

Quantitative Games on Probabilistic Timed Automata

Marta Kwiatkowska, Gethin Norman, and Ashutosh Trivedi

Oxford University Computing Laboratory, Parks Road, Oxford, OX1 3QD

Abstract. Two-player zero-sum games are a well-established model for synthesising controllers that optimise some performance criterion. In such games one player represents the controller, while the other describes the (adversarial) environment, and controller synthesis corresponds to computing the optimal strategies of the controller for a given criterion. Asarin and Maler initiated the study of quantitative games on (non-probabilistic) timed automata by synthesising controllers which optimise the time to reach a final state. The correctness and termination of their approach was dependent on exploiting the properties of a special class of functions, called *simple functions*, that can be finitely represented. In this paper we consider quantitative games over probabilistic timed automata. Since the concept of simple functions is not sufficient to solve games in this setting, we generalise simple functions to so-called *quasi-simple functions*. Then, using this class of functions, we demonstrate that the problem of solving games with either expected reachability-time or expected discounted-time criteria on probabilistic timed automata are in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$. As a side result, we also show that the problem of solving discounted-time games on timed automata is EXPTIME-complete.

1 Introduction

Two-player zero-sum games on finite automata, as a mechanism for supervisory controller synthesis of discrete event systems, were introduced by Ramadge and Wonham [24]. In this setting the two players—called Min and Max—represent the ‘controller’ and the ‘environment’ and control-program synthesis corresponds to finding a winning (or optimal) strategy of the ‘controller’ for some given performance objective. If the objectives are dependent on time, e.g. when the objective corresponds to completing a given set of tasks within some deadline, then games on timed automata are a well-established approach for controller synthesis, see for example [3,1,8,6].

In this paper we extend this approach to systems which are quantitative in terms of time *and* probabilistic behaviour. Probabilistic information is important for modelling, e.g., faulty or unreliable components, the random coin flips of distributed communication and security protocols, and performance characteristics. We consider games on probabilistic timed automata [22,15,5], a modelling framework for real-time systems exhibiting both nondeterministic and probabilistic behaviour. We concentrate on expected reachability-time games, where the performance objective concerns the expected minimum time the controller can ensure for the system to reach a target, regardless of uncontrollable (environmental) events. This approach has many practical applications, including job-shop scheduling, where machines can be faulty or have

variable execution times, and both routing and task graph scheduling problems, where both time and stochastic behaviour is also relevant. We also discuss discounted-time games where, intuitively, at each transition the system breaks down with some non-zero probability, and the players try to optimise the expected time to breakdown.

Contributions. Our approach is inspired by the work of Asarin and Maler [3] who initiated the study of quantitative games on (non-probabilistic) timed automata. Their results were dependent on exploiting the properties of a special class of functions, called *simple functions*, that can be finitely represented. Since the concept of simple functions is not sufficient to solve games in this setting, we generalise simple functions to so-called *quasi-simple functions*. Using this class of functions and the boundary region graph construction [20], we demonstrate that the problem of solving games with either expected reachability-time or expected discounted-time criteria on probabilistic timed automata are in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$. As a side result, we also show that the problem of solving discounted-time games on timed automata is EXPTIME-complete.

Related Work. Hoffman and Wong-Toi [14] were the first to define and solve optimal controller synthesis problem for timed automata. For a detailed introduction to the topic of qualitative games on timed automata, see e.g. [4]. Asarin and Maler [3] initiated the study of quantitative games on timed automata by providing a symbolic algorithm to solve reachability-time games. The works of [10] and [18] showed that the decision version of the reachability-time game is EXPTIME-complete for timed automata with at least two clocks. For average-time objectives, Jurdziński and Trivedi [19] showed the EXPTIME-completeness of the problem for timed automata with two or more clocks.

A natural extension of reachability-time games for timed automata is reachability-price games for priced timed automata. Alur, Bernadsky, and Madhusudan [1] and Bouyer et al. [8] gave semi-algorithms to compute the value of reachability-price games on linearly-priced timed automata. In [11] and [7] it was shown that checking the existence of optimal strategies in a reachability-price game is undecidable for automata with three clocks and stopwatch prices.

We are not aware of any previous work studying games on probabilistic timed automata. For a significantly different model of stochastic timed games, [9] show that deciding whether a target is reachable within a given probability bound is undecidable. Regarding one-player games on probabilistic timed automata, [16] shows that a number of one-player optimisation problems on concavely-priced probabilistic timed automata can be reduced to solving corresponding problems on the boundary region graph. We also mention [21], based on the digital clocks approach [13], which solves expected-time (and expected-cost) reachability for a subclass of probabilistic timed automata.

2 Preliminaries

We begin by presenting the background material required in the remainder of the paper. We assume, the sets \mathbb{N} of non-negative integers, \mathbb{R} of reals and \mathbb{R}_{\oplus} of non-negative reals. For $n \in \mathbb{N}$, let $\llbracket n \rrbracket_{\mathbb{N}}$ and $\llbracket n \rrbracket_{\mathbb{R}}$ denote the sets $\{0, 1, \dots, n\}$, and $\{r \in \mathbb{R} \mid 0 \leq r \leq n\}$ respectively. For $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we define $\|x\|_{\infty} = \max \{|x_i| \mid 1 \leq i \leq n\}$.

Probability distributions. A *discrete probability distribution* over a countable set Q is a function $\mu : Q \rightarrow [0, 1]$ such that $\sum_{q \in Q} \mu(q) = 1$. For a possible uncountable set Q' , we define $\mathcal{D}(Q')$ to be the set of functions $\mu : Q' \rightarrow [0, 1]$ such that the set $\text{supp}(\mu) = \{q \in Q' \mid \mu(q) > 0\}$ is countable and, over $\text{supp}(\mu)$, μ is a distribution. We say that $\mu \in \mathcal{D}(Q)$ is a *point distribution* if $\mu(q) = 1$ for some $q \in Q$.

Markov decision processes. We next introduce Markov decision processes a modelling formalism for systems exhibiting nondeterministic and probabilistic behaviour.

Definition 1. A Markov decision process (MDP) is a tuple $M = (S, F, A, p, \pi)$ where:

- S is the set of states including a set of final states F ;
- A is the set of actions;
- $p : S \times A \rightarrow \mathcal{D}(S)$ is a partial function called the probabilistic transition function;
- $\pi : S \times A \rightarrow \mathbb{R}_{\oplus}$ is the reward function.

We write $A(s)$ for the set of actions available at s , i.e., the set of actions a for which $p(s, a)$ is defined. In an MDP \mathcal{M} , if the current state is s , then there is a non-deterministic choice between the actions in $A(s)$ and if action a is chosen the probability of reaching the state $s' \in S$ equals $p(s'|s, a) \stackrel{\text{def}}{=} p(s, a)(s')$.

Clocks, clock valuations, regions and zones. We fix a constant $k \in \mathbb{N}$ and finite set of *clocks* \mathcal{C} . A (k -bounded) *clock valuation* is a function $\nu : \mathcal{C} \rightarrow \llbracket k \rrbracket_{\mathbb{R}}$ and we write V for the set of clock valuations.

Assumption 1. Although clocks are usually allowed to take arbitrary non-negative values, we have restricted their values to be bounded by the constant k . This restriction is for technical convenience and comes without significant loss of generality.

If $\nu \in V$ and $t \in \mathbb{R}_{\oplus}$ then we write $\nu + t$ for the clock valuation defined by $(\nu + t)(c) = \nu(c) + t$, for all $c \in \mathcal{C}$. For $C \subseteq \mathcal{C}$, we write $\nu[C:=0]$ for the clock valuation where $\nu[C:=0](c) = 0$ if $c \in C$, and $\nu[C:=0](c) = \nu(c)$ otherwise. Let $X \subseteq V$ be a convex subset of clock valuations and let $F : X \rightarrow \mathbb{R}$ be a continuous function. We write \bar{F} for the unique continuous function $\bar{F} : \bar{X} \rightarrow \mathbb{R}$, such that $\bar{F}(\nu) = F(\nu)$ for all $\nu \in X$.

The set of *clock constraints* over \mathcal{C} is the set of conjunctions of *simple constraints*, which are constraints of the form $c \bowtie i$ or $c - c' \bowtie i$, where $c, c' \in \mathcal{C}$, $i \in \llbracket k \rrbracket_{\mathbb{N}}$, and $\bowtie \in \{<, >, =, \leq, \geq\}$. For every $\nu \in V$, let $\text{SCC}(\nu)$ be the set of simple constraints which hold in ν . A *clock region* is a maximal set $\zeta \subseteq V$, such that $\text{SCC}(\nu) = \text{SCC}(\nu')$ for all $\nu, \nu' \in \zeta$. Every clock region is an equivalence class of the indistinguishability-by-clock-constraints relation, and vice versa. Note that ν and ν' are in the same clock region if and only if the integer parts of the clocks and the partial orders of the clocks, determined by their fractional parts, are the same in ν and ν' . We write $[\nu]$ for the clock region of ν and, if $\zeta = [\nu]$, write $\zeta[C:=0]$ for the clock region $[\nu[C:=0]]$.

A *clock zone* is a convex set of clock valuations, which is a union of a set of clock regions. We write \mathcal{Z} for the set of clock zones. For any clock zone W and clock valuation ν , we use the notation $\nu \in W$ to denote that $[\nu] \in W$. A set of clock valuations is a clock zone if and only if it is definable by a clock constraint. For $W \subseteq V$, we write \bar{W} for the smallest closed set in V containing W . Observe that, for every clock zone W , the set \bar{W} is also a clock zone.

3 Stochastic Games on Probabilistic Timed Automata

In this section we introduce stochastic games played on probabilistic timed automata.

Probabilistic timed automata. Probabilistic timed automata are a modelling framework for real-time systems exhibiting both nondeterministic and probabilistic behaviour. The formalism is derived by extending classical timed automata [2] with discrete probability distributions over edges.

Definition 2 (PTA syntax). A probabilistic timed automaton (PTA) is a tuple $\mathsf{T} = (L, L_F, \mathcal{C}, \text{Inv}, \text{Act}, E, \delta)$ where:

- L is the finite set of locations including the set of final locations L_F ;
- \mathcal{C} is the finite set of clocks;
- $\text{Inv} : L \rightarrow \mathcal{Z}$ is the invariant condition;
- Act is the finite set of actions;
- $E : L \times \text{Act} \rightarrow \mathcal{Z}$ is the action enabledness function;
- $\delta : (L \times \text{Act}) \rightarrow \mathcal{D}(2^{\mathcal{C}} \times L)$ is the transition probability function.

A *timed automaton* is a PTA with the property that $\delta(\ell, a)$ is a point distribution for all $\ell \in L$ and $a \in \text{Act}$. When we consider a PTA as an input of an algorithm, its size should be understood as the sum of the sizes of encodings of L , \mathcal{C} , Inv , Act , E , and δ . As usual [17], we assume that probabilities are expressed as ratios of two natural numbers, each written in binary. In addition, we assume the following standard restriction on PTAs which ensures *time divergent* behaviour.

Assumption 2. We restrict attention to *structurally non-Zeno* PTAs [26,17].

A *configuration* of a PTA T is a pair (ℓ, ν) , where $\ell \in L$ is a location and $\nu \in V$ is a clock valuation over \mathcal{C} such that $\nu \in \text{Inv}(\ell)$. For any $t \in \mathbb{R}$, we let $(\ell, \nu) + t$ equal the configuration $(\ell, \nu + t)$. Informally, the behaviour of a PTA is as follows. In configuration (ℓ, ν) time passes before an available action is triggered, after which a discrete probabilistic transition occurs. Time passage is available only if the invariant condition $\text{Inv}(\ell)$ is satisfied while time elapses, and an action a can be chosen after time t elapses only if it is enabled after time elapse, i.e., if $\nu + t \in E(\ell, a)$. Both the time and the action chosen are nondeterministic. If the action a is chosen, then the probability of moving to the location ℓ' and resetting all of the clocks in \mathcal{C} to 0 is given by $\delta[\ell, a](\mathcal{C}, \ell')$.

Formally, the semantics of a PTA is given by an MDP which has both an infinite number of states and an infinite number of transitions.

Definition 3 (PTA semantics). Let $\mathsf{T} = (L, L_F, \mathcal{C}, \text{Inv}, \text{Act}, E, \delta)$ be a PTA. The semantics of T is the MDP $\llbracket \mathsf{T} \rrbracket = (S, F, A, p, \pi)$ where

- $S \subseteq L \times V$, the set of states, is such that $(\ell, \nu) \in S$ if and only if $\nu \in \text{Inv}(\ell)$;
- $F = S \cap (L_F \times V)$ is the set of final states;
- $A = \mathbb{R}_{\oplus} \times \text{Act}$ is the set of timed actions;
- $p : S \times A \rightarrow \mathcal{D}(S)$ is the probabilistic transition function such that for $(\ell, \nu) \in S$ and $(t, a) \in A$, we have $p((\ell, \nu), (t, a)) = \mu$ if and only if

- $\nu+t' \in \text{Inv}(\ell)$ for all $t' \in [0, t]$;
 - $\nu+t \in E(\ell, a)$;
 - $\mu(\ell', \nu') = \sum_{C \subseteq \mathcal{C} \wedge (\nu+t) \mid_{C:=0} = \nu'} \delta[\ell, a](C, \nu')$ for all $(\ell', \nu') \in S$.
- $\pi : S \times A \rightarrow \mathbb{R}$ is the reward function where $\pi(s, (t, a)) = t$ for $s \in S$ and $(t, a) \in A$.

Probabilistic Timed Game Arena. We are now in a position to introduce probabilistic timed game arenas.

Definition 4. A probabilistic timed game arena is a triplet $\mathcal{T} = (\mathbb{T}, L_{\text{Min}}, L_{\text{Max}})$ where $\mathbb{T} = (L, L_F, \mathcal{C}, \text{Inv}, \text{Act}, E, \delta)$ is a PTA and $(L_{\text{Min}}, L_{\text{Max}})$ is a partition of L .

The semantics of a probabilistic timed game arena \mathcal{T} is the stochastic game arena $\llbracket \mathcal{T} \rrbracket = (\llbracket \mathbb{T} \rrbracket, S_{\text{Min}}, S_{\text{Max}})$ where $\llbracket \mathbb{T} \rrbracket = (S, A, F, p, \pi)$ is the semantics of \mathbb{T} , and $S_{\text{Min}} = S \cap (L_{\text{Min}} \times V)$ and $S_{\text{Max}} = S \setminus S_{\text{Min}}$. Intuitively S_{Min} is the set of states controlled by player Min, and S_{Max} is the set of states controlled by player Max.

In a turn-based game on \mathcal{T} players Min and Max move a token along the states of the PTA in the following manner. If the current state is s , then the player controlling the state chooses an action $(t, a) \in A(s)$ after which state $s' \in S$ is reached with probability $p(s'|s, a)$. In the next turn the player controlling the state s' chooses an action in $A(s')$ and a probabilistic transition is made accordingly.

We say that $(s, (t, a), s')$ is a transition in \mathcal{T} if $p(s'|s, (t, a)) > 0$ and a *play* of \mathcal{T} is a sequence $\langle s_0, (t_1, a_1), s_1, \dots \rangle \in S \times (A \times S)^*$ such that $(s_i, (t_{i+1}, a_{i+1}), s_{i+1})$ is a transition for all $i \geq 0$. We write *Play* (*FPlay*) for the sets of infinite (finite) plays and Play_s (FPlay_s) for the sets of infinite (finite) plays starting from state s . For a finite play r let $\text{last}(r)$ denote the last state of the play. Let X_i and Y_i denote the random variables corresponding to i^{th} state and action of a play.

A *strategy* of player Min in \mathcal{T} is a partial function $\mu : \text{FPlay} \rightarrow \mathcal{D}(A)$, defined for $r \in \text{FPlay}$ if and only if $\text{last}(r) \in S_{\text{Min}}$, such that $\text{supp}(\mu(r)) \subseteq A(\text{last}(r))$. A strategy of player Max can be defined analogously. We write Σ_{Min} and Σ_{Max} for the set of strategies of players Min and Max, respectively. Let $\text{Play}_s^{\mu, \chi}$ denote the subset of Play_s which corresponds to the set of plays in which players play according to μ and χ , respectively. A strategy σ is *pure* if $\sigma(r)$ is a point distribution for all $r \in \text{FPlay}$ for which it is defined, while it is *stationary* if $\text{last}(r) = \text{last}(r')$ implies $\sigma(r) = \sigma(r')$ for all $r, r' \in \text{FPlay}$.

To analyse the behaviour of a stochastic game on \mathcal{T} under a strategy pair (μ, χ) , for every state s of \mathcal{T} , we define a probability space $(\text{Play}_s^{\mu, \chi}, \mathcal{F}_{\text{Play}_s^{\mu, \chi}}, \text{Prob}_s^{\mu, \chi})$ over the set of infinite plays under strategies μ and χ with s as the initial state. Given a *real-valued random variable* $f : \text{Play} \rightarrow \mathbb{R}$, we can then define the expectation of this variable $\mathbb{E}_s^{\mu, \chi} \{f\}$ with respect to strategy pair (μ, χ) when starting in s .

For technical convenience we make the following standard [23] assumption (a similar assumption is required for optimal expected reachability price problem for finite MDP [12]):

Assumption 3. For every strategy pair $\mu \in \Sigma_{\text{Min}}, \chi \in \Sigma_{\text{Max}}$, and state $s \in S$ we have that $\lim_{i \rightarrow \infty} \text{Prob}_s^{\mu, \chi}(X_i \in F) = 1$.

Expected Reachability-Time Game. In an expected reachability-time game on $\mathcal{T} = (\mathbb{T}, L_{\text{Min}}, L_{\text{Max}})$ player Min attempts to reach the final states as quickly as possible, while the objective of player Max is the opposite. More precisely, Min is interested in minimising her losses, while player Max is interested in maximising his winnings where, if player Min uses the strategy $\mu \in \Sigma_{\text{Min}}$ and player Max uses the strategy $\chi \in \Sigma_{\text{Max}}$, player Min loses the following amount to player Max:

$$\text{EReach}(s, \mu, \chi) \stackrel{\text{def}}{=} \mathbb{E}_s^{\mu, \chi} \left\{ \sum_{i=1}^{\min\{i \mid X_i \in F\}} \pi(X_{i-1}, Y_i) \right\}.$$

Observe that player Max can choose his actions to win at least an amount arbitrarily close to $\sup_{\chi \in \Sigma_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}(s, \mu, \chi)$. This is called the *lower value* $\underline{\text{Val}}(s)$ of the expected reachability-time game starting at s :

$$\underline{\text{Val}}(s) \stackrel{\text{def}}{=} \sup_{\chi \in \Sigma_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}(s, \mu, \chi).$$

Similarly, player Min can choose to lose at most an amount arbitrarily close to $\inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \text{EReach}(s, \mu, \chi)$. This is called the *upper value* $\overline{\text{Val}}(s)$ of the game:

$$\overline{\text{Val}}(s) \stackrel{\text{def}}{=} \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \text{EReach}(s, \mu, \chi).$$

It is straightforward to verify that $\underline{\text{Val}}(s) \leq \overline{\text{Val}}(s)$ for all $s \in S$. We say that the expected reachability-time game is determined if $\underline{\text{Val}}(s) = \overline{\text{Val}}(s)$ for all $s \in S$. In this case we also say that the value of the game exists and denote it by $\text{Val}(s) = \underline{\text{Val}}(s) = \overline{\text{Val}}(s)$ for all $s \in S$. The results of this paper present an alternative proof of the following proposition.

Proposition 5. *Expected reachability-time games are determined.*

For $\mu \in \Sigma_{\text{Min}}$ and $\chi \in \Sigma_{\text{Max}}$ we define $\text{Val}^\mu(s) = \sup_{\chi \in \Sigma_{\text{Max}}} \text{EReach}(s, \mu, \chi)$ and $\text{Val}_\chi(s) = \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}(s, \mu, \chi)$. For an $\varepsilon > 0$, we say that $\mu \in \Sigma_{\text{Min}}$ or $\chi \in \Sigma_{\text{Max}}$ is ε -optimal if $\text{Val}^\mu(s) \leq \text{Val}(s) + \varepsilon$ or $\text{Val}_\chi(s) \geq \text{Val}(s) - \varepsilon$, respectively, for all $s \in S$. If an expected reachability-time game is determined, then for every $\varepsilon > 0$, both players have ε -optimal strategies.

Optimality Equations. We now review optimality equations for characterising the value in an expected reachability-time game. Let \mathcal{T} be a probabilistic timed game arena and let $P : S \rightarrow \mathbb{R}_\oplus$. We say that P is a solution of optimality equations $\text{Opt}(\mathcal{T})$, and we write $P \models \text{Opt}(\mathcal{T})$ if, for all $s \in S$:

$$P(s) = \begin{cases} 0 & \text{if } s \in F \\ \inf_{(t,a) \in A(s)} \{t + \sum_{s' \in S} p(s'|s, (t,a)) \cdot P(s')\} & \text{if } s \in S_{\text{Min}} \setminus F \\ \sup_{(t,a) \in A(s)} \{t + \sum_{s' \in S} p(s'|s, (t,a)) \cdot P(s')\} & \text{if } s \in S_{\text{Max}} \setminus F \end{cases}$$

Under Assumption 3, the proof of the following proposition is routine and for details see, for example, [18].

Proposition 6. *If $P \models \text{Opt}(\mathcal{T})$, then $\text{Val}(s) = P(s)$ for all $s \in S$ and, for every $\varepsilon > 0$, both players have pure ε -optimal strategies.*

Using Proposition 6, it follows that the problem of solving an expected reachability-time game on \mathcal{T} can be reduced to solving the optimality equations $\text{Opt}(\mathcal{T})$. In the non-probabilistic setting, Jurdziński and Trivedi [18] showed that solving optimality equations for a reachability-time game on a (non-probabilistic) timed automaton \mathcal{T} can be reduced to solving a reachability-price game on an abstraction, called the boundary region graph. Recently [16], we extended this result reducing a number of one-player optimisation problems on probabilistic timed automata to solving corresponding problems on boundary region graphs. In the next section, we review boundary region graph abstraction for probabilistic timed automata and, in Section 6, we argue that boundary region graph abstraction for probabilistic timed automata is sufficient to solve expected reachability-time games. In Section 7, we explain that expected discounted-time games on probabilistic timed automata can also be reduced to solving discounted-price games on their boundary region graph. In Section 8, we discuss some implications of these reductions on the complexity of the decision problems related to these games.

4 The Boundary Region Graph Abstraction

In this section we review the boundary region graph for PTAs introduced in [16].

Regions. A *region* is a pair (ℓ, ζ) , where ℓ is a location and ζ is a clock region such that $\zeta \subseteq \text{Inv}(\ell)$. For any $s=(\ell, \nu)$, we write $[s]$ for the region $(\ell, [\nu])$ and \mathcal{R} for the set of regions. A set $Z \subseteq L \times V$ is a *zone* if, for every $\ell \in L$, there is a clock zone W_ℓ (possibly empty), such that $Z = \{(\ell, \nu) \mid \ell \in L \wedge \nu \in W_\ell\}$. For a region $R=(\ell, \zeta) \in \mathcal{R}$, we write \overline{R} for the zone $\{(\ell, \nu) \mid \nu \in \overline{\zeta}\}$, recall $\overline{\zeta}$ is the smallest closed set in V containing ζ .

For $R, R' \in \mathcal{R}$, we say that R' is in the future of R , or that R is in the past of R' , if there is $s \in R$, $s' \in R'$ and $t \in \mathbb{R}_{\oplus}$ such that $s' = s+t$; we then write $R \rightarrow_* R'$. We say that R' is the *time successor* of R if $R \rightarrow_* R'$, $R \neq R'$, and $R \rightarrow_* R'' \rightarrow_* R'$ implies $R''=R$ or $R''=R'$ and write $R \rightarrow_{+1} R'$ and $R' \leftarrow_{+1} R$.

We say that a region $R \in \mathcal{R}$ is *thin* if $[s] \neq [s+\varepsilon]$ for every $s \in R$ and $\varepsilon > 0$; other regions are called *thick*. We write $\mathcal{R}_{\text{Thin}}$ and $\mathcal{R}_{\text{Thick}}$ for the sets of thin and thick regions, respectively. Note that if $R \in \mathcal{R}_{\text{Thick}}$ then, for every $s \in R$, there is an $\varepsilon > 0$, such that $[s] = [s+\varepsilon]$. Observe that the time successor of a thin region is thick, and vice versa.

We say $(\ell, \nu) \in L \times V$ is in the *closure of the region* (ℓ, ζ) , and we write $(\ell, \nu) \in \overline{(\ell, \zeta)}$, if $\nu \in \overline{\zeta}$. For any $\nu \in V$, $b \in \llbracket k \rrbracket_{\mathbb{N}}$ and $c \in \mathcal{C}$ such that $\nu(c) \leq b$, we let $\text{time}(\nu, (b, c)) \stackrel{\text{def}}{=} b - \nu(c)$. Intuitively, $\text{time}(\nu, (b, c))$ returns the amount of time that must elapse in ν before the clock c reaches the integer value b . Note that, for any $(\ell, \nu) \in L \times V$ and $a \in \text{Act}$, if $t = \text{time}(\nu, (b, c))$ is defined, then $(\ell, [\nu+t]) \in \mathcal{R}_{\text{Thin}}$ and $\text{supp}(p_{\mathcal{T}}(\cdot \mid (\ell, \nu), (t, a))) \subseteq \mathcal{R}_{\text{Thin}}$. Observe that, for every $R' \in \mathcal{R}_{\text{Thin}}$, there is a number $b \in \llbracket k \rrbracket_{\mathbb{N}}$ and a clock $c \in \mathcal{C}$, such that, for every $R \in \mathcal{R}$ in the past of R' , we have $s \in R$ implies $(s+(b-s(c)) \in R'$; and we write $R \rightarrow_{b,c} R'$.

The Boundary Region Graph. The boundary region graph is motivated by the following. Consider any $a \in \text{Act}$, $s = (\ell, \nu)$ and $R = (\ell, \zeta) \rightarrow_* R' = (\ell, \zeta')$ such that $s \in R$ and $R' \in E(\ell, a)$.

- If $R' \in \mathcal{R}_{\text{Thick}}$, then there are infinitely many $t \in \mathbb{R}_{\oplus}$ such that $s+t \in R'$. However, amongst all such t 's, for one of the boundaries of ζ' , the closer $\nu+t$ is to this boundary, the ‘better’ the timed action (t, a) becomes for a player’s objective. However, since R' is a thick region, the set $\{t \in \mathbb{R}_{\oplus} \mid s+t \in R'\}$ is an open interval, and hence does not contain its boundary values. Observe that the infimum equals $b_- - \nu(c_-)$ where $R \xrightarrow{b_-, c_-} R_- \xrightarrow{+1} R'$ and the supremum equals $b_+ - \nu(c_+)$ where $R \xrightarrow{b_+, c_+} R_+ \xrightarrow{-1} R'$. In the boundary region graph we include these ‘best’ timed actions through the actions $((b_-, c_-, a), R')$ and $((b_+, c_+, a), R')$.
- If $R' \in \mathcal{R}_{\text{Thin}}$, then there exists a unique $t \in \mathbb{R}_{\oplus}$ such that $(\ell, \nu+t) \in R'$. Moreover since R' is a thin region, there exists a clock $c \in C$ and a number $b \in \mathbb{N}$ such that $R \xrightarrow{b, c} R'$ and $t = b - \nu(c)$. In the boundary region graph we summarise this ‘best’ timed action from region R via region R' through the action $((b, c, a), R')$.

Based on this intuition the boundary region graph is defined as follows.

Definition 7. Let $\mathsf{T} = (L, L_F, C, Inv, Act, E, \delta)$ be a PTA. The boundary region graph of T is defined as the MDP $\widehat{\mathsf{T}} = (\widehat{S}, \widehat{F}, \widehat{A}, \widehat{p}, \widehat{\pi})$ where

- $\widehat{S} = \{((\ell, \nu), (\ell, \zeta)) \mid (\ell, \zeta) \in \mathcal{R} \wedge \nu \in \bar{\zeta}\}$ and $\widehat{F} = \{((\ell, \nu), (\ell, \zeta)) \in \widehat{S} \mid \ell \in L_F\}$;
- the finite set of boundary actions $\widehat{A} \subseteq (\llbracket k \rrbracket_{\mathbb{N}} \times C \times Act) \times \mathcal{R}$ and for $R \in \mathcal{R}$ we let¹ $\widehat{A}(R) = \{\alpha \in \widehat{A}((\ell, \nu), R) \mid ((\ell, \nu), R) \in \widehat{S}\}$;
- for any state $((\ell, \nu), (\ell, \zeta)) \in \widehat{S}$ and action $((b, c, a), (\ell, \zeta_a)) \in \widehat{A}$ we have $\widehat{p}((\ell, \nu), (\ell, \zeta), ((b, c, a), (\ell, \zeta_a))) = \mu$ if and only if

$$\mu((\ell', \nu'), (\ell', \zeta')) = \sum_{C \subseteq C \wedge \nu_a[C:=0] = \nu' \wedge \zeta_a[C:=0] = \zeta'} \delta[\ell, a](C, \ell')$$

for all $((\ell', \nu'), (\ell', \zeta')) \in \widehat{S}$ where $\nu_a = \nu + \text{time}(\nu, (b, c))$ and one of the following conditions holds:

- $(\ell, \zeta) \xrightarrow{b, c} (\ell, \zeta_a)$ and $\zeta_a \in E(\ell, a)$
 - $(\ell, \zeta) \xrightarrow{b, c} (\ell, \zeta_-) \xrightarrow{+1} (\ell, \zeta_a)$ for some (ℓ, ζ_-) and $\zeta_a \in E(\ell, a)$
 - $(\ell, \zeta) \xrightarrow{b, c} (\ell, \zeta_+) \xrightarrow{-1} (\ell, \zeta_a)$ for some (ℓ, ζ_+) and $\zeta_a \in E(\ell, a)$.
- $\widehat{\pi} : \widehat{S} \times \widehat{A} \rightarrow \mathbb{R}$ is such that for $((\ell, \nu), (\ell, \zeta)) \in \widehat{S}$ and $((b, c, a), R) \in \widehat{A}((\ell, \nu), (\ell, \zeta))$ we have $\widehat{\pi}((\ell, \nu), (\ell, \zeta), ((b, c, a), R)) = b - \nu(c)$.

Although the boundary region graph is infinite, for a fixed initial state we can restrict attention to a finite state subgraph, thanks to the following observation [20].

Lemma 8. For any state of a boundary region graph, its reachable sub-graph is finite.

5 Solving PTA Games on the Boundary Region Graph.

We now show that the boundary region graph abstraction for PTAs is sufficient to solve the expected reachability-time games. The partition of the locations of a probabilistic timed game arena $\mathcal{T} = (\mathsf{T}, L_{\text{Min}}, L_{\text{Max}})$ gives rise to the partition $(\widehat{S}_{\text{Min}}, \widehat{S}_{\text{Max}})$ of the set of states \widehat{S} of its boundary region graph and let $\widehat{\mathcal{T}} = (\widehat{\mathsf{T}}, \widehat{S}_{\text{Min}}, \widehat{S}_{\text{Max}})$.

¹ Notice that $\widehat{A}(R) = \widehat{A}(s)$ for all $s = ((\ell, \nu), R) \in \widehat{S}$.

We begin by reviewing the optimality equations for an expected reachability-time game on a boundary region graph \widehat{T} . Let $P : \widehat{S} \rightarrow \mathbb{R}_\oplus$. We say that P is a solution of optimality equations $\text{Opt}(\widehat{T})$, and we write $P \models \text{Opt}(\widehat{T})$, if for any $s \in \widehat{S}$:

$$P(s) = \begin{cases} 0 & \text{if } s \in \widehat{F} \\ \min_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + \sum_{s' \in S} p(s'|s, \alpha) \cdot P(s')\} & \text{if } s \in \widehat{S}_{\text{Min}} \setminus \widehat{F} \\ \max_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + \sum_{s' \in S} p(s'|s, \alpha) \cdot P(s')\} & \text{if } s \in \widehat{S}_{\text{Max}} \setminus \widehat{F}. \end{cases}$$

Before trying to solve $\text{Opt}(\widehat{T})$ for a probabilistic timed game, let us consider the simpler case when T is a timed game.

The non-probabilistic case. For a timed game T we define $\text{Succ} : \widehat{S} \times \widehat{A} \rightarrow \widehat{S}$ as follows:

$$\text{Succ}(((\ell, \nu), R), ((b, c, a), (\ell, \zeta))) = ((\ell', (\nu + b - \nu(c))[C:=0]), (\ell', \zeta[C:=0])),$$

where $(C, \ell') \in 2^C \times L$ is such that $\delta(\ell, a)(C, \ell') = 1$. Now, using this function, the optimality equations $\text{Opt}(\widehat{T})$ can be rewritten as:

$$P(s) = \begin{cases} 0 & \text{if } s \in \widehat{F} \\ \min_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + P(\text{Succ}(s, \alpha))\} & \text{if } s \in \widehat{S}_{\text{Min}} \setminus \widehat{F} \\ \max_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + P(\text{Succ}(s, \alpha))\} & \text{if } s \in \widehat{S}_{\text{Max}} \setminus \widehat{F}. \end{cases}$$

Based on these equations [3] introduced the following value iteration algorithm.

Algorithm 9. Value iteration algorithm for (non-probabilistic) $\text{Opt}(\widehat{T})$.

1. Set $i := 0$, $p_0(s) := 0$ if $s \in \widehat{F}$ and $p_0(s) := \infty$ otherwise.
2. Set $p_{i+1} := \Phi(p_i)$.
3. If $p_{i+1} = p_i$ then return p_i , else set $i := i+1$ and goto step 2.

where $\Phi : [\widehat{S} \rightarrow \mathbb{R}_\oplus] \rightarrow [\widehat{S} \rightarrow \mathbb{R}_\oplus]$ is such that for any $f : \widehat{S} \rightarrow \mathbb{R}_\oplus$ and $s \in \widehat{S}$:

$$\Phi(f)(s) = \begin{cases} 0 & \text{if } s \in \widehat{F} \\ \min_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + f(\text{Succ}(s, \alpha))\} & \text{if } s \in \widehat{S}_{\text{Min}} \setminus \widehat{F} \\ \max_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + f(\text{Succ}(s, \alpha))\} & \text{if } s \in \widehat{S}_{\text{Max}} \setminus \widehat{F}. \end{cases} \quad (1)$$

The proof of correctness of this algorithm is reliant on the concept of *simple functions*, and certain closure properties of these functions, which we now review.

Definition 10 (Simple Functions). Let $X \subseteq V$. A function $F : X \rightarrow \mathbb{R}$ is simple if either: there is $e \in \mathbb{Z}$, such that $F(\nu) = e$ for every $\nu \in X$; or there are $e \in \mathbb{Z}$ and $c \in C$, such that $F(\nu) = e - \nu(c)$ for all $\nu \in X$.

We say a function $F : \widehat{S} \rightarrow \mathbb{R}_\oplus$ is regionally simple if for every region $(\ell, \zeta) \in \mathcal{R}$ the function $F((\ell, \cdot), (\ell, \zeta))$ is simple. Asarin and Maler [3] showed that the following properties hold for simple functions.

Proposition 11 (Properties of simple functions).

1. If $F : X \rightarrow \mathbb{R}$ is simple, then $\overline{F} : \overline{X} \rightarrow \mathbb{R}$ is simple.
2. If $F, F' : \widehat{S} \rightarrow \mathbb{R}$ are regionally simple functions, then $\min(F, F')$ and $\max(F, F')$ are also regionally simple.
3. If F be regionally simple, then, for every region $R = (\ell, \zeta)$ and $\alpha \in \widehat{A}(R)$, the function $t((\ell, \cdot), R), \alpha) + F(\text{Succ}((\ell, \cdot), R), \alpha)$ is simple.
4. Any decreasing sequence of regionally simple functions is finite.

Using the first three closure properties of simple functions, it is easy to see that the function Φ in (1) is such that, if f is regionally simple, then so is $\Phi(f)$. Since the initial function p_0 is regionally simple, it is immediate that, if Algorithm 9 terminates, it will return a regionally simple solution of $\text{Opt}(\widehat{\mathcal{T}})$. Now, since the function Φ is monotonic, the sequence $\langle p_0, p_1, p_2, \dots \rangle$ of intermediate value functions in Algorithm 9 is a decreasing sequence of regionally simple functions. Proposition 11(4) then guarantees the termination of the value iteration algorithm. Jurdziński and Trivedi [18] show that if a solution of $\text{Opt}(\widehat{\mathcal{T}})$ is regionally simple then it gives a solution of optimality equations for the original timed automaton.

The probabilistic case. In this section we consider extending the above approach to solve $\text{Opt}(\widehat{\mathcal{T}})$ when \mathcal{T} is a probabilistic timed automaton. Based on the optimality equations, we define the value improvement function $\Psi : [\widehat{S} \rightarrow \mathbb{R}_\oplus] \rightarrow [\widehat{S} \rightarrow \mathbb{R}_\oplus]$ such that for any $f : \widehat{S} \rightarrow \mathbb{R}_\oplus$ and $s \in \widehat{S}$:

$$\Psi(f)(s) = \begin{cases} 0 & \text{if } s \in \widehat{F} \\ \min_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + \sum_{s' \in S} p(s'|s, \alpha) \cdot f(s')\} & \text{if } s \in \widehat{S}_{\text{Min}} \setminus \widehat{F} \\ \max_{\alpha \in \widehat{A}(s)} \{t(s, \alpha) + \sum_{s' \in S} p(s'|s, \alpha) \cdot f(s')\} & \text{if } s \in \widehat{S}_{\text{Max}} \setminus \widehat{F}. \end{cases} \quad (2)$$

It is straightforward to verify that a fixpoint of the function Ψ is the solution of $\text{Opt}(\widehat{\mathcal{T}})$. By Assumption 3 and Lemma 8, it is immediate that Ψ is a contraction, and therefore Ψ can be used in a straightforward value iteration algorithm to approximate $\text{Opt}(\widehat{\mathcal{T}})$. However, trying to extend the approach of Asarin and Maler [3] fails since the intermediate functions in the value iteration algorithms no longer remain regionally simple. To overcome this problem, we present a generalisation of simple functions, which we call *quasi-simple* functions.

Before introducing quasi-simple functions, we require the partial order $\trianglelefteq \subseteq V \times V$, where for any valuations ν and ν' we have $\nu \trianglelefteq \nu'$ if and only if there exists a $t \in \mathbb{R}_\oplus$ such that for each clock $c \in C$ either $\nu'(c) - \nu(c) = t$ or $\nu(c) = \nu'(c)$, and $\nu'(c) - \nu(c) = t$ for at least one clock $c \in C$. In this case we also write $(\nu' - \nu) = t$.

Definition 12 (Quasi-Simple Functions). Let $X \subseteq V$ be a subset of valuations. A function $F : X \rightarrow \mathbb{R}$ is quasi-simple if:

- (Lipschitz Continuous) there exists $K \geq 0$ such that $|F(\nu) - F(\nu')| \leq K \cdot \|\nu - \nu'\|_\infty$ for all $\nu, \nu' \in X$;
- (Monotonically decreasing and nonexpansive w.r.t. \trianglelefteq) $\nu \trianglelefteq \nu'$ implies $F(\nu) \geq F(\nu')$ and $F(\nu) - F(\nu') \leq \nu' - \nu$ for all $\nu, \nu' \in X$.

Proposition 13 (Quasi-simple functions generalise simple functions). *Every simple function is also quasi-simple.*

Proof. Let $X \subseteq V$ be a subset of valuations and $F : X \rightarrow \mathbb{R}$ a simple function. If F is constant then the proposition trivially follows. Otherwise, there exists $b \in \mathbb{Z}$ and $c \in C$ such that $F(\nu) = b - \nu(c)$ for all $\nu \in X$. We need to show that F is Lipschitz continuous, and monotonically decreasing and nonexpansive w.r.t \preceq .

1. To prove that F is Lipschitz continuous, notice that $|F(\nu) - F(\nu')| = |b - \nu(c) - b + \nu'(c)| = |\nu'(c) - \nu(c)| \leq \|\nu - \nu'\|_\infty$.
2. For $\nu, \nu' \in X$ such that $\nu \preceq \nu'$, we have $F(\nu) = b - \nu(c) \geq b - \nu'(c) = F(\nu')$.
From the first part of this proof, it trivially follows that $F(\nu) - F(\nu') \leq \nu - \nu'$. \square

We say a function $F : \widehat{S} \rightarrow \mathbb{R}_\oplus$ is regionally quasi-simple if for every region $(\ell, \zeta) \in \mathcal{R}$ the function $F((\ell, \cdot), (\ell, \zeta))$ is quasi-simple.

Lemma 14 (Properties of Quasi-Simple Functions).

1. If $F : X \rightarrow \mathbb{R}$ is quasi-simple, then $\overline{F} : \overline{X} \rightarrow \mathbb{R}$ is quasi-simple.
2. If $F, F' : \widehat{S} \rightarrow \mathbb{R}$ are regionally quasi-simple functions, then $\max(F, F')$ and $\min(F, F')$ are also regionally quasi-simple.
3. If F is regionally quasi-simple, then, for any $R = (\ell, \zeta)$ and $\alpha \in \widehat{A}(R)$, the function $t((\ell, \cdot), R), \alpha) + \sum_{s' \in S} p(s' | ((\ell, \cdot), R), \alpha) \cdot F(s')$ is quasi-simple.
4. The limit of a sequence of quasi-simple functions is quasi-simple.

From the first three properties, it follows that the intermediate functions of Ψ in (2) are regionally quasi-simple. In addition, the fourth property implies that its fixpoint is also regionally quasi-simple, and hence the solution of $\text{Opt}(\widehat{T})$ is regionally quasi-simple.

Proposition 15. *Let \mathcal{T} be a probabilistic timed game. If $P \models \text{Opt}(\widehat{T})$, then P is regionally quasi-simple.*

6 Correctness of the Reduction

We now demonstrate the correctness of our results, showing that the problem of expected reachability-time games on PTAs can be reduced to expected reachability-price games over the boundary region graph. For a given function $f : \widehat{S} \rightarrow \mathbb{R}$, we define $\widetilde{f} : S \rightarrow \mathbb{R}$ by $\widetilde{f}(\ell, \nu) = f((\ell, \nu), (\ell, [\nu]))$. Formally we have the following result.

Theorem 16. *Let \mathcal{T} be a probabilistic timed game. If $P \models \text{Opt}(\widehat{T})$, then $\widetilde{P} \models \text{Opt}(\mathcal{T})$.*

Before giving the proof we require the following property of quasi-simple functions.

Lemma 17. *Let $s = (\ell, \nu) \in S$ and $(\ell, \zeta) \in \mathcal{R}$ such that $(\ell, [\nu]) \rightarrow_* (\ell, \zeta)$. If $F : \widehat{S} \rightarrow \mathbb{R}$ is regionally quasi-simple, then the function $F_{s, \zeta, a}^\oplus : I \rightarrow \mathbb{R}$ defined as*

$$F_{s, \zeta, a}^\oplus(t) = t + \sum_{(C, \ell') \in 2^C \times L} \delta[\ell, a](C, \ell') \cdot F((\ell', \nu_C^t), (\ell', \zeta^C))$$

is continuous and nondecreasing, where $I = \{t \in \mathbb{R}_\oplus \mid \nu + t \in \zeta\}$, $\nu_C^t = \nu + t[C:=0]$ and $\zeta^C = \zeta[C:=0]$.

Proof (of Theorem 16). Suppose that $P \models \text{Opt}(\widehat{T})$, to prove this theorem it is sufficient to show that for any $s=(\ell, \nu) \in S_{\text{Min}}$ we have:

$$\tilde{P}(s) = \inf_{(t,a) \in A(s)} \left\{ t + \sum_{(C,\ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot \tilde{P}(\ell', (\nu+t)[C:=0]) \right\} \quad (3)$$

and for any $s=(\ell, \nu) \in S_{\text{Max}}$ we have:

$$\tilde{P}(s) = \sup_{(t,a) \in A(s)} \left\{ t + \sum_{(C,\ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot \tilde{P}(\ell', (\nu+t)[C:=0]) \right\}. \quad (4)$$

In the remainder of the proof we restrict attention to Min states as the case for Max states follows similarly. Therefore we fix $s=(\ell, \nu) \in S_{\text{Min}}$ for the remainder of the proof. For $a \in \text{Act}$, let $\mathcal{R}_{\text{Thin}}^a$ and $\mathcal{R}_{\text{Thick}}^a$ denote the set of thin and thick regions respectively that are successors of $[\nu]$ and are subsets of $E(\ell, a)$. Considering the RHS of (3) we have:

$$\text{RHS of (3)} = \min_{a \in \text{Act}} \{T_{\text{Thin}}(s, a), T_{\text{Thick}}(s, a)\}, \quad (5)$$

where $T_{\text{Thin}}(s, a)$ ($T_{\text{Thick}}(s, a)$) is the infimum (supremum) of the RHS of (3) over all actions (t, a) such that $[\nu+t] \in \mathcal{R}_{\text{Thin}}^a$ ($[\nu+t] \in \mathcal{R}_{\text{Thick}}^a$). For the first term we have:

$$\begin{aligned} T_{\text{Thin}}(s, a) &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thin}}^a} \inf_{\substack{t \in \mathbb{R}^+ \\ \nu+t \in \zeta}} \left\{ t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot \tilde{P}(\ell', \nu_C^t) \right\} \\ &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thin}}^a} \inf_{\substack{t \in \mathbb{R}^+ \\ \nu+t \in \zeta}} \left\{ t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^t), (\ell', \zeta^C)) \right\} \\ &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thin}}^a} \left\{ t^{(\ell, \zeta)} + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^{t^{(\ell, \zeta)}}), (\ell', \zeta^C)) \right\} \end{aligned}$$

where ν_C^t denotes the clock valuation $(\nu+t)[C:=0]$, $t^{(\ell, \zeta)}$ the time to reach the region R from s and ζ^C the region $\zeta[C:=0]$. Considering the second term of (5) we have

$$\begin{aligned} T_{\text{Thick}}(s, a) &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thick}}^a} \inf_{\substack{t \in \mathbb{R}^+ \\ \nu+t \in \zeta}} \left\{ t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot \tilde{P}(\ell', \nu_C^t) \right\} \\ &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thick}}^a} \inf_{\substack{t \in \mathbb{R}^+ \\ \nu+t \in \zeta}} \left\{ t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^t), (\ell', \zeta^C)) \right\} \\ &= \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thick}}^a} \inf_{\substack{t_{R_-}^s < t < t_{R_+}^s \\ R \xrightarrow{+1} R_- \\ R \xrightarrow{+1} R_+}} \left\{ t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^t), (\ell', \zeta^C)) \right\} \end{aligned}$$

From Proposition 15 it follows that P is regionally quasi-simple and, from Lemma 17, the function:

$$t + \sum_{(C, \ell') \in 2^c \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^t), (\ell', \zeta^C))$$

is continuous and nondecreasing over $\{t \mid \nu + t \in \zeta\}$. Therefore it follows that

$$T_{\text{Thick}}(s, a) = \min_{(\ell, \zeta) \in \mathcal{R}_{\text{Thick}}^a} \min_{\substack{t = t_{R_-}^s, t_{R_+}^s \\ (\ell, \zeta) \xleftarrow{+1} R_- \\ (\ell, \zeta) \xrightarrow{+1} R_+}} \left\{ t + \sum_{(C, \ell') \in 2^C \times L} \delta[\ell, a](C, \ell') \cdot P((\ell', \nu_C^t), (\ell', \zeta^C)) \right\}$$

Substituting the values of $T_{\text{Thin}}(s, a)$ and $T_{\text{Thick}}(s, a)$ into (5) and observing that, for any thin region $(\ell, \zeta) \in \mathcal{R}_{\text{Thin}}^a$, there exist $b \in \mathbb{Z}$ and $c \in C$ such that $\nu + (b - \nu(c)) \in \zeta$, it follows from Definition 7 that RHS of (3) equals:

$$\min_{\alpha \in \widehat{A}(s, [s])} \left\{ t((s, [s]), \alpha) + \sum_{(s', R') \in \widehat{S}} \widehat{p}((s', R') \mid (s, [s]), \alpha) \cdot P(s', R') \right\}$$

which by definition equals $\widetilde{P}(s)$ as required. \square

7 Expected Discounted-Time Games

Let $\mathcal{T} = (\mathbb{T}, L_{\text{Min}}, L_{\text{Max}})$ be a probabilistic timed game arena and $\lambda \in [0, 1)$ be a discount factor. In an expected discounted-time game starting from state s and for a strategy pair μ, χ , player Min loses the following amount to player Max:

$$\text{EDisct}(s, \mu, \chi) \stackrel{\text{def}}{=} \mathbb{E}_s^{\mu, \chi} \left\{ \sum_{i=1}^{\infty} \lambda^i \cdot \pi(X_{i-1}, Y_i) \right\} .$$

The concepts for the expected discounted-time game are defined in an analogous manner to that of expected reachability-time games. A reduction from expected discounted-time game to expected reachability-time game is standard [25]. Therefore, using the techniques presented in this paper one can reduce the problem of solving expected discounted-time games on \mathcal{T} to solving corresponding problem on $\widetilde{\mathcal{T}}$. Similarly a (non-probabilistic) discounted-time game on timed automata can be reduced to solving discounted-price games on (non-probabilistic) boundary region graph.

Proposition 18 (Discounted-Time Games).

1. *Expected discounted-time games on probabilistic timed automata can be reduced to expected discounted-price games on the corresponding boundary region graph.*
2. *Discounted-time games on timed automata can be reduced to discounted-price games on the corresponding (non-probabilistic) boundary region graph.*

8 Complexity

Theorem 19. *Discounted-time games are EXPTIME-complete.*

Proof. Using the ideas from [27] the problem of solving average-time games on timed automata can be reduced, in polynomial time, to solving discounted-time games on

timed automata with an equal number of clocks. Since the problem of solving average-time games is EXPTIME-complete [19] on timed automata with at least two clocks; discounted-time games are EXPTIME-hard on timed automata with at least two clocks.

From Proposition 18 it follows that, to solve a discounted-time game on a timed automaton, it is sufficient to solve the discounted-price game on the finite (non-probabilistic) boundary region graph of the automaton. Observe that every region, and hence also every configuration of the game, can be written down in polynomial space, and that every move of the game can be simulated in polynomial time. Therefore, the value of the game can be computed by a straightforward alternating PSPACE algorithm, and hence the problem is in EXPTIME since $\text{APSPACE} = \text{EXPTIME}$. \square

Theorem 20. *The expected reachability-time games and the expected discounted-time games are EXPTIME-hard and they are in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$.*

Proof. The EXPTIME-hardness of expected reachability-time games and expected discounted-time games on probabilistic timed automata with two or more clocks follows from the fact that discounted-time games are EXPTIME-complete. The membership in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$ follows from the reduction to the boundary region graph, and the observations that: size of the boundary region graph is exponential in the size of the PTA; and the complexity of solving expected reachability-price games and expected discounted-price games on finite MDP is in $\text{NP} \cap \text{co-NP}$. \square

9 Conclusion

In this paper we have employed the boundary region graph to solve quantitative games over probabilistic timed automata. The approach is based on extending the class of simple functions introduced in [3] to quasi-simple functions. Our results demonstrate that the problem of solving games with either expected reachability-time or expected discounted-time criteria on PTA are in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$. Future work includes finding practical symbolic zone-based algorithms to solve quantitative games on timed automata and, perhaps more ambitiously, games on PTA. Regarding the other quantitative games on PTA, we conjecture that it is possible to reduce expected average-time games on PTA to mean payoff games on boundary region graph. However, the techniques presented in this paper are insufficient to demonstrate such a reduction.

Acknowledgements. The authors are supported in part by EPSRC grants EP/D076625 and EP/F001096. In addition, the third author wishes to thank Vojtěch Forejt for some stimulating discussions on this class of games.

References

1. R. Alur, M. Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In *Proc. ICALP'04*, volume 3142 of *LNCS*. Springer, 2004.
2. R. Alur and D. Dill. A theory of timed automata. In *Theoretical Computer Science*, volume 126, 1994.

3. E. Asarin and O. Maler. As soon as possible: Time optimal control for timed automata. In *Proc. HSCC'99*, volume 1569 of *LNCS*. Springer, 1999.
4. E. Asarin, O. Maler, and A. Pnueli. Symbolic controller synthesis for discrete and timed systems. In *Hybrid Systems II*, volume 999 of *LNCS*. Springer, 1995.
5. D. Beauquier. Probabilistic timed automata. *Theoretical Computer Science*, 292(1), 2003.
6. G. Behrmann, A. Cougnard, A. David, E. Fleury, K. Larsen, and D. Lime. UPPAAL-Tiga: Time for playing games! In *Proc. CAV'07*, volume 4590 of *LNCS*. Springer, 2007.
7. P. Bouyer, T. Brihaye, and N. Markey. Improved undecidability results on weighted timed automata. In *Information Processing Letters*, volume 98. Elsevier, 2006.
8. P. Bouyer, F. Cassez, E. Fleury, and K. G. Larsen. Optimal strategies in priced timed game automata. In *FSTTCS'04*, volume 3328 of *LNCS*. Springer, 2004.
9. P. Bouyer and V. Forejt. Reachability in stochastic timed games. In *Proc. ICALP'09*, volume 5556 of *LNCS*. Springer, 2009.
10. T. Brihaye, T. A. Henzinger, V. S. Prabhu, and J. Raskin. Minimum-time reachability in timed games. In *Proc. ICALP'07*, volume 4596 of *LNCS*. Springer, 2007.
11. Thomas Brihaye, Véronique Bruyère, and Jean-François Raskin. On optimal timed strategies. In *Proc. FORMATS'05*, volume 3829 of *LNCS*. Springer, 2005.
12. L. de Alfaro. Computing minimum and maximum reachability times in probabilistic systems. In *Proc. CONCUR'99*, volume 1664 of *LNCS*. Springer, 1999.
13. T. A. Henzinger, Z. Manna, and A. Pnueli. What good are digital clocks? In *Proc. ICALP'92*, volume 623 of *LNCS*. Springer, 1992.
14. G. Hoffmann and H Wong-Toi. The input-output control of real-time discrete event systems. In *In the Proceedings of 13th IEEE Real-Time Systems Symposium*, 1992.
15. H. Jensen. Model checking probabilistic real time systems. In *Proc. 7th Nordic Workshop on Programming Theory*, Report 86:247–261. Chalmers University of Technology, 1996.
16. M. Jurdziński, M. Kwiatkowska, G. Norman, and A. Trivedi. Concavely-priced probabilistic timed automata. In *Proc. CONCUR'09*, volume 5710 of *LNCS*, pages 415–430. Springer, 2009.
17. M. Jurdziński, J. Sproston, and F. Laroussinie. Model checking probabilistic timed automata with one or two clocks. *Logical Methods in Computer Science*, 4(3), 2008.
18. M. Jurdziński and A. Trivedi. Reachability-time games on timed automata. In *Proc. ICALP'07*, volume 4596 of *LNCS*. Springer, 2007.
19. M. Jurdziński and A. Trivedi. Average-time games. In *Proc. FSTTCS'08*, volume 2 of *Leibniz International Proceedings in Informatics*. Schloss Dagstuhl, 2008.
20. M. Jurdziński and A. Trivedi. Concavely-priced timed automata. In *Proc. FORMATS'08*, volume 5215 of *LNCS*. Springer, 2008.
21. M. Kwiatkowska, G. Norman, D. Parker, and J. Sproston. Performance analysis of probabilistic timed automata using digital clocks. *Formal Methods in System Design*, 29, 2006.
22. M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282, 2002.
23. A. Neyman and S. Sorin, editors. *Stochastic Games and Applications*, volume 570 of *NATO Science Series C*. Kluwer Academic Publishers, 2004.
24. P. J. Ramadge and W. M. Wonham. The control of discrete event systems. In *In Proceedings of the IEEE*, volume 77(1), 1989.
25. L. S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. U.S.A.*, 39, 1953.
26. S. Tripakis. Verifying progress in timed systems. In *Proc. ARTS'99*, volume 1601 of *LNCS*. Springer, 1999.
27. U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158, 1996.