{\mathcal{A}}(s \models \varphi) = 1 \quad\text{iff}\quad \mathbb{P}_{\mathsf{R}(\mathcal{A})}(\iota(s) \models \varphi) = 1\,.$$

Similarly, applying Corollary 17, we get that

$$\{\varrho \in \mathsf{Runs}(\mathcal{A}, s) \mid \varrho \text{ fair}\} \text{ is large} \quad \text{iff} \quad \{\varrho \in \mathsf{Runs}(\mathsf{R}(\mathcal{A}), \iota(s)) \mid \varrho \text{ fair}\} \text{ is large.}$$

Thus, the theorem is equivalent to the same result for $\mathsf{R}(\mathcal{A})$ instead of $\mathcal{A}$. We thus now assume w.l.o.g. that $\mathcal{A} = \mathsf{R}(\mathcal{A})$.

The proof then partly relies on Theorem 28. We indeed prove that these two properties are equivalent:

(1) the set of fair infinite runs satisfying $\varphi$ is large,
(2) for every fair infinite symbolic path $\pi$ such that $\dim_{\mathcal{A}}(\pi) = \top$, $\pi \models \varphi$.[15]

First note that the set of fair infinite runs satisfying $\varphi$ is the union of the fair infinite symbolic paths satisfying $\varphi$. We write $[\![\varphi]\!]_{\mathit{fair}}$ for the set of runs that are read over fair infinite symbolic paths starting from $s$, with defined dimension, and satisfying $\varphi$.

We first prove that (2) implies (1). We prove that $[\![\varphi]\!]^c_{\mathit{fair}}$ (the complement of $[\![\varphi]\!]_{\mathit{fair}}$) is meager using Banach-Mazur games (see definition on page 6, or [Oxt57]), which will imply that $[\![\varphi]\!]_{\mathit{fair}}$ is large, hence point (1) holds. We define as earlier $\mathcal{B}$, the family we play with, as the set of all basic open sets. We will prove that Player 2 has a winning strategy to avoid $[\![\varphi]\!]^c_{\mathit{fair}}$. The strategy of Player 2 is as follows:

- Player 1 has chosen a cylinder $\mathsf{Cyl}(\alpha_0)$ which ends up in state $q_0$ (with $\dim(\alpha_0) = \top$),
- Player 2 then chooses $\mathsf{Cyl}(\alpha_1)$ such that $\alpha_0$ is a strict prefix of $\alpha_1$ with defined dimension, and $\alpha_1$ ends up in a BSCC $B$ of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$.
- Then, whatever Player 1 chooses, Player 2 can ensure that all possible edges with defined dimension of the BSCC $B$ are visited infinitely often.

Under that strategy, the outcome is an infinite symbolic path,[16] which is fair, has defined dimension (because all chosen cylinders have defined dimension), and hence satisfies $\varphi$ by hypothesis. Hence, its intersection with $[\![\varphi]\!]^c_{\mathit{fair}}$ is empty, which yields the expected result.

We now prove that (1) implies (2) (or more precisely its contrapositive). We assume that there exists a fair infinite path $\pi$ such that $\dim(\pi) = \top$ and $\pi \not\models \varphi$, and show that the set $[\![\varphi]\!]_{\mathit{fair}}$ is not large. This fair infinite path $\pi$ ends up in a BSCC of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. Let $\pi'$ be the shortest prefix of $\pi$ which ends in this BSCC. Then, as $\varphi$ is location-based, every fair infinite path with prefix $\pi'$ (*i.e.*, in $\mathsf{Cyl}(\pi')$) will not satisfy $\varphi$. Hence $[\![\varphi]\!]^c_{\mathit{fair}}$ is non-meager (because $(\mathsf{Runs}(\mathcal{A}, s), \mathcal{T}^s_{\mathcal{A}})$ is a Baire space) and $[\![\varphi]\!]_{\mathit{fair}}$ is not large. □

[15] Thanks to Theorem 28, this is equivalent to $\mathbb{P}_{\mathcal{A}}(s \models \varphi) = 1$.

[16] Formally, it would be included in such an infinite symbolic path, as all its prefixes are supposed to be constrained. Note that if the obtained constrained symbolic path is empty, Player 2 automatically wins the game, as desired.

*Remark 33.* Theorem 32 does not hold for general timed automata. Indeed for the two-clock example $\mathcal{A}$ of Fig. 4, with $s_0 = (\ell_0, 0, 0)$, $\mathcal{A}, s_0 \mathrel{|\!\!\approx_{\mathcal{T}}} \mathbf{F}\,\neg p$ but $\mathcal{A}, s_0 \mathrel{\not|\!\!\approx_{\mathbb{P}}} \mathbf{F}\,\neg p$.

Indeed, the set of runs starting in $s_0$ and satisfying $\mathbf{F}\,\neg p$ is large. Indeed we show that $[\![\mathbf{G}\,p]\!]$ is meager, using (once more!) Banach-Mazur games. Player 2 clearly has a strategy to ensure $\cap_i B_i \cap [\![\mathbf{G}\,p]\!] = \emptyset$, by visiting locations from the right-hand side loop, whatever the first move of Player 1 is. Thus, in a single round of the game, Player 2 wins, and $[\![\mathbf{G}\,p]\!]$ is meager. However, $\mathbb{P}(s_0 \models \mathbf{F}\,\neg p) < 1$ as a consequence of Proposition 25 (and Remark 26).

### 4.4 Decidability of the model-checking problems

Gathering the results of this section and using an “optimized version” of one-dimensional regions [LMS04], we get the following results for the two model-checking problems:

**Corollary 34.** *The almost-sure and large model-checking problems for one-clock timed automata are*

*(i)* NLOGSPACE-*Complete for location-based $\omega$-regular properties, and*
*(ii)* PSPACE-*Complete for* LTL *properties.*

*Proof.* In everything that precedes, we have assumed that regions of the timed automaton are intervals of length one, by splitting constraints with respect to all constants $c$ below $M$, the maximal constant appearing in the timed automaton. This can be improved [LMS04] by considering regions of the form $(c, d)$ or $\{c\}$ where $c$ and $d$ are two consecutive constants appearing in the automaton. This leads to a region automaton whose size is no more exponential in the size of the original automaton, but whose size becomes only polynomial in the size of the automaton. All results we have obtained in the refined framework can be obtained as well in this optimized construction. It is hence sufficient to look for a reachable BSCC in the blue optimized graph, also denoted $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, which satisfies the $\omega$-regular property we want to verify. We focus here on a Streett accepting condition of the form $\bigwedge_{1 \le i \le p} (\Box\Diamond\alpha_i \to \Box\Diamond\beta_i)$ where $\alpha_i$ and $\beta_i$ are sets of states of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$.[17] We assume that given a finite graph $\mathcal{G}$, $q$ a state of $\mathcal{G}$, and $S$ a set of states of $\mathcal{G}$, $\mathsf{Reachability}_{\mathcal{G}}(q, S)$ decides whether a state of $S$ is reachable from $q$ in $\mathcal{G}$. Such a procedure runs in NLOGSPACE (in the size of $\mathcal{G}$). We propose the following non-deterministic procedure to decide the negation of the above Streett accepting condition on $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ (interpreted in a probabilistic manner):

```
# Guess an index $1 \le i \le p$
# Guess a state $q$ (of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$) in $\alpha_i$ such that $\mathsf{Reachability}_{\mathcal{G}_{\mathsf{b}}(\mathcal{A})}(q_0, \{q\})$
    (where $q_0$ is the initial state of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$)
# If not $\mathsf{Reachability}_{\mathcal{G}_{\mathsf{b}}(\mathcal{A})}(q, \beta_i)$,
```

[17] This means that for every $1 \le i \le p$, if $\alpha_i$ is visited infinitely often, then so is $\beta_i$.

```
#    then if not $\mathsf{notBSCC}_{\mathcal{G}_{\mathsf{b}}(\mathcal{A})}(q)$
#        then return false
```

where $\mathsf{notBSCC}_{\mathcal{G}}(q)$ decides whether $q$ is not in a BSCC of a finite graph $\mathcal{G}$ as follows:

```
# Guess a state $q'$ of $\mathcal{G}$
# If $\mathsf{Reachability}_{\mathcal{G}}(q, \{q'\})$
#    then if not $\mathsf{Reachability}_{\mathcal{G}}(q', \{q\})$
#        then return false
```

Globally, this algorithm can be turned (applying Immerman-Szelepcsényi Theorem, stating that co-NLOGSPACE coincides with NLOGSPACE) into an NLOGSPACE algorithm. Note that there was no need to first construct $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, as edges with defined dimension can be guessed locally. The NLOGSPACE lower bound is trivial from that of the reachability problem in a finite graph.

The complexity result for LTL properties comes from the exponential blowup due to the product of a deterministic automaton for the formula with the timed automaton. □

# 5 A Note on Zeno Behaviours

In timed automata, and more generally in continuous-time models, some runs are *Zeno*.Recall that a run $\varrho = s_0 \xrightarrow{\tau_1 \cdot e_1} s_1 \xrightarrow{\tau_2 \cdot e_2} \cdots$ of a timed automaton is *Zeno* if $\sum_{i=1}^{\infty} \tau_i < \infty$ (*i.e.*, infinitely many actions happen in a finite amount of time). Zeno behaviours are problematic since they most of the time have no physical interpretation. As argued in [DP03], some fairness constraints are often put on executions, enforcing non-Zeno behaviours, but in probabilistic systems, probabilities are supposed to replace fairness assumptions, and it is actually the case in continuous-time Markov chains in which Zeno runs have probability 0 [BHHK03]. In our framework, it is hopeless to get a similar result because some timed automata are *inherently* Zeno. For instance, all runs are Zeno in the automaton consisting of a single location with a non-resetting loop guarded by $x \leq 1$. However, we show that we can decide whether the probability of the set of Zeno runs in a (one-clock) timed automaton is 0. We also give a nice characterization of the one-clock timed automata for which Zeno behaviours are negligible. This class is natural, since it corresponds to those automata which have no *'inherently Zeno components'* (reachable with a positive probability). Finally, we will see that the so-defined class encompasses classical definitions of *non-Zeno* timed automata.

## 5.1 Checking probabilistic non-Zenoness

We write $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno})$ for the probability of the set of Zeno runs in $\mathcal{A}$ from $s$. This set is measurable (in $\Omega_{\mathcal{A}}^{s}$), as it can be written as $\bigcup_{M \in \mathbb{N}} \bigcap_{n \in \mathbb{N}} \mathsf{Cyl}(\pi_{\mathcal{C}_{n,M}}(s, e_1, \ldots, e_n))$ where $\mathcal{C}_{n,M}$ is the constraint $\sum_{1 \leq i \leq n} \tau_i \leq\ M$.

**Theorem 35.** *Given a single-clock (non-blocking) timed automaton $\mathcal{A}$ and a state $s$ of $\mathcal{A}$, one can decide in* $\mathsf{NLOGSPACE}$ *whether* $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = 0$.

*Proof.* Thanks to Lemma 8, w.l.o.g. we can assume that $\mathcal{A} = R(\mathcal{A})$. Note that we can first remove syntactically all resets from edges labelled by $x = 0$. Fix a state $s$ in $\mathcal{A}$. As in the previous section we decompose the set of infinite runs into:

($F_1$) the set of runs with infinitely many resets,

($F_2$) the set of runs with finitely many resets, and which are ultimately in the unbounded region $(M, +\infty)$,

($F_3$) the set of runs with finitely many resets, and which ultimately stay forever in a bounded region, either $\{c\}$ with $0 \leq c \leq M$, or $(c, c+1)$ with $0 \leq c < M$.

We borrow notations from the previous section, and in that case, we also have that

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = \sum_{i=1,2,3} \mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_i) \cdot \mathbb{P}_{\mathcal{A}}(s, F_i) \qquad (2)$$

when these conditional probabilities are well-defined (otherwise it is correct to remove the term from the sum).

The proof of the theorem is then decomposed into two parts, first we prove that the two first terms of the above sum always equal to 0, and then that we can decide whether the last term is equal to 0.

**Lemma 36.** $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_1) = 0$ *and* $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_2) = 0$.

*Proof.* We distinguish between the two cases.

**Case** $F_1$ We consider the set of runs with infinitely many resets. This set can be decomposed according to the states $(q, 0)$ (where $q \in Q$ is a region) that are visited infinitely often. We show that $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid \Box\Diamond(q, 0)) = 0$. In order to prove this, we distinguish the four following subcases depending on the set $I((q, 0))$: either (*i*) $I((q, 0)) \cap [0, 1) = \emptyset$, or (*ii*) $(0, 1) \subseteq I((q, 0))$, or (*iii*) $\{0\} \subsetneq I((q, 0))$, or (*iv*) $\{0\} = I((q, 0))$.

Let us first treat the easy case (*i*). If $I((q, 0)) \cap [0, 1) = \emptyset$, since the timed automaton is non-blocking, this means that each time the automaton arrives in state $(q, 0)$ at least 1 time unit elapses before the next transition. Hence a run visiting infinitely often such state $(q, 0)$ is necessarily non-Zeno.

Let us now consider case (*ii*), *i.e.*, we assume that $(0, 1) \subseteq I((q, 0))$. Since the probability distribution over the delays is then equivalent to the Lebesgue measure (see hypothesis ($\star$)), the probability of waiting a time delay $\tau \leq \frac{1}{2}$ in $(q, 0)$ is positive and strictly smaller than 1 (we write $\lambda_{(q,0)}$ for this value: $0 < \lambda_{(q,0)} < 1$). Let $E_k$ be the set of runs starting from $s$, visiting $(q, 0)$ infinitely often, and such that from the $k$-th passage on, the time elapsed from state $(q, 0)$ (before taking an action) is less than $\frac{1}{2}$. We have $\mathbb{P}_{\mathcal{A}}(E_k) \leq \prod_k^\infty \lambda_{(q,0)} = 0$, and as a consequence

$$\mathbb{P}_{\mathcal{A}}\big(s \models \mathsf{Zeno} \mid \Box\Diamond(q, 0) \wedge (ii)\big) \leq \sum_{k=0}^{\infty} \mathbb{P}(E_k) = 0\,.$$

In case $(iii)$, we assume that $\{0\} \subsetneq I((q,0))$. If $(0,1) \subseteq I((q,0))$, we are done by case $(ii)$. We can thus suppose that if $0 \neq \tau \in I((q,0))$, we have that $\tau \geq 1$. If $I((q,0))$ reduces to a finite union of points, the probability $\lambda_0$ of waiting a delay greater than or equal to 1 is positive and strictly smaller than 1 (because the measure is then equivalent to the uniform measure over those points, see hypothesis $(\star)$). When going infinitely often through $(q,0)$, we will thus wait infinitely often a time greater than or equal to 1. If $I((q,0))$ contains an open interval, the probability of waiting a delay greater or equal than 1 from $(q,0)$ is 1 (by hypothesis $(\star)$). From this we can easily derive that:

$$\mathbb{P}\big(s \models \mathsf{Zeno} \mid \Box\Diamond(q,0) \wedge (iii)\big) = 0\,.$$

Let us conclude with case $(iv)$ where $I((q,0)) = \{0\}$. Since no positive delay can elapse from $(q,0)$, the probability of taking any edge enabled in $(q,0)$ is positive (the distribution over edges indeed becomes uniform). Hence, any state $(q_e,0)$ reachable from $(q,0)$ taking edge $e$, is almost surely infinitely often visited (as soon as $(q,0)$ is). From $(q_e,0)$, again two situations are possible: either $I((q_e,0)) = \{0\}$ or not. In the first case, note that it is necessarily the case that such a chain $(q,0) \xrightarrow{0,e_1} (q_1,0) \xrightarrow{0,e_2} (q_2,0)\cdots$ is finite, otherwise the run would contain only finitely many resets[18]. Thus we surely reach infinitely often a state $(q',0)$ such that $I((q',0)) \neq \{0\}$ allowing us to rely on the previous cases to obtain the desired results.

Gathering the four cases, we conclude that $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid \Box\Diamond(q,0)) = 0$. Hence

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_1) = 0\,.$$

**Case** $F_2$ We consider the set of runs with finitely many resets and which end up in the unbounded region. From any state $s$ in the unbounded region, the set of potential delays is necessarily of the form $[0,+\infty)$[19]. From hypothesis $(\dagger)$ on the distributions over delays, the probability of waiting a time delay $\tau \leq \frac{1}{2}$ from $s$, denoted $\lambda_s$, can be bounded by a constant: $0 < \lambda_s \leq \lambda_0 < 1$. Let $E_k$ denote the set of executions which, at the $k$-th step, are in the unbounded resgion without leaving it afterwards, and such that all delays afterwards are less than $\frac{1}{2}$. The probability of being Zeno when in $E_k$ satisfies: $\mathbb{P}(E_k) \leq \prod_{i>k} \lambda_0 = 0$, from which we derive:

$$\mathbb{P}(s \models \mathsf{Zeno} \mid F_2) \leq \sum_{k=0}^{\infty} \mathbb{P}(E_k) = 0\,.$$

This conclude the proof of the first lemma. □

The case of condition $F_3$ is not similar to the two previous cases. Indeed, it is worth noticing that every execution satisfying the condition $F_3$ is Zeno. Hence, if

[18] Recall that edges labelled with $x = 0$ are not labelled with a reset.

[19] Otherwise the clock would be compared to a constant greater than the maximal one

$\mathbb{P}_{\mathcal{A}}(s \models F_3) \neq 0$ (otherwise the term $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_3) \cdot \mathbb{P}_{\mathcal{A}}(s \models F_3)$ does not appear in the sum 2), then $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno} \mid F_3) = 1$. It remains to compute or characterize the value $\mathbb{P}_{\mathcal{A}}(s \models F_3)$.

A BSCC $B$ in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ is called a *Zeno BSCC* if it is bounded and contains no resetting edges. Note that in a Zeno BSCC the value of the clock lies in a unique interval $(c, c+1)$ (with $0 \leq c < M$) or $\{c\}$ (with $0 \leq c \leq M$).

**Lemma 37.** $\mathbb{P}_{\mathcal{A}}(s \models F_3) = \displaystyle\sum_{B \text{ Zeno BSCC}} \mathbb{P}_{\mathcal{G}_b(\mathcal{A})}(s \models \Diamond B)$.

*Proof.* Runs in $\mathcal{A}$ are almost surely fair (thanks to Lemma 20), hence $\mathbb{P}_{\mathcal{A}}(s \models F_3) = \mathbb{P}_{\mathcal{A}}(s \models F_3 \wedge \mathsf{fair})$. By definition of a fair run, if a blue edge is enabled infinitely often, then this edge appears infinitely often along that run. Now any fair path in $\mathcal{A}$ which takes a red edge has probability 0, hence it is sufficient to consider fair paths in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. In that case, fair runs correspond to paths in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ which end up in a BSCC. It is now sufficient to remark that for fair runs in $F_3$, the BSCC should be bounded and without resetting edges. Indeed, if one of these condition does not hold, the run would not be in $F_3$ (either it would end up in an unbounded region, or have infinitely many resets). Conversely, any run ending up in a Zeno BSCC satisfy $F_3$. Hence, the mentioned equality holds. □

From all these results, we get that

$$\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = 0 \quad \text{iff} \quad \forall B \text{ Zeno BSCC, } \mathbb{P}_{\mathcal{G}_{\mathsf{b}}(\mathcal{A})}(s \models \Diamond B) = 0$$

Applying results from [BBB$^+$07], it is easy to decide the right-hand side of the equivalence. It reduces to checking whether there exists a Zeno BSCC in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, reachable in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, *i.e.*, reachable in $\mathcal{A}$ *via* blue edges. This can be done in $\mathsf{NLOGSPACE}$, if we take the optimized one-dimensional region automaton already mentioned in the proof of Theorem 28. □

### 5.2 Topological characterization of probabilistic non-Zenoness

In Section 4, we gave a topological characterization of the probability of sets of runs defined by an $\mathsf{LTL}$ formula. Although Zeno runs cannot be defined in $\mathsf{LTL}$, we obtain a similar result.

**Theorem 38.** *Let $\mathcal{A}$ be a one-clock (non-blocking) timed automaton and $s$ a state of $\mathcal{A}$. Then, $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = 0$ iff the set of Zeno runs starting in $s$ is meager.*

The proof of this theorem once more relies on an application of Banach-Mazur games.

*Proof.* Assume first that $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = 0$. Then no BSCC of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$ is Zeno. We once more play a Banach-Mazur game using the basic open sets. Player 1 plays some move $\alpha_0$ (possibly with some constraint), and player 2 then plays a move $\alpha_1$ leading to a BSCC $B$ of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. By hypothesis, $B$ is not a Zeno BSCC, hence either it is not bounded, or it contains resetting edges.

- We first consider the case where $B$ contains no resetting edges. In that case, it means that the clock value when in $B$ is always above the maximal constant. Hence, the game can keep going on, and each time Player 2 chooses a move, he first chooses a move which constrains the cylinder saying that the delay has to be larger than 1. This is always possible, due to the form of the constraints, which all include $(M, +\infty)$. In that case, it is not difficult to check that the resulting runs are all non-Zeno.
- We now consider the case where $B$ has resetting edges. Note that the clock can become larger than 0. In that case, Player 2 can always choose a move so that it terminates with a resetting edge, but has visited a positive region, and has enforced that the value of the clock in that precise region was larger than 1/2. In that case also, all runs resulting from that play are non-Zeno.

Hence, we get that Player 2 has a strategy to avoid the set of Zeno runs, hence this set is meager.

Conversely assume that the set of Zeno runs is meager, but assume also that $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) > 0$. Once more, let's play the Banach-Mazur game. Player 2 has a strategy to avoid Zeno behaviours. However, as $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) > 0$, Player 1 can play a first move leading to a Zeno BSCC $B$ of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. Then $B$ has no resetting edges and lies within an interval $(c; c+1)$ or $\{c\}$. Then whatever move chooses Player 2, the resulting runs will all be Zeno, hence contradicting the assumption that the set of Zeno runs is meager. The claim follows. □

### 5.3 Relation with classical non-Zenoness assumptions

The proof of Theorem 35 gives a characterization of automata for which the probability of Zeno runs is 0: they are those timed automata $\mathcal{A}$ in which there are no Zeno BSCCs in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. In the literature, several assumptions can be found, to handle Zeno runs. We pick two such assumptions, and show that our framework gives probability zero to Zeno runs under those restrictions.
In their seminal paper about timed automata, Alur and Dill [AD94] want to decide the existence of non-Zeno accepted behaviours. They prove it is equivalent to having, in the region automaton, a reachable SCC (strongly connected component) satisfying the following *progress condition*: the SCC is either not bounded or with a reset of a clock. This condition is looser than the strong non-Zenoness of [AMPS98] (a witness is the simple automaton $\mathcal{A}_0$ depicted on Fig. 5) but is stronger than our condition on Zeno BSCC in $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$. Indeed our condition only constrains bottom SCCs of $\mathcal{G}_{\mathsf{b}}(\mathcal{A})$, and only those reachable by finite paths of defined dimension. Hence, any automaton $\mathcal{A}$ such that the region automaton contains no bounded SCC without resetting edges satisfies $\mathbb{P}_{\mathcal{A}}(s \models \mathsf{Zeno}) = 0$. The automaton $\mathcal{A}_1$ (resp. $\mathcal{A}_2$) of Fig. 6 (resp. Fig. 7) is not *strongly non-Zeno* and does not satisfy the *progress condition*, however it satisfies $\mathbb{P}_{\mathcal{A}_1}(s \models \mathsf{Zeno}) = 0$ (resp. $\mathbb{P}_{\mathcal{A}_2}(s \models \mathsf{Zeno}) = 0$). Let us notice

that $\mathcal{A}_1$ does not satisfy the *progress condition* since its region automaton contains a bounded SCC without resetting edges (which is not a bottom SCC). The $\mathcal{A}_2$ does not satisfy the *progress condition* for the same reason, however in this case, the bounded SCC without resetting edges is a bottom SCC, but it is only reachable by finite paths of undefined dimension.

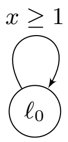

**Fig. 5.** $\mathcal{A}_0$

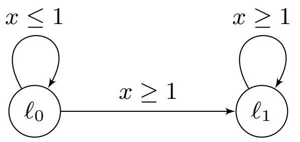

**Fig. 6.** $\mathcal{A}_1$

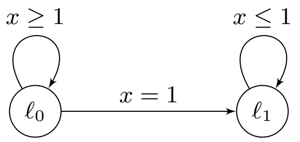

**Fig. 7.** $\mathcal{A}_2$

Let us give a final example of timed automata which does not satisfy the *progress condition* although the probability of its set of Zeno runs is zero. As for $\mathcal{A}_1$ the reason $\mathcal{A}_3$ (see Fig. 8) does not satisfy the *progress condition* is that its region automaton (Fig. 9) contains a bounded SCC without resetting edges (which is not a bottom SCC).

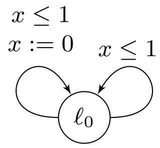

**Fig. 8.** $\mathcal{A}_3$

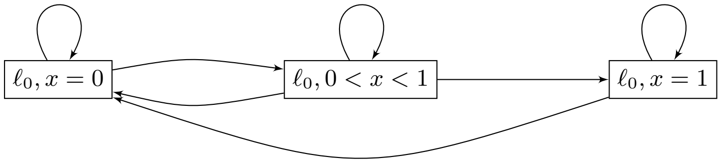

**Fig. 9.** The region automaton of $\mathcal{A}_3$

## 6 Conclusion

The goal of this paper was to present non-standard semantics for LTL interpreted over timed automaton that rule out "unlikely" events, but do not affect the decidability and complexity of the model checking problem. For this purpose, we introduced a probabilistic almost-sure semantics that relies on some mild stochastic assumptions about the delays and the resolution of the nondeterministic choices, and a topological semantics based on the notion of largeness. For one-clock timed automata we proved the equivalence of the two semantics. The topological characterization of the almost-sure semantics has several important consequences: first, it shows that the precise choice of the measures used in the definition of the almost-sure semantics are irrelevant and second, as the topology is defined by the local conditions (using the notion of dimension), it yields a graph-based model-checking algorithm.

Although the formal definitions of the probabilistic and topological semantics reuse concepts of [BBB$^+$07], where similar questions have been studied when interpreting LTL over finite words, the results for LTL over infinite words presented in this paper cannot be viewed as consequences of [BBB$^+$07]. This becomes clear from the observation that the almost-sure and topological semantics for LTL over infinite words do not agree for timed automata with two or more clocks, while the approach of [BBB$^+$07] does not impose any restrictions on the number of clocks. In fact, our proof for the topological dimension-based characterization of the almost-sure semantics LTL over infinite words in one-clock timed automata relies on a combination of techniques for the analysis of probabilistic systems with properties that are specific for timed automata with a single clock. Moreover, for one-clock timed automata, we obtain a nice characterization of timed automata having non-Zeno behaviours with probability one, and show that it can be decided in NLOGSPACE if an automaton has this property.

As future works, we obviously plan to study the general case of $n$-clock timed automata. We will also look at timed games and see how probabilities can help simplify the techniques (used for instance in [dFH$^+$03,BHPR07]) for handling Zeno behaviours.

**Acknowledgment:** We are grateful to Nicolas Markey for insightful discussions about complexity classes.

## References

[AD94] R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

[ALM05] R. Alur, S. La Torre, and P. Madhusudan. Perturbed timed automata. In *Proc. 8th International Workshop on Hybrid Systems: Computation and Control (HSCC'05)*, vol. 3414 of *LNCS*, pp. 70–85. Springer, 2005.

[AMPS98] E. Asarin, O. Maler, A. Pnueli, and J. Sifakis. Controller synthesis for timed automata. In *Proc. IFAC Symposium on System Structure and Control*, pp. 469–474. Elsevier Science, 1998.

[BBB$^+$07] 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)*, LNCS. Springer, 2007. To appear.

[BHHK03] 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.

[BHPR07] Th. Brihaye, Th. A. Henzinger, V. Prabhu, and J.-F. Raskin. Minimum-time reachability in timed games. In *Proc. 34th International Colloquium on Automata, Languages and Programming (ICALP'07)*, vol. 4596 of *LNCS*, pp. 825–837. Springer, 2007.

[BMR06] 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)*, vol. 3887 of *LNCS*, pp. 238–249. Springer, 2006.

[CHR02] F. Cassez, Th. A. Henzinger, and J.-F. Raskin. A comparison of control problems for timed and hybrid systems. In *Proc. 5th International Workshop on Hybrid Systems: Computation and Control (HSCC'02)*, vol. 2289 of *LNCS*, pp. 134–148. Springer, 2002.

[DDMR04] 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)*, vol. 3253 of *LNCS*, pp. 118–133. Springer, 2004.

[DDR04] 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)*, vol. 2993 of *LNCS*, pp. 296–310. Springer, 2004.

[dFH$^+$03] L. de Alfaro, M. Faella, Th. A. Henzinger, R. Majumdar, and M. Stoelinga. The element of surprise in timed games. In *Proc. 14th International Conference on Concurrency Theory (CONCUR'03)*, vol. 2761 of *LNCS*, pp. 142–156. Springer, 2003.

[DP03] J. Desharnais and P. Panangaden. Continuous stochastic logic characterizes bisimulation of continuous-time Markov processes. *Journal of Logic and Algebraic Programming*, 56:99–115, 2003.

[GHJ97] V. Gupta, Th. A. Henzinger, and R. Jagadeesan. Robust timed automata. In *Proc. International Workshop on Hybrid and Real-Time Systems (HART'97)*, vol. 1201 of *LNCS*, pp. 331–345. Springer, 1997.

[HR00] 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)*, vol. 1790 of *LNCS*, pp. 145–159. Springer, 2000.

[KSK76] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains.* Graduate Texts in Mathematics. Springer, 1976.

[LMS04] F. Laroussinie, N. Markey, and Ph. Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. 15th International Conference on Concurrency Theory (CONCUR'04)*, vol. 3170 of *LNCS*, pp. 387–401. Springer, 2004.

[Mun00] J. R. Munkres. *Topology.* Prentice Hall, 2nd edition, 2000.

[Oxt57] J. C. Oxtoby. The Banach-Mazur game and Banach category theorem. *Annals of Mathematical Studies*, 39:159–163, 1957. Contributions to the Theory of Games, volume 3.

[Pnu77] A. Pnueli. The temporal logic of programs. In *Proc. 18th Annual Symposium on Foundations of Computer Science (FOCS'77)*, pp. 46–57. IEEE Comp. Soc. Press, 1977.

[Pnu83] A. Pnueli. On the extremely fair treatment of probabilistic algorithms. In *Proc. 15th Annual Symposium on Theory of Computing (STOC'83)*, pp. 278–290. ACM Press, 1983.

[Var85] M. Vardi. Automatic verification of probabilistic concurrent finite-state programs. In *Proc. 26th Symposium on Foundations of Computer Science (FOCS'85)*, pp. 327–338. IEEE Comp. Soc. Press, 1985.

[VV06] D. Varacca and H. Völzer. Temporal logics and model checking for fairly correct systems. In *Proc. 21st Annual Symposium on Logic in Computer Science (LICS'06)*, pp. 389–398. IEEE Comp. Soc. Press, 2006.

## Technical appendix

The proof of the next proposition will also justify the construction for the probability measure $\mathbb{P}_{\mathcal{A}}$.

**Proposition 6.** *Let $\mathcal{A}$ be a timed automaton. For every state $s$, $\mathbb{P}_{\mathcal{A}}$ is a probability measure over $(\mathsf{Runs}(\mathcal{A}, s), \Omega^s_{\mathcal{A}})$.*

*Proof.* We first recall a basic property in measure theory [KSK76].

**Proposition A** *Let $\nu$ be a non-negative additive set function defined on some set space $\mathcal{F}$ such that for every $A \in \mathcal{F}$, $\nu(A) < \infty$. The three following properties are equivalent:*

1. *$\nu$ is $\sigma$-additive,*
2. *for every sequence $(A_n)_n$ of elements of $\mathcal{F}$ such that $A_0 \subseteq A_1 \subseteq A_2 \subseteq \cdots$ and $A = \bigcup_n A_n \in \mathcal{F}$, $\lim_n \nu(A_n) = \nu(A)$,*
3. *for every sequence $(B_n)_n$ of elements of $\mathcal{F}$ such that $B_0 \supseteq B_1 \supseteq B_2 \supseteq \cdots$ and $\bigcap_n B_n = \emptyset$, $\lim_n \nu(B_n) = 0$.*

For every $n \in \mathbb{N}$, we write $\mathcal{F}_n(s)$ for the ring[20] generated by the set of (basic) cylinders from $s$ of length $n$, *i.e.*, all $\mathsf{Cyl}(\pi_{\mathcal{C}}(s, e_1 \dots e_n))$. The elements of $\mathcal{F}_n(s)$ are thus finite unions of basic cylinders of length $n$. We then define

$$\mathcal{F}(s) = \bigcup_n \mathcal{F}_n(s)$$

**Lemma B** *For every $n$, $\mathbb{P}_{\mathcal{A}}$ is a probability measure on $\mathcal{F}_n(s)$.*

*Proof.* First, by induction on $n$, it is not difficult to prove that for every $n \in \mathbb{N}$,

$$\sum_{(e_1,\dots,e_n)} \mathbb{P}_{\mathcal{A}}(\pi(s, e_1 \dots e_n)) = \mathbb{P}_{\mathcal{A}}(\pi(s)) = 1 \qquad (3)$$

We fix $n \in \mathbb{N}$. $\mathbb{P}_{\mathcal{A}}$ is obviously additive, non-negative and finite over $\mathcal{F}_n(s)$. Take a sequence $(A_i)_i$ of elements of $\mathcal{F}_n(s)$ such that $A_0 \subseteq A_1 \subseteq A_2 \subseteq \cdots$ and $A = \bigcup_i A_i \in \mathcal{F}_n(s)$. There are finitely many distinct sequences of edges of length $n$. Hence, by intersectiong each of the $A_i$'s with each of the symbolic paths $\pi(s, e_1 \dots e_n)$ of length $n$, we assume w.l.o.g. that each $A_i$ is a single constrained finite symbolic path.

Let $e_1 \dots e_n$ be the sequence of edges underlying all constrained symbolic paths $A_i$, and write $\mathcal{C}_i$ for the tightest constraint defining $A_i$ (*i.e.*, $A_i = \pi_{\mathcal{C}_i}(s, e_1 \dots e_n)$).

[20] A *ring* $R \subseteq 2^S$ is such that $\emptyset \in R$, $R$ is closed by finite union and by complement.

We have that $\mathcal{C}_i \subseteq \mathcal{C}_{i+1}$, and $(\mathcal{C}_i)_i$ converges to $\mathcal{C}$, which corresponds to the constraint associated with $A$. We can write, if $\mathbb{1}_\alpha$ is the characteristic function of set $\alpha$, that:

$$\begin{aligned}
\lim_i \mathbb{P}_{\mathcal{A}}(A_i) &= \lim_i \int_{\tau_1 \in I(s,e_1)} p_{s+\tau_1}(e_1) \int_{\tau_2 \in I(s_{\tau_1} \in I(s_{\tau_1},e_2))} p_{s_{\tau_1}+\tau_2}(e_2) \cdots \\
&\qquad \int_{\tau_n \in I(s_{\tau_1 \cdots \tau_{n-1}},e_n)} p_{s_{\tau_1 \cdots \tau_{n-1}}+\tau_n}(e_n)\, \mathbb{1}_{\mathcal{C}_i}(\tau_1,\dots,\tau_n)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{n-1}}}(\tau_n) \cdots \mathrm{d}\mu_s(\tau_1) \\
&= \int_{\tau_1 \in I(s,e_1)} p_{s+\tau_1}(e_1) \int_{\tau_2 \in I(s_{\tau_1} \in I(s_{\tau_1},e_2))} p_{s_{\tau_1}+\tau_2}(e_2) \cdots \\
&\qquad \int_{\tau_n \in I(s_{\tau_1 \cdots \tau_{n-1}},e_n)} p_{s_{\tau_1 \cdots \tau_{n-1}}+\tau_n}(e_n) \Big( \lim_i \mathbb{1}_{\mathcal{C}_i}(\tau_1,\dots,\tau_n) \Big)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{n-1}}}(\tau_n) \cdots \mathrm{d}\mu_s(\tau_1) \\
&\qquad \text{(by dominated convergence and equation (3))} \\
&= \int_{\tau_1 \in I(s,e_1)} p_{s+\tau_1}(e_1) \int_{\tau_2 \in I(s_{\tau_1} \in I(s_{\tau_1},e_2))} p_{s_{\tau_1}+\tau_2}(e_2) \cdots \\
&\qquad \int_{\tau_n \in I(s_{\tau_1 \cdots \tau_{n-1}},e_n)} p_{s_{\tau_1 \cdots \tau_{n-1}}+\tau_n}(e_n)\, \mathbb{1}_{\mathcal{C}}(\tau_1,\dots,\tau_n)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{n-1}}}(\tau_n) \cdots \mathrm{d}\mu_s(\tau_1) \\
&= \mathbb{P}_{\mathcal{A}}(A)
\end{aligned}$$

This shows that $\mathbb{P}_{\mathcal{A}}$ is a measure on $\mathcal{F}_n(s)$, for all $n \in \mathbb{N}$. It is moreover a probability measure since $\mathbb{P}_{\mathcal{A}}(\mathcal{F}_n(s)) = \mathbb{P}_{\mathcal{A}}(\pi(s)) = 1$. □

**Lemma C** *$\mathbb{P}_{\mathcal{A}}$ is a probability measure on $\mathcal{F}(s)$.*

*Proof.* Obviously $\mathbb{P}_{\mathcal{A}}$ is non-negative on $\mathcal{F}(s)$, additive (because $\mathcal{F}_n(s) \subseteq \mathcal{F}_{n+1}(s)$ for every $n \in \mathbb{N}$) and finite over $\mathcal{F}(s)$. It remains to prove that it is $\sigma$-additive. For this, we use Proposition A, and consider a sequence $(B_n)_n$ of sets in $\mathcal{F}(s)$ such that $B_0 \supseteq B_1 \supseteq B_2 \supseteq \cdots$ and $\bigcap_n B_n = \emptyset$. W.l.o.g. we assume that for every $i$, $B_n \in \mathcal{F}_n(s)$. We want to prove that $\lim_n \mathbb{P}_{\mathcal{A}}(B_n) = 0$. Applying a reasoning similar to that of [KSK76, Lemmas 2.1, 2.2, 2.3], it is sufficient to do the proof when $B_n$ is some $\mathsf{Cyl}(\pi_n)$ where $\pi_n$ is a finite (constrained) symbolic path of length $n$. We write $\mathcal{C}_n$ for the tightest constraint over variables $(\tau_i)_{i \le n}$ corresponding to $\pi_n$. We define $p_i$ the constraint from $\mathbb{R}_+^{i+1}$ onto the $i$ first components (thus in $\mathbb{R}_+^i$). Note that this projection is continuous (for the product topologies). In $\pi_n$, if $i < n$, the $i$ first variables are constrained by $\mathcal{C}_n^i = p_i(\mathcal{C}_n^{i+1})$. Moreover, for every $i \le n$, we have that

$$\mathcal{C}_{n+1}^i \subseteq \mathcal{C}_n^i \quad \text{and} \quad \mathcal{C}_n^i \subseteq \mathcal{C}_n^{i-1}$$

Fix some $i$, the sequence $(\mathcal{C}_n^i)_n$ is nested, hence converges to $\mathcal{C}^i$, and $\mathcal{C}^i \subseteq \mathcal{C}^{i-1}$. By continuity of the projection over the $i$ first components, we have that $\mathcal{C}^i = p_i(\mathcal{C}^{i+1})$. If none of the $\mathcal{C}^i$ is empty, we can thus construct an element in $\bigcap_i \mathcal{C}^i$ as follows: we take some $\tau_1$ satisfying the constraint $\mathcal{C}^1$; we have that $\mathcal{C}^1 = p_1(\mathcal{C}^2)$ (and $\mathcal{C}^2$ is a constraint over $\tau_1$ and $\tau_2$), hence there exists $\tau_2$ such that $(\tau_1, \tau_2)$ satisfies $\mathcal{C}^2$

(while $\tau_1$ still satisfies $\mathcal{C}^1$); we do the same step-by-step for all $\tau_i$ and construct a sequence $(\tau_i)_i$ which satisfies all constraints $\mathcal{C}^i$. This sequence corresponds to a run in $\bigcap_i \mathsf{Cyl}(\pi_i)$. As we assumed at the beginning of the paragraph that $\bigcap_i \mathsf{Cyl}(\pi_i) = \emptyset$, it thus means that there exists some $i \in \mathbb{N}$ such that $\mathcal{C}^i = \emptyset$.

We will use the fact that $\mathcal{C}^i = \bigcap_{n \geq i} \mathcal{C}^i_n$ is empty to prove that $\lim_n \mathbb{P}_{\mathcal{A}}(\pi_n) = 0$. We write, still with the notation that $\mathbb{1}_\alpha$ is the characteristic function of set $\alpha$:

$$\begin{aligned}\mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi_n)) = &\int_{\tau_1 \in I(s,e_1)} p_{s+\tau_1}(e_1) \int_{\tau_2 \in I(s_{\tau_1} \in I(s_{\tau_1},e_2))} p_{s_{\tau_1}+\tau_2}(e_2) \cdots \\ &\quad \int_{\tau_n \in I(s_{\tau_1 \cdots \tau_{n-1}},e_n)} p_{s_{\tau_1 \cdots \tau_{n-1}}+\tau_n}(e_n)\, \mathbb{1}_{\mathcal{C}_n}(\tau_1, \ldots, \tau_n)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{n-1}}}(\tau_n) \cdots \mathrm{d}\mu_s(\tau_1) \\ \leq &\int_{\tau_1 \in I(s,e_1)} \int_{\tau_2 \in I(s_{\tau_1} \in I(s_{\tau_1},e_2))} \cdots \int_{\tau_i \in I(s_{\tau_1 \cdots \tau_{i-1}},e_i)} \mathbb{1}_{\mathcal{C}^i_n}(\tau_1, \ldots, \tau_i)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{i-1}}}(\tau_i) \cdots \mathrm{d}\mu_s(\tau_1)\end{aligned}$$

Applying the dominated convergence theorem, we get that:

$$\begin{aligned}\lim_n \mathbb{P}_{\mathcal{A}}(\mathsf{Cyl}(\pi_n)) &= \int_{\cdots} \int_{\cdots} \cdots \int_{\cdots} \Big( \lim_n \mathbb{1}_{\mathcal{C}^i_n}(\tau_1, \ldots, \tau_i) \Big)\, \mathrm{d}\mu_{s_{\tau_1 \cdots \tau_{i-1}}}(\tau_i) \cdots \mathrm{d}\mu_s(\tau_1) \\ &= 0\end{aligned}$$

This concludes the proof that $\mathbb{P}_{\mathcal{A}}$ is $\sigma$-additive on $\mathcal{F}$, and thus the proof that $\mathbb{P}_{\mathcal{A}}$ is a probability measure on $\mathcal{F}(s)$. □

We conclude the proof using the following classical measure extension theorem:

**Theorem D (Carathéodory's extension theorem)** *Let $S$ be a set, and $\nu$ a $\sigma$-finite measure defined on a ring $R \subseteq 2^S$. Then, $\nu$ can be extended in a unique manner to the $\sigma$-algebra generated by $R$.*

We apply Theorem D to the set $S = \mathsf{Runs}(\mathcal{A}, s)$, $R = \mathcal{F}(s)$, and $\nu = \mathbb{P}_{\mathcal{A}}$ which is a $\sigma$-finite measure on $\mathcal{F}(s)$. Hence, there is a unique extension of $\mathbb{P}_{\mathcal{A}}$ on $\Omega^s_{\mathcal{A}}$, the $\sigma$-algebra generated by the cylinders, which is a probability measure. □

---

We now detail a proof omitted in the main text and which serves proving Proposition 25 (see page 22).

**Lemma E** *If* $\gamma_i^N = \frac{1}{1-t_{i-1}} \int_{t_i=t_{i-1}}^1 \frac{1}{2-t_i} \cdot \frac{1}{1-t_i} \int_{t_{i+1}=t_i}^1 \cdots \frac{1}{2-t_{N-1}} \cdot \frac{1}{1-t_{N-1}} \int_{t_N=t_{N-1}}^1 \mathrm{d}t_N \ldots \mathrm{d}t_i$, *then:*

$$\gamma_i^N \;\geq\; \frac{2^{N+1-i}-1}{2^{N-i}} - \frac{2^{N-i}-1}{2^{N-i}} \cdot (2 - t_{i-1}).$$

*Proof.* The base case is when $i = N$. In that case,

$$\gamma_N^N = \frac{1}{1-t_{N-1}} \int_{t_N=t_{N-1}}^{1} \mathrm{d}t_N = 1$$

which proves the property.

We assume we have proved the property for $i+1$, and want to prove it for $i$.

$$\begin{aligned}
\gamma_i^N &= \frac{1}{1-t_{i-1}} \int_{t_i=t_{i-1}}^{1} \frac{1}{2-t_i} \cdot \gamma_{i+1}^N \,\mathrm{d}t_i \\
&\geq \frac{1}{1-t_{i-1}} \int_{t_i=t_{i-1}}^{1} \frac{1}{2-t_i} \cdot \left( \frac{2^{N-i}-1}{2^{N-i-1}} - \frac{2^{N-i-1}-1}{2^{N-i-1}} \cdot (2-t_i) \right) \mathrm{d}t_i \qquad \text{(by i.h.)} \\
&\geq \frac{1}{1-t_{i-1}} \left[ -\frac{2^{N-i}-1}{2^{N-i-1}} \cdot \log(2-t_i) - \frac{2^{N-i-1}-1}{2^{N-i-1}} \cdot t_i \right]_{t_i=t_{i-1}}^{1} \\
&\geq \frac{1}{1-t_{i-1}} \left( \frac{2^{N-i}-1}{2^{N-i-1}} \cdot \log(2-t_{i-1}) - \frac{2^{N-i-1}-1}{2^{N-i-1}} \cdot (1-t_{i-1}) \right)
\end{aligned}$$

Now, when $0 \leq x \leq 1$ we know that $\log(1+x) \geq x - \frac{x^2}{2}$ (see Lemma F). Applying this inequality to $x = 1 - t_{i-1}$, we get the following inequality:

$$\begin{aligned}
\gamma_i^N &\geq \frac{1}{1-t_{i-1}} \left( \frac{2^{N-i}-1}{2^{N-i-1}} \cdot \left( (1-t_{i-1}) - \frac{(1-t_{i-1})^2}{2} \right) - \frac{2^{N-i-1}-1}{2^{N-i-1}} \cdot (1-t_{i-1}) \right) \\
&\geq \left( \frac{2^{N-i}-1}{2^{N-i-1}} - \frac{2^{N-i-1}-1}{2^{N-i-1}} \right) - \frac{2^{N-i}-1}{2^{N-i}} \cdot (1-t_{i-1}) \\
&\geq \frac{2^{N-i+1}-1}{2^{N-i}} - \frac{2^{N-i}-1}{2^{N-i}} \cdot (2-t_{i-1})
\end{aligned}$$

This concludes the inductive case. □

**Lemma F** *Let* $0 \leq x \leq 1$. *Then* $\log(1+x) \geq x - \frac{x^2}{2}$.

*Proof.* First observe that functions $t \rightarrow \frac{1}{1+t}$ and $t \rightarrow 1 - t + \frac{t^2}{1+t}$ coincide on $\mathbb{R} \setminus \{-1\}$. This can easily be checked by developping the second function. Let now $0 \leq x \leq 1$; the integrales of both functions on the interval $[0, x]$ are equal:

$$\int_{t=0}^{x} \frac{1}{1+t} \mathrm{d}t = \int_{t=0}^{x} (1 - t + \frac{t^2}{1+t}) \mathrm{d}t\,.$$

Computing the first integral, and simplifying the second, we deduce:

$$\log(1+x) = x - \frac{x^2}{2} + \int_{t=0}^{x} \frac{t^2}{1+t} \mathrm{d}t\,.$$

For all $0 \leq t \leq x$, we have $\frac{1}{1+t} \geq \frac{1}{1+x}$. Hence $\int_{t=0}^{x} \frac{t^2}{1+t} \mathrm{d}t \geq \frac{1}{1+x} \int_{t=0}^{x} t^2 \mathrm{d}t = \frac{x^3}{3(1+x)}$. Since $\frac{x^3}{3(1+x)} \geq 0$ for all $x \in [0,1]$, we obtain the desired inequality: $\log(1+x) \geq x - \frac{x^2}{2}$. □