

Expected Reachability-Time Games

Vojtěch Forejt¹, Marta Kwiatkowska², Gethin Norman³, and Ashutosh Trivedi²

¹ Faculty of Informatics, Masaryk University, Botanická 68a, Brno, Czech Republic

² Oxford University Computing Laboratory, Parks Road, Oxford, OX1 3QD, UK

³ Department of Computing Science, University of Glasgow, Glasgow, G12 8RZ, UK

Abstract. In an expected reachability-time game (ERTG) two players, Min and Max, move a token along the transitions of a probabilistic timed automaton, so as to minimise and maximise, respectively, the expected time to reach a target. These games are concurrent since at each step of the game both players choose a timed move (a time delay and action under their control), and the transition of the game is determined by the timed move of the player who proposes the shorter delay. A game is turn-based if at any step of the game, all available actions are under the control of precisely one player. We show that while concurrent ERTGs are not always determined, turn-based ERTGs are positionally determined. Using the boundary region graph abstraction, and a generalisation of Asarin and Maler’s simple function, we show that the decision problems related to computing the upper/lower values of concurrent ERTGs, and computing the value of turn-based ERTGs are decidable and their complexity is in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$.

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 [1]. 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 e.g. [2,3,4,5,6].

In this paper we extend this approach to objectives that are quantitative both in terms of timed *and* probabilistic behaviour. Probabilistic behaviour is important in 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 (PTAs) [7,8,9], a model for real-time systems exhibiting nondeterministic and probabilistic behaviour. We concentrate on *expected reachability-time games* (ERTGs), which are games on PTAs where the performance objective concerns the minimum expected time the controller can ensure for the system to reach a target, regardless of the uncontrollable (environmental) events that occur. This approach has many practical applications, e.g., in job-shop scheduling, where machines can be faulty or have variable execution time, and both routing and task graph scheduling problems. For real-life examples relevant to our setting, see e.g. [10,6].

In the games that we study, a token is placed on a configuration of a PTA and a play of the game corresponds to both players proposing a timed move of the PTA, i.e. a time delay and action under their control (we assume each action of the PTA is under the control of precisely one of the players). Once the players have made their choices, the timed move with the shorter delay⁴ is performed and the token is moved according to the probabilistic transition function of the PTA. Players Min and Max choose their moves in order to minimise and maximise, respectively, the payoff function (the time till the first visit of a target in the case of ERTGs). It is well known, see, e.g. [11], that concurrent timed games are not *determined*, which means the *upper value* of the game (the minimum expected time to reach a target that Min can ensure) is strictly greater than the *lower value* of the game (the maximum expected time to reach a target that Max can ensure). A game is determined if the lower and upper values are equal, and in this case, the *optimal value* of the game exists and equals the upper and lower values. We show that a subclass of ERTGs, called *turn-based* ERTGs, where at each step of the game only one of the players has available actions are *positionally determined*, i.e. both players have ε -optimal (optimal up to a given precision $\varepsilon > 0$) positional (history-independent and non-randomised) strategies.

The problem we consider is inspired by Asarin and Maler [2] who studied the *brachystochronic problem* for timed automata. This work focused on reachability-time games, i.e. games on a timed automata where the objective concerns the time to reach a target. The techniques of [2] exploit properties of a special class of functions called *simple functions*. The importance of simple functions is also observed in [12] in the context of one-player games. Simple functions have also enabled the computation of a uniform solution for (turn-based) reachability-time games [13] and the proof of correctness of game-reduction for turn-based average-time games [14]. However, we show that the concept of simple functions is not sufficient in the setting of PTAs.

Contribution. We show that the problem of deciding whether the upper (lower, or the optimal when it exists) value of an ERTG is at most a given bound is decidable. An important contribution of the paper is the generalisation of simple functions to *quasi-simple functions*. By using this class of functions and the boundary region abstraction [15,16], we give a novel proof of positional determinacy of *turn-based* ERTGs. We demonstrate that the problem of finding the upper and lower value of general ERTGs is in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$. An EXPTIME-hardness lower bound follows from the EXPTIME-completeness of the corresponding optimisation problem [16]. From [17] it follows that the problem is not NEXPTIME-hard, unless NP equals co-NP. Extending this work we get the similar results for expected discounted-time games.

Related Work. Hoffman and Wong-Toi [18] 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. [19]. Asarin and Maler [2] initiated the study of quantitative games on timed automata by providing a symbolic algorithm to solve reachability-time games. The work of [20] and [13] show that the decision version of the reachability-time game is EXPTIME-complete for timed automata with

⁴ Min and Max represent two different forms of non-determinism called *angelic* and *demonic*. To prevent the introduction of a third form, we assume the move of Max (the environment) is taken if the delays are equal. The converse can be used without changing the presented results.

at least two clocks. The tool UPPAAL Tiga [5] is capable of solving reachability and safety objectives for games on timed automata. Jurdziński and Trivedi [14] show the EXPTIME-completeness for average-time games on automata with two or more clocks.

A natural extension of reachability-time games are games on priced timed automata where the objective concerns the cumulated price of reaching a target. Both [3] and [4] present semi-algorithms for computing the value of such games for linear prices. In [21] the problem of checking the existence of optimal strategies is shown to be undecidable with [22] showing undecidability holds even for three clocks and stopwatch prices.

We are not aware of any previous work studying two-player quantitative games on PTAs. For a significantly different model of stochastic timed games, deciding whether a target is reachable within a given bound is undecidable [23]. Regarding one-player games on PTAs, [16] reduce a number of optimisation problems on concavely-priced PTAs to solving the corresponding problems on the boundary region abstraction and [24] solve expected reachability-price problems for linearly-priced PTAs using digital clocks. In [25] the problem of deciding whether a target can be reached within a given price and probability bound is shown to be undecidable for priced PTAs with three clocks and stopwatch prices. By a simple modification of the proofs in [25] it can be demonstrated that checking the existence of optimal strategies is undecidable for reachability-price turn-based games on priced (probabilistic) timed automata with three clocks and stopwatch prices.

A full version of this paper, including proofs is also available [26], while a preliminary version appeared as [27].

2 Expected Reachability Games

Expected reachability games (ERGs) are played between two players Min and Max on a state-transition graph, whose transitions are nondeterministic and probabilistic, by jointly resolving the nondeterminism to move a token along the transitions of the graph. The objective for player Min in the game is to reach the final states with the smallest accumulated reward, while Max tries to do the opposite.

Before we give a formal definition, we need to introduce the concept of discrete probability distributions. A *discrete distribution* over a (possibly uncountable) set Q is a function $d : Q \rightarrow [0, 1]$ such that $\text{supp}(d) = \{q \in Q \mid d(q) > 0\}$ is at most countable and $\sum_{q \in Q} d(q) = 1$. Let $\mathcal{D}(Q)$ denote the set of all discrete distributions over Q . We say a distribution $d \in \mathcal{D}(Q)$ is a *point distribution* if $d(q) = 1$ for some $q \in Q$.

Definition 1. An ERG is a tuple $G = (S, F, A_{\text{Min}}, A_{\text{Max}}, p_{\text{Min}}, p_{\text{Max}}, \pi_{\text{Min}}, \pi_{\text{Max}})$ where:

- S is a (possibly uncountable) set of states including a set of final states F ;
- A_{Min} and A_{Max} are (possibly uncountable) sets of actions controlled by players Min and Max and \perp is a distinguished action such that $A_{\text{Min}} \cap A_{\text{Max}} = \{\perp\}$;
- $p_{\text{Min}} : S \times A_{\text{Min}} \rightarrow \mathcal{D}(S)$ and $p_{\text{Max}} : S \times A_{\text{Max}} \rightarrow \mathcal{D}(S)$ are the partial probabilistic transition functions for players Min and Max such that $p_{\text{Min}}(s, \perp)$ and $p_{\text{Max}}(s, \perp)$ are undefined for all $s \in S$;
- $\pi_{\text{Min}} : S \times A_{\text{Min}} \rightarrow \mathbb{R}_{\geq 0}$ and $\pi_{\text{Max}} : S \times A_{\text{Max}} \rightarrow \mathbb{R}_{\geq 0}$ are the reward functions for players Min and Max.

We say that the ERG is *finite* if both S and A are finite. For any state s , we let $A_{\text{Min}}(s)$ denote the set of actions available to player Min in s , i.e., the actions $a \in A_{\text{Min}}$ for which $p_{\text{Min}}(s, a)$ is defined, letting $A_{\text{Min}}(s) = \perp$ if no such action exists. Similarly, $A_{\text{Max}}(s)$ denotes the actions available to player Max in s and we let $A(s) = A_{\text{Min}}(s) \times A_{\text{Max}}(s)$. We say that s is *controlled* by Min (Max) if $A_{\text{Max}}(s) = \{\perp\}$ ($A_{\text{Min}}(s) = \{\perp\}$) and the game G is *turn-based* if there is a partition $(S_{\text{Min}}, S_{\text{Max}})$ of S such that all states in S_{Min} (S_{Max}) are controlled by Min (Max).

A game G starts with a token in some initial state and players Min and Max construct an infinite play by repeatedly choosing enabled actions, and then moving the token to a successor state determined by their probabilistic transition functions where the reward of the move is determined by their reward functions. More precisely, if in state s players Min and Max choose actions a and b respectively, then if $\pi_{\text{Min}}(s, a) < \pi_{\text{Max}}(s, b)$ or $b = \perp$ the probabilistic transition function and reward value are determined by Min's choice, i.e. by the transition function $p_{\text{Min}}(s, a)$ and reward value $\pi_{\text{Min}}(s, a)$, and otherwise are determined by Max's choice. Formally we introduce the following auxiliary functions of an ERG which return the transition function and reward value of the game.

Definition 2. Let G be an ERG. The probabilistic transition and reward functions $p : S \times A_{\text{Min}} \times A_{\text{Max}} \rightarrow \mathcal{D}(S)$ and $\pi : S \times A_{\text{Min}} \times A_{\text{Max}} \rightarrow \mathbb{R}_{\geq 0}$ of G are such that for any $s \in S$ and $(a, b) \in A_{\text{Min}}$:

$$p(s, a, b) = \begin{cases} \text{undefined} & \text{if } a = b = \perp \\ p_{\text{Min}}(s, a) & \text{if } a \neq \perp \text{ and either } b = \perp \text{ or } \pi_{\text{Min}}(s, a) < \pi_{\text{Max}}(s, b) \\ p_{\text{Max}}(s, b) & \text{otherwise} \end{cases}$$

$$\pi(s, a, b) = \begin{cases} \pi_{\text{Min}}(s, a) & \text{if } b = \perp \text{ or } \pi_{\text{Min}}(s, a) < \pi_{\text{Max}}(s, b) \\ \pi_{\text{Max}}(s, b) & \text{otherwise.} \end{cases}$$

From the conditions imposed on the probabilistic transition function, it follows that $(a, b) \in A(s)$ if and only if $p(s, a, b)$ is defined. Using these definitions, if in state s the action pair $(a, b) \in A(s)$ is chosen, then the probability of making a transition to s' equals $p(s'|s, a, b) \stackrel{\text{def}}{=} p(s, a, b)(s')$ and the reward equals $\pi(s, a, b)$.

A transition of G is a tuple $(s, (a, b), s')$ such that $p(s'|s, a, b) > 0$ and a play is an finite or infinite sequence $\langle s_0, (a_1, b_1), s_1, \dots \rangle$ such that $(s_i, (a_{i+1}, b_{i+1}), s_{i+1})$ is a transition for all $i \geq 0$. For a finite play $\rho = \langle s_0, (a_1, b_1), s_1, \dots, s_k \rangle$, let $\text{last}(\rho)$ denote the last state s_k of the play. We write Play (Play_{fin}) for the sets of (finite) plays in G and $\text{Play}(s)$ ($\text{Play}_{\text{fin}}(s)$) for the sets of (finite) plays starting from $s \in S$.

A *strategy* of Min is a function $\mu : \text{Play}_{\text{fin}} \rightarrow \mathcal{D}(A_{\text{Min}})$ such that $\text{supp}(\mu(\rho)) \subseteq A_{\text{Min}}(\text{last}(\rho))$ for all finite plays $\rho \in \text{Play}_{\text{fin}}$, i.e. for any finite play, a strategy returns a distribution over actions available to Min in the last state of the play. A strategy χ of Max is defined analogously and we let Σ_{Min} and Σ_{Max} denote the sets of strategies of Min and Max, respectively. A strategy σ is *pure* if $\sigma(\rho)$ is a point distribution for all $\rho \in \text{Play}_{\text{fin}}$, while it is *stationary* if $\text{last}(\rho) = \text{last}(\rho')$ implies $\sigma(\rho) = \sigma(\rho')$ for all $\rho, \rho' \in \text{Play}_{\text{fin}}$. A strategy is *positional* if it is pure and stationary and let Π_{Min} and Π_{Max} denote the set of positional strategies of Min and Max, respectively.

For any state s and strategy pair $(\mu, \chi) \in \Sigma_{\text{Min}} \times \Sigma_{\text{Max}}$, let $\text{Play}^{\mu, \chi}(s)$ denote the infinite plays in which Min and Max play according to μ and χ , respectively. Using

standard results from probability theory, see e.g., [28], we can construct a probability measure $Prob_s^{\mu,\chi}$ over the set $Play^{\mu,\chi}(s)$. Let X_i and Y_i denote the random variables corresponding to i^{th} state and action of a play (i.e., for play $\langle s_0, (a_1, b_1), s_1, \dots \rangle$ we have $X_i = s_i$ and $Y_{i+1} = (a_{i+1}, b_{i+1})$), and given a *real-valued random variable* $f : Play \rightarrow \mathbb{R}$, let $\mathbb{E}_s^{\mu,\chi}\{f\}$ denote the expected value of f with respect to the probability measure $Prob_s^{\mu,\chi}$. To keep the presentation simple, for the rest of the paper we only consider *transient stochastic games* [30, Chapter 4] (games where every play is finite with probability 1) and for this reason we make the following assumption⁵.

Assumption 1. For any strategy pair $(\mu, \chi) \in \Sigma_{\text{Min}} \times \Sigma_{\text{Max}}$, and state $s \in S$ there is $q > 0$ and $n \in \mathbb{N}$ such that $Prob_s^{\mu,\chi}(X_n \in F) \geq q$.

Recall that the objective for Min is to reach a final state with the smallest accumulated reward, while for Max it is the opposite. Starting from s , if Min uses the strategy μ and Max χ , then the expected reward accumulated before reaching a final state is given by:

$$\text{EReach}^{\mu,\chi}(s) \stackrel{\text{def}}{=} \mathbb{E}_s^{\mu,\chi} \left\{ \sum_{i=0}^{\min\{k-1 \mid X_k \in F\}} \pi(X_i, Y_{i+1}) \right\}.$$

Observe when starting at state s , Max can choose actions such that the expected reward is *at least* a value arbitrarily close to $\sup_{\chi \in \Sigma_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}^{\mu,\chi}(s)$. This is called the *lower value* $\text{Val}_*(s)$ of the game when starting at state s . For $\chi \in \Sigma_{\text{Max}}$ let $\text{Val}_\chi(s) = \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}^{\mu,\chi}(s)$. We say χ is *optimal* (ε -*optimal*), if $\text{Val}_\chi(s) = \text{Val}_*(s)$ ($\text{Val}_\chi(s) \geq \text{Val}_*(s) - \varepsilon$) for all $s \in S$. Similarly, Min can make choices such that the expected reward is *at most* a value arbitrarily close to the *upper value* $\text{Val}^*(s) = \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \text{EReach}^{\mu,\chi}(s)$. In addition, for $\mu \in \Sigma_{\text{Min}}$, we can define $\text{Val}^\mu(s)$ and say when μ is *optimal* or ε -*optimal*.

A game G is *determined* if $\text{Val}_*(s) = \text{Val}^*(s)$ for all $s \in S$ and then we say that the value of the game exists and equals $\text{Val}(s) = \text{Val}_*(s) = \text{Val}^*(s)$. If G is determined, then each player has an ε -optimal strategy for all $\varepsilon > 0$. A game is *positionally determined* if

$$\text{Val}(s) = \inf_{\mu \in \Pi_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \text{EReach}^{\mu,\chi}(s) = \sup_{\chi \in \Pi_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \text{EReach}^{\mu,\chi}(s)$$

for all $s \in S$. It is straightforward to see that if a game is positionally determined, then both players have *positional* ε -optimal strategies for all $\varepsilon > 0$.

Optimality Equations. We complete this section by introducing optimality equations for ERGs. For a game G and function $P : S \rightarrow \mathbb{R}_{\geq 0}$, we say that P is a solution of the optimality equations $\text{Opt}^*(G)$, and write $P \models \text{Opt}^*(G)$, if for any $s \in S$:

$$P(s) = \begin{cases} 0 & \text{if } s \in F \\ \inf_{\alpha \in A_{\text{Min}}(s)} \left\{ \sup_{\beta \in A_{\text{Max}}(s)} \left\{ \pi(s, \alpha, \beta) + \sum_{s' \in S} p(s'|s, \alpha, \beta) \cdot P(s') \right\} \right\} & \text{if } s \notin F. \end{cases}$$

and P is a solution of the optimality equations $\text{Opt}_*(G)$, and write $P \models \text{Opt}_*(G)$, if for any $s \in S$:

$$P(s) = \begin{cases} 0 & \text{if } s \in F \\ \sup_{\beta \in A_{\text{Max}}(s)} \left\{ \inf_{\alpha \in A_{\text{Min}}(s)} \left\{ \pi(s, \alpha, \beta) + \sum_{s' \in S} p(s'|s, \alpha, \beta) \cdot P(s') \right\} \right\} & \text{if } s \notin F. \end{cases}$$

⁵ Techniques (see, e.g., *positive stochastic games* [30, Chapter 4]) for lifting such an assumption are orthogonal to the main idea presented in this paper.

The following result demonstrate the correspondence between these equations and the lower and upper values of the expected reachability game.

Proposition 3. *For any ERG G and bounded function $P : S \rightarrow \mathbb{R}_{\geq 0}$:*

- *if $P \models \text{Opt}^*(G)$, then $\text{Val}^*(s) = P(s)$ for all $s \in S$ and for any $\varepsilon > 0$ player Min has a positional strategy μ_ε such that $\text{Val}^{\mu_\varepsilon}(s) \leq P(s) + \varepsilon$ for all $s \in S$;*
- *if $P \models \text{Opt}_*(G)$, then $\text{Val}_*(s) = P(s)$ for all $s \in S$ and for any $\varepsilon > 0$ player Max has a positional strategy χ_ε such that $\text{Val}_{\chi_\varepsilon}(s) \geq P(s) - \varepsilon$ for all $s \in S$.*

If G is turn-based, then the equations $\text{Opt}^*(G)$ and $\text{Opt}_*(G)$ are the same and we write $\text{Opt}(G)$ for these equations. The following is a direct consequence of Proposition 3.

Proposition 4. *If G is a turn-based, $P : S \rightarrow \mathbb{R}_{\geq 0}$ is a bounded and $P \models \text{Opt}(G)$, then $\text{Val}(s) = P(s)$ for all $s \in S$ and for any $\varepsilon > 0$ both players have ε -optimal strategies.*

3 Expected Reachability-Time Games

Expected reachability-time games (ERTGs) are played on the infinite graph of a probabilistic timed automaton where Min and Max choose their moves so that the expected time to reach a final state is minimised or maximised, respectively. Before defining ERTGs, we introduce the concept of clocks, constraints, regions, and zones.

Clocks. Let \mathcal{C} be a finite set of *clocks*. A *clock valuation* on \mathcal{C} is a function $\nu : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ and we write V for the set of clock valuations. Abusing notation, we also treat a valuation ν as a point in $\mathbb{R}^{|\mathcal{C}|}$. If $\nu \in V$ and $t \in \mathbb{R}_{\geq 0}$ 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 valuation where $\nu[C:=0](c)$ equals 0 if $c \in C$ and $\nu(c)$ otherwise. For $X \subseteq V$, we write \overline{X} for the smallest closed set in V containing X . Although clocks are usually allowed to take arbitrary non-negative values, w.l.o.g [31] we assume that there is an upper bound K such that for every clock $c \in \mathcal{C}$ we have that $\nu(c) \leq K$.

Clock constraints. A *clock constraint* over \mathcal{C} is a conjunction of *simple constraints* of the form $c \bowtie i$ or $c - c' \bowtie i$, where $c, c' \in \mathcal{C}$, $i \in \mathbb{N}$, $i \leq K$, and $\bowtie \in \{<, >, =, \leq, \geq\}$. For $\nu \in V$, let $\text{SCC}(\nu)$ be the finite set of simple constraints which hold in ν .

Clock regions. A *clock region* is a maximal set $\zeta \subseteq V$ such that $\text{SCC}(\nu) = \text{SCC}(\nu')$ for all $\nu, \nu' \in \zeta$. We write \mathcal{R} for the finite set of clock regions. Every clock region is an equivalence class of the indistinguishability-by-clock-constraints relation, and vice versa. We write $[\nu]$ for the clock region of ν and, if $\zeta = [\nu]$, write $\zeta[C:=0]$ for $[\nu[C:=0]]$.

Clock zones. 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. A set of clock valuations is a clock zone if and only if it is definable by a clock constraint. Observe that, for every clock zone W , the set \overline{W} is also a clock zone.

We now introduce ERTGs which extend classical timed automata [32] with discrete distributions and a partition of the actions between two players Min and Max.

Definition 5 (ERTG Syntax). *A (concurrent) expected reachability-time game (ERTG) is a tuple $\mathcal{T} = (L, L_F, \mathcal{C}, \text{Inv}, \text{Act}, \text{Act}_{\text{Min}}, \text{Act}_{\text{Max}}, E, \delta)$ where*

- L is a finite set of locations including a set of final locations L_F ;
- \mathcal{C} is a finite set of clocks;
- $Inv : L \rightarrow \mathcal{Z}$ is an invariant condition;
- Act is a finite set of actions and $\{Act_{\text{Min}}, Act_{\text{Max}}\}$ is a partition of Act ;
- $E : L \times Act \rightarrow \mathcal{Z}$ is an action enabledness function;
- $\delta : L \times Act \rightarrow \mathcal{D}(2^{\mathcal{C}} \times L)$ is a probabilistic transition function.

When we consider an ERTG as an input of an algorithm, its size is understood as the sum of the sizes of encodings of L , \mathcal{C} , Inv , Act , E , and δ . As usual [33], we assume that probabilities are expressed as ratios of two natural numbers, each written in binary.

An ERTG is *turn-based* if for each location ℓ , only one player has enabled actions, i.e. $E(\ell, a) = \emptyset$ for all $a \in Act_{\text{Min}}$ or $a \in Act_{\text{Max}}$. In this case, we write L_{Min} and L_{Max} for the set of locations where players Min and Max, respectively, have an enabled action. A *one-player ERTG* is a turn-based ERTG where one of the player does not control any location, i.e., either $L_{\text{Min}} = \emptyset$ or $L_{\text{Max}} = \emptyset$. A (non-probabilistic) *reachability-timed game* is an ERTG such that $\delta(\ell, a)$ is a point distribution for all $\ell \in L$ and $a \in Act$.

A *configuration* of an ERTG is a pair (ℓ, ν) , where ℓ is a location and ν a clock valuation such that $\nu \in Inv(\ell)$. For any $t \in \mathbb{R}$, we let $(\ell, \nu) + t$ equal the configuration $(\ell, \nu + t)$. In a configuration (ℓ, ν) , a timed action (time-action pair) (t, a) is available if and only if the invariant condition $Inv(\ell)$ is continuously satisfied while t time units elapse, and a is enabled (i.e. the enabling condition $E(\ell, a)$ is satisfied) after t time units have elapsed. Furthermore, if the timed action (t, a) is performed, then the next configuration is determined by the probabilistic transition relation δ , i.e. with probability $\delta[\ell, a](C, \ell')$ the clocks in C are reset and we move to the location ℓ' .

An ERTG starts at some initial configuration and Min and Max construct an infinite play by repeatedly choosing available timed actions $(t_a, a) \in \mathbb{R}_{\geq 0} \times Act_{\text{Min}}$ and $(t_b, b) \in \mathbb{R}_{\geq 0} \times Act_{\text{Max}}$ proposing \perp if no timed action is available. The player responsible for the move is Min if the time delay of Min's choice is less than that of Max's choice or Max chooses \perp , and otherwise Max is responsible. We assume the players cannot simultaneously choose \perp . We now present the formal semantics which is an ERG with potentially infinite number of states and actions. It is straightforward to show the semantics of a turn-based ERTG is a turn-based ERG.

Definition 6 (ERTG Semantics). Let \mathcal{T} be an ERTG. The semantics of \mathcal{T} is given the ERG $\llbracket \mathcal{T} \rrbracket = (S, F, A_{\text{Min}}, A_{\text{Max}}, p_{\text{Min}}, p_{\text{Max}}, \pi_{\text{Min}}, \pi_{\text{Max}})$ where

- $S \subseteq L \times V$ is the (possibly uncountable) set of states such that $(\ell, \nu) \in S$ if and only if $\nu \in Inv(\ell)$ and $F = \{(\ell, \nu) \in S \mid \ell \in L_F\}$ is the set of final states;
- $A_{\text{Min}} = (\mathbb{R}_{\geq 0} \times Act_{\text{Min}}) \cup \{\perp\}$ and $A_{\text{Max}} = (\mathbb{R}_{\geq 0} \times Act_{\text{Max}}) \cup \{\perp\}$ are the sets of timed actions of players Min and Max;
- for $\star \in \{\text{Min}, \text{Max}\}$, $(\ell, \nu) \in S$ and $(t, a) \in A_{\star}$ the probabilistic transition function p_{\star} is defined when $\nu + t' \in Inv(\ell)$ for all $t' \leq t$ and $\nu + t \in E(\ell, a)$ and for any (ℓ, ν') :

$$p_{\star}((\ell, \nu), (t, a))((\ell', \nu')) = \sum_{C \subseteq \mathcal{C} \wedge (\nu + t)[C := 0] = \nu'} \delta[\ell, a](C, \ell');$$

- for $\star \in \{\text{Min}, \text{Max}\}$, $s \in S$ and $(t, a) \in A_{\text{Min}}$ the reward function π_{\star} is given by $\pi_{\star}(s, (t, a)) = t$.




Fig. 1. An expected reachability-time game.

The sum in the definitions of p_{Min} and p_{Max} is used to capture the fact that resetting different subsets of \mathcal{C} may result in the same clock valuation (e.g. if all clocks are initially zero, then we end up with the same valuation, no matter which clocks we reset). Also, notice that the reward function of the ERG corresponds to the elapsed time of each move. For any ERTG \mathcal{T} , to ensure Assumption 1 holds on the ERG $\llbracket \mathcal{T} \rrbracket$, we require only that the following weaker assumption holds on $\llbracket \mathcal{T} \rrbracket$.

Assumption 2. For any strategy pair $(\mu, \chi) \in \Sigma_{\text{Min}} \times \Sigma_{\text{Max}}$, and state $s \in S$ we have that $\lim_{n \rightarrow \infty} \text{Prob}_s^{\mu, \chi}(X_n \in F) = 1$.

Example 7. Consider the ERTG in Figure 1; we use solid and dashed lines to indicate actions controlled by Min and Max respectively. The shaded circle denotes the final location. Considering location ℓ_1 , the invariant condition is $0 < y \leq 2 \wedge x \leq 2$, actions a and c are enabled when $y > 1$ and, if a is taken, we move to ℓ_F , while if c is taken, with probability 0.2 we move to ℓ_0 and reset y , and with probability 0.8 move to ℓ_F .

Starting in the configuration⁶ $(\ell_0, (0, 0))$ and supposing Min's strategy is to choose $(1.1, b)$ (i.e., wait 1.1 time units before performing action b) in location ℓ_0 and then choose $(0.5, a)$ in location ℓ_1 , while Max's strategy in location ℓ_1 is to choose $(0.2, c)$. One possible play under this strategy pair is $\langle (\ell_0, (0, 0)), ((1.1, b), \perp), (\ell_1, (0, 1.1)), ((0.5, a), (0.2, c)), (\ell_0, (0.2, 0)), ((1.1, b), \perp), (\ell_F, (1.3, 1.1)) \rangle$ which has probability $0.5 \cdot 0.2 \cdot 0.5 = 0.05$ and time $1.1 + 0.2 + 1.1 = 2.4$ of reaching the final location. Using the optimality equations $\text{Opt}^*(G)$ and $\text{Opt}_*(G)$, we obtain upper and lower value in state $(\ell_0, (0, 0))$ of $\frac{10}{9}$ and 1, respectively. For details of the equations see [26].

Example 7 above demonstrates that in general expected reachability-time games are not determined. However, our results yield a novel proof of the following fundamental result for turn-based expected reachability-time games.

Theorem 8. *Turn-based ERTGs are positionally determined.*

Since the general ERTG are not determined, we study the following decision problem related to computing the upper-value of a configuration. All presented results also apply to the corresponding lower value problem, and the value problem, if the value exists.

Definition 9 (ERTG Decision Problem). *The decision problem for an ERTG \mathcal{T} , a state s of $\llbracket \mathcal{T} \rrbracket$, and a bound $T \in \mathbb{Q}$ is to decide whether $\text{Val}^*(s) \leq T$.*

We now present the second fundamental result of the paper.

⁶ We suppose the first (second) coordinate in a clock valuation correspond to the clock x (y).

Theorem 10. *The ERTG decision problem is in $NEXPTIME \cap co-NEXPTIME$.*

From [16] we know that the ERTG problem is EXPTIME-complete even for one player ERTGs with two or more clocks. Hence the ERTG problem for general (two-player, concurrent) ERTG is at least EXPTIME-hard. Moreover, from the results of [17] and [29] it follows that ERTG problem is not NEXPTIME-hard, unless $NP = co-NP$.

4 Proofs of Theorems 8 and 10

This section is dedicated to the correctness of Theorems 8 and 10. We begin by defining *boundary region abstraction* (BRA) (an instance of an ERG) of an ERTG. In Section 4.2 we show that the solution of the optimality equations for a BRA always exists and is unique. While Section 4.3 demonstrates (Theorem 15) that the solution of the optimality equations of the BRA can be used to construct a solution of the optimality equations of the ERTG. Using these results we can then prove our main results.

Proof outline of Theorem 8. Using Theorem 15, a bounded solution of the equations for the upper and lower values of a ERTG always exists, and hence Proposition 3 implies both players have positional ε -optimal strategies. Since for turn-based ERTGs both equations are equivalent, from Proposition 4 positional determinacy of turn-based ERTGs follows.

Proof outline of Theorem 10. From Theorem 15 the upper value of a state of a ERTG can be derived from that of the boundary region abstraction. Since in the BRA the sub-graph of reachable states from any state is finite (Lemma 12) and its size is at most exponential in size of its ERTG, the upper value of a state in BRA can be computed by analysing an ERG of exponential size. The membership of the ERTG problem in $NEXPTIME \cap co-NEXPTIME$ then follows from the fact that a non-deterministic Turing machine needs to guess a (rational) solution of optimality equations only for exponentially many states, and it can verify in exponential time whether it is indeed a solution.

4.1 Boundary region abstraction

The region graph [32] is useful for solving time-abstract optimisation problems on timed automata. The region graph, however, is not suitable for solving timed optimisation problems and games on timed automata as it abstracts away the timing information. The corner-point abstraction [34] is an abstraction of timed automata which retains some timing information, but it is not convenient for the dynamic programming based proof techniques used in this paper. The boundary region abstraction (BRA) [13], a generalisation of the corner-point abstraction, is more suitable for such proof techniques. More precisely, we need to prove certain properties of values in ERTG, which we can do only when reasoning about all states of the ERTG. In the corner point abstraction we cannot do this since it represents only states corresponding to corner points of regions. Here, we generalise the BRA of [13] to handle ERTG.

Timed Successor Regions. Recall that \mathcal{R} is the set of clock regions. For $\zeta, \zeta' \in \mathcal{R}$, we say that ζ' is in the future of ζ , denoted $\zeta \xrightarrow{*} \zeta'$, if there exist $\nu \in \zeta, \nu' \in \zeta'$ and

$t \in \mathbb{R}_{\geq 0}$ such that $\nu' = \nu + t$ and say ζ' is the *time successor* of ζ if $\nu + t' \in \zeta \cup \zeta'$ for all $t' \leq t$ and write $\zeta \rightarrow \zeta'$, or equivalently $\zeta' \leftarrow \zeta$, to denote this fact. For regions $\zeta, \zeta' \in \mathcal{R}$ such that $\zeta \xrightarrow{*} \zeta'$ we write $[\zeta, \zeta']$ for the zone $\cup\{\zeta'' \mid \zeta \xrightarrow{*} \zeta'' \wedge \zeta'' \xrightarrow{*} \zeta'\}$.

Thin and Thick Regions. We say that a region ζ is *thin* if $[\nu] \neq [\nu + \varepsilon]$ for every $\nu \in \zeta$ and $\varepsilon > 0$ and *thick* otherwise. We write $\mathcal{R}_{\text{Thin}}$ and $\mathcal{R}_{\text{Thick}}$ for the sets of thin and thick regions, respectively. Observe that if $\zeta \in \mathcal{R}_{\text{Thick}}$ then, for any $\nu \in \zeta$, there exists $\varepsilon > 0$, such that $[\nu] = [\nu + \varepsilon]$ and the time successor of a thin region is thick, and vice versa.

Intuition for the Boundary Region Graph. Recall K is an upper bound on clock values and let $\llbracket K \rrbracket_{\mathbb{N}} = \{0, 1, \dots, K\}$. For any $\nu \in V$, $b \in \llbracket K \rrbracket_{\mathbb{N}}$ and $c \in \mathcal{C}$ we define $\text{time}(\nu, (b, c)) \stackrel{\text{def}}{=} b - \nu(c)$ if $\nu(c) \leq b$, and $\text{time}(\nu, (b, c)) \stackrel{\text{def}}{=} 0$ if $\nu(c) > b$. 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 . Observe that, for any $\zeta' \in \mathcal{R}_{\text{Thin}}$, there exists $b \in \llbracket K \rrbracket_{\mathbb{N}}$ and $c \in \mathcal{C}$, such that $\nu \in \zeta$ implies $(\nu + (b - \nu(c))) \in \zeta'$ for all $\zeta \in \mathcal{R}$ in the past of ζ' and write $\zeta \rightarrow_{b,c} \zeta'$. The boundary region abstraction is motivated by the following. Consider $a \in \text{Act}$, (ℓ, ν) and $\zeta \xrightarrow{*} \zeta'$ such that $\nu \in \zeta$, $[\zeta, \zeta'] \subseteq \text{Inv}(\ell)$ and $\nu' \in E(\ell, a)$.

- If $\zeta' \in \mathcal{R}_{\text{Thick}}$, then there are infinitely many $t \in \mathbb{R}_{\geq 0}$ such that $\nu + t \in \zeta'$. 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 ζ' is a thick region, the set $\{t \in \mathbb{R}_{\geq 0} \mid \nu + t \in \zeta'\}$ is an open interval, and hence does not contain its boundary values. Observe that the infimum equals $b_{\text{inf}} - \nu(c_{\text{inf}})$ where $\zeta \rightarrow_{b_{\text{inf}}, c_{\text{inf}}} \zeta_{\text{inf}} \rightarrow \zeta'$ and the supremum equals $b_{\text{sup}} - \nu(c_{\text{sup}})$ where $\zeta \rightarrow_{b_{\text{sup}}, c_{\text{sup}}} \zeta_{\text{sup}} \leftarrow \zeta'$. In the boundary region abstraction we include these ‘best’ timed actions through the actions $(b_{\text{inf}}, c_{\text{inf}}, a, \zeta')$ and $(b_{\text{sup}}, c_{\text{sup}}, a, \zeta')$.
- If $\zeta' \in \mathcal{R}_{\text{Thin}}$, then there exists a unique $t \in \mathbb{R}_{\geq 0}$ such that $\nu + t \in \zeta'$. Moreover since ζ' is a thin region, there exists a clock $c \in \mathcal{C}$ and a number $b \in \mathbb{N}$ such that $\zeta \rightarrow_{b,c} \zeta'$ and $t = b - \nu(c)$. In the boundary region abstraction we summarise this ‘best’ timed action from region ζ via region ζ' through the action (b, c, a, ζ') .

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

Definition 11. For an ERTG $\mathcal{T} = (L, L_F, \mathcal{C}, \text{Inv}, \text{Act}, \text{Act}_{\text{Min}}, \text{Act}_{\text{Max}}, E, \delta)$ the BRA of \mathcal{T} is given by the ERG $\widehat{\mathcal{T}} = (\widehat{S}, \widehat{F}, \widehat{A}_{\text{Min}}, \widehat{A}_{\text{Max}}, \widehat{p}_{\text{Min}}, \widehat{p}_{\text{Max}}, \widehat{\pi}_{\text{Min}}, \widehat{\pi}_{\text{Max}})$ where

- $\widehat{S} \subseteq L \times V \times \mathcal{R}$ is the (possibly uncountable) set of states such that $(\ell, \nu, \zeta) \in \widehat{S}$ if and only if $\zeta \in \mathcal{R}$, $\zeta \subseteq \text{Inv}(\ell)$, and $\nu \in \zeta$;
- $\widehat{F} = \{(\ell, \nu, \zeta) \in \widehat{S} \mid \ell \in L_F\}$ is the set of final states;
- $\widehat{A}_{\text{Min}} \subseteq (\llbracket K \rrbracket_{\mathbb{N}} \times \mathcal{C} \times \text{Act}_{\text{Min}} \times \mathcal{R}) \cup \{\perp\}$ is the set of actions of player Min;
- $\widehat{A}_{\text{Max}} \subseteq (\llbracket K \rrbracket_{\mathbb{N}} \times \mathcal{C} \times \text{Act}_{\text{Max}} \times \mathcal{R}) \cup \{\perp\}$ is the set of actions of player Max;
- for $\star \in \{\text{Min}, \text{Max}\}$, $s = (\ell, \nu, \zeta) \in \widehat{S}$ and $\alpha = (b_\alpha, c_\alpha, a_\alpha, \zeta_\alpha) \in \widehat{A}_\star$ the probabilistic transition function p_\star is defined if $[\zeta, \zeta_\alpha] \subseteq \text{Inv}(\ell)$ and $\zeta_\alpha \subseteq E(\ell, a_\alpha)$ and for any $(\ell', \nu', \zeta') \in \widehat{S}$:

$$\widehat{p}_\star(s, \alpha)((\ell', \nu', \zeta')) = \sum_{C \subseteq \mathcal{C} \wedge \nu_\alpha[C:=0] = \nu' \wedge \zeta_\alpha[C:=0] = \zeta'} \delta[\ell, a_\alpha](C, \ell')$$

where $\nu_\alpha = \nu + \text{time}(\nu, (b_\alpha, c_\alpha))$ and one of the following conditions holds:

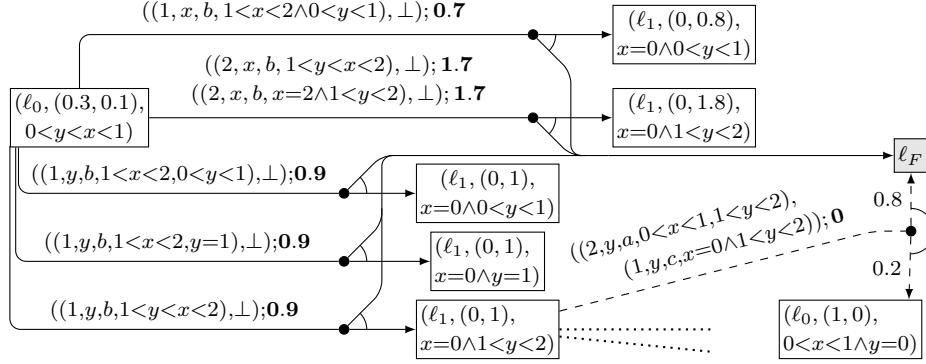


Fig. 2. Sub-graph of the boundary region abstraction for the ERTG of Figure 1.

- $\zeta \rightarrow b_\alpha, c_\alpha \zeta_\alpha$,
 - $\zeta \rightarrow b_\alpha, c_\alpha \zeta_{\text{inf}} \rightarrow \zeta_\alpha$ for some $\zeta_{\text{inf}} \in \mathcal{R}$
 - $\zeta \rightarrow b_\alpha, c_\alpha \zeta_{\text{sup}} \leftarrow \zeta_\alpha$ for some $\zeta_{\text{sup}} \in \mathcal{R}$;
- for $\star \in \{\text{Min}, \text{Max}\}$, $(\ell, \nu, \zeta) \in \widehat{S}$ and $(b_\alpha, c_\alpha, a_\alpha, \zeta_\alpha) \in \widehat{A}_\star$ the reward function $\widehat{\pi}_\star$ is given by $\widehat{\pi}_\star((\ell, \nu, \zeta), (b_\alpha, c_\alpha, a_\alpha, \zeta_\alpha)) = b_\alpha - \nu(c_\alpha)$.

Although the boundary region abstraction is not a finite ERG, for a fixed initial state we can restrict attention to a finite ERG, thanks to the following result of [15,16].

Lemma 12. *For any state of a boundary region abstraction, its reachable sub-graph is finite and is constructible in time exponential in the size of corresponding ERTG.*

Example 13. Sub-graph of BRA reachable from $(\ell_0, (0.3, 0.1), 0 < y < x < 1)$ for the ERTG of Figure 1 is shown in Figure 2. Edges are labelled (b, c, a, ζ) whose intuitive meaning is to wait until clock c attains the value b and then fire action a . The rewards of edges (indicated in bold) correspond to the time delay before the action is fired. Figure 2 includes the actions available in the initial state and one of action pairs available in $(\ell_1, (0, 1), x=0 \wedge 1 < y < 2)$. To simplify, the states with location ℓ_F are merged together into a single state labelled ℓ_F and probabilities that are equal to 0.5 are omitted.

4.2 Solving optimality equations of a boundary region abstraction

Based on the optimality equations $\text{Opt}^*(\widehat{\mathcal{T}})$ (see Section 2), we define the value improvement function $\Psi : [\widehat{S} \rightarrow \mathbb{R}_{\geq 0}] \rightarrow [\widehat{S} \rightarrow \mathbb{R}_{\geq 0}]$ where for $f : \widehat{S} \rightarrow \mathbb{R}_{\geq 0}$ and $s \in \widehat{S}$:

$$\Psi(f)(s) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } s \in \widehat{F} \\ \min_{\alpha \in \widehat{A}_{\text{Min}}(s)} \left\{ \max_{\beta \in \widehat{A}_{\text{Max}}(s)} \left\{ \widehat{\pi}(s, \alpha, \beta) + \sum_{s' \in \widehat{S}} \widehat{p}(s'|s, \alpha, \beta) \cdot f(s') \right\} \right\} & \text{if } s \notin \widehat{F} \end{cases}$$

By construction, a fixpoint of Ψ is a solution of $\text{Opt}^*(\widehat{\mathcal{T}})$. The following demonstrates the existence and uniqueness of a fixpoint of Ψ , and thus also the solution of $\text{Opt}^*(\widehat{\mathcal{T}})$.




Fig. 3. Example demonstrating optimal strategies are not regionally positional.

Proposition 14. For any ERTG \mathcal{T} , the value improvement function Ψ on the BRA $\widehat{\mathcal{T}}$ has a unique fixed point and equals $\lim_{i \rightarrow \infty} \Psi^i(f)$ for an arbitrary $f \in [\widehat{S} \rightarrow \mathbb{R}_{\geq 0}]$.

Proof. From Assumption 2 and Lemma 12 follows that every $|L \times \mathcal{R}|$ -th iterate of Ψ is a contraction. Hence using Banach fixed point theorem the result is immediate. \square

4.3 Correctness of the boundary region abstraction reduction

In this section we show how the optimality equations for the boundary region abstraction can be used to solve optimality equations for its ERTG. Given an ERTG \mathcal{T} and real-valued function $f : \widehat{S} \rightarrow \mathbb{R}$ on the states of the BRA $\widehat{\mathcal{T}}$, we define $\tilde{f} : S \rightarrow \mathbb{R}$ by $\tilde{f}(\ell, \nu) = f(\ell, \nu, [\nu])$ which gives a real-valued function on the states of \mathcal{T} . The following theorem states that, by applying this mapping, the solution of optimality equations for an ERTG is given by that of the optimality equations for its BRA.

Theorem 15. Let \mathcal{T} be an ERTG. If $P \models \text{Opt}^*(\widehat{\mathcal{T}})$, then $\tilde{P} \models \text{Opt}^*(\mathcal{T})$.

To prove Theorem 15 we first introduce quasi-simple functions and state some of their key properties. Next, we show that for any BRA $\widehat{\mathcal{T}}$ the solution of $\text{Opt}^*(\widehat{\mathcal{T}})$ is regionally quasi-simple (a quasi-simple function for every region). Finally, we sketch how Theorem 15 follows from this fact (Proposition 19 and Theorem 21).

Quasi-simple functions. Asarin and Maler [2] introduced simple functions, a finitely representable class of functions, with the property that every decreasing sequence is finite. Given $X \subseteq V$, a function $F : X \rightarrow \mathbb{R}$ is *simple* if there exists $e \in \mathbb{N}$ and either $F(\nu) = e$ for all $\nu \in X$, or there exists $c \in C$ and $F(\nu) = e - \nu(c)$ for all $\nu \in X$. A function $F : \widehat{S} \rightarrow \mathbb{R}_{\geq 0}$ is *regionally simple* if $F(\ell, \cdot, \zeta)$ is simple for all $\ell \in L$ and $\zeta \in \mathcal{R}$.

For timed games, Asarin and Maler showed that if $f : \widehat{S} \rightarrow \mathbb{R}_{\geq 0}$ is regionally simple, then $\Psi(f)$ is regionally simple. Therefore, since Ψ is a decreasing function, it follows that starting from a regionally simple function in finitely many iterations of Ψ a fixed point is reached and the upper value in reachability-time games is regionally simple. Also, using the properties of simple functions, [13] shows that for a non-probabilistic reachability-time game, the optimal strategies are *regionally positional*, i.e., in every state of a region the strategy chooses the same action. Unfortunately, in the case of ERTGs, $\Psi(f)$ is not necessarily regionally simple for any given regionally simple function f . Moreover, as the example below demonstrates, neither is the value of the game necessarily regionally-simple nor optimal strategies regionally positional.

Example 16. Consider the ERTG shown in Figure 3. Observe that for every state (ℓ_0, ν) in the region $(\ell_0, 0 < \nu < 1)$, the optimal expected time to reach ℓ_F equals $\min\{\inf_{t > 0} \{t + 0.5 \cdot 1 + 0.5 \cdot 0\}, 1 - \nu(x)\} = \min\{0.5, 1 - \nu(x)\}$. Hence optimal expected reachability-time is not regionally simple. Moreover, the optimal strategy is

not regionally positional, since if $\nu(x) \leq 0.5$, then the optimal strategy is to fire a immediately, while otherwise the optimal strategy is to wait until $\nu(x)=1$ and fire b .

Due to these results it is not possible to work with simple function, and we define quasi-simple functions. Let $\preceq \subseteq V \times V$ be the partial order on clock valuations, where $\nu \preceq \nu'$ if and only if there exists a $t \in \mathbb{R}_{\geq 0}$ such that for each clock $c \in C$ either $\nu'(c) - \nu(c) = t$ or $\nu(c) = \nu'(c)$. For $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we let $\|x\|_\infty = \max \{|x_i| \mid 1 \leq i \leq n\}$.

Definition 17. Let $X \subseteq V$. A function $F : X \rightarrow \mathbb{R}$ is quasi-simple if for all $\nu, \nu' \in X$:

- (Lipschitz Continuous) there exists $k \geq 0$ such that $|F(\nu) - F(\nu')| \leq k \cdot \|\nu - \nu'\|_\infty$;
- (Monotonically decreasing and nonexpansive w.r.t. \preceq) $\nu \preceq \nu'$ implies $F(\nu) \geq F(\nu')$ and $F(\nu) - F(\nu') \leq \|\nu - \nu'\|_\infty$.

For a convex set $X \subseteq V$ and continuous function $F : X \rightarrow \mathbb{R}$, we let $\bar{F} : \bar{X} \rightarrow \mathbb{R}$ denote the unique continuous function satisfying $\bar{F}(\nu) = F(\nu)$ for all $\nu \in X$.

Theorem 18 (Properties of Quasi-simple Functions). Let $X \subseteq V$.

1. Every simple function is also quasi-simple.
2. If $F : X \rightarrow \mathbb{R}$ is quasi-simple, then $\bar{F} : \bar{X} \rightarrow \mathbb{R}$ is quasi-simple.
3. If $F, F' : X \rightarrow \mathbb{R}$ are quasi-simple functions, then both the pointwise minimum and maximum of F and F' are quasi-simple.
4. The limit of a sequence of quasi-simple functions is quasi-simple.

We say that $f : \hat{S} \rightarrow \mathbb{R}_{\geq 0}$ is regionally quasi-simple if $f(\ell, \cdot, \zeta)$ is quasi-simple for all $\ell \in L$ and $\zeta \in \mathcal{R}$. Using Theorem 18 and Definition 11 we get the following result.

Proposition 19. If f is regionally quasi-simple, then $\Psi(f)$ is regionally quasi-simple.

From Proposition 14 it follows that for an arbitrary function $f : \hat{S} \rightarrow \mathbb{R}_{\geq 0}$ the limit of the sequence $\langle f, \Psi(f), \Psi^2(f), \dots \rangle$ is the solution of $\text{Opt}^*(\hat{\mathcal{T}})$. From Proposition 19 it follows that, if we start from a regionally quasi-simple function f , then all the functions in the sequence $\langle f, \Psi(f), \Psi^2(f), \dots \rangle$ are regionally quasi-simple. Since the limit of quasi-simple functions is quasi-simple, the following proposition is immediate.

Proposition 20. For any ERTG \mathcal{T} , if $P \models \text{Opt}^*(\hat{\mathcal{T}})$, then P is regionally quasi-simple.

The following result states that, from a regionally quasi-simple solution of the optimality equations for the boundary region abstraction, one can derive the solution of the optimality equations for the expected reachability time-game.

Theorem 21. For any ERTG \mathcal{T} , if $P \models \text{Opt}^*(\hat{\mathcal{T}})$ and P is regionally quasi-simple, then $\tilde{P} \models \text{Opt}^*(\mathcal{T})$.

The following observation is crucial for the proof of Theorem 21.

Lemma 22. Let $s = (\ell, \nu) \in S$ and $\zeta \in \mathcal{R}$ such that $[\nu] \xrightarrow{*} \zeta$. If $P : \hat{S} \rightarrow \mathbb{R}$ is regionally quasi-simple, then the functions:

$$t \mapsto t + \sum_{s' \in S} p(s'|s, (t, a), \perp) \cdot \tilde{P}(s') \quad \text{and} \quad t \mapsto t + \sum_{s' \in S} p(s'|s, \perp, (t, b)) \cdot \tilde{P}(s')$$

are continuous and nondecreasing on the interval $\{t \in \mathbb{R}_{\geq 0} \mid \nu + t \in \zeta\}$.

5 Conclusions

We introduced expected reachability-time games and showed that the natural decision problem is decidable and in $\text{NEXPTIME} \cap \text{co-NEXPTIME}$. Furthermore, we proved that the turn-based subclass of these games is positionally determined. We believe that the main contribution of this paper is the concept of quasi-simple function that generalise simple functions to the context of probabilistic timed games. In fact, the techniques introduced in this paper extend to expected discounted-time games (EDTGs)⁷ in a straightforward manner, since every expected discounted-time game can be reduced to an expected reachability-time game. Hence all the result presented for ERTGs are valid for EDTGs as well. Regarding other games on probabilistic timed automata, we conjecture that it is possible to reduce expected average-time games to mean payoff games on the boundary region abstraction. However, the techniques presented in this paper are insufficient to demonstrate such a reduction.

Although the computational complexity of solving games on timed automata is high, UPPAAL Tiga [5] is able to solve practical [6,10] reachability and safety properties for timed games by using efficient symbolic zone-based algorithms. A natural future work is to investigate the possibility of extending similar algorithms for probabilistic timed games.

Acknowledgements. We would like to thank anonymous reviewers for their helpful comments and for finding an error in the submitted version. The authors are supported in part by EPSRC grants EP/D076625, EP/F001096, EP/D07956X/2, and by research center ITI, project No. 1M0545.

References

1. Ramadge, P., Wonham, W.: The control of discrete event systems. In: Proc. IEEE. Volume 77(1). (1989)
2. Asarin, E., Maler, O.: As soon as possible: Time optimal control for timed automata. In: Proc. HSCC'99. Volume 1569 of LNCS., Springer (1999)
3. Alur, R., Bernadsky, M., Madhusudan, P.: Optimal reachability for weighted timed games. In: Proc. ICALP'04. Volume 3142 of LNCS., Springer (2004)
4. Bouyer, P., Cassez, F., Fleury, E., Larsen, K.G.: Optimal strategies in priced timed game automata. In: FSTTCS'04. Volume 3328 of LNCS., Springer (2004)
5. Behrmann, G., Cougnard, A., David, A., Fleury, E., Larsen, K., Lime, D.: UPPAAL-Tiga: Time for playing games! In: Proc. CAV'07. Volume 4590 of LNCS., Springer (2007)
6. Cassez, F., Jessen, J., Larsen, K., Raskin, J., Reynier, P.: Automatic synthesis of robust and optimal controllers: an industrial case study. In: HSCC'09. Volume 5469 of LNCS. (2009)
7. Kwiatkowska, M., Norman, G., Segala, R., Sproston, J.: Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science* **282** (2002)
8. Jensen, H.: Model checking probabilistic real time systems. In: Proc. 7th Nordic Workshop on Programming Theory. Report 86:247–261, Chalmers University of Technology (1996)
9. Beauquier, D.: Probabilistic timed automata. *Theoretical Computer Science* **292**(1) (2003)

⁷ In EDTG there is a fixed discount factor $\lambda \in [0,1)$, and when players follow strategies $\mu \in \Sigma_{\text{Min}}$ and $\chi \in \Sigma_{\text{Max}}$ the reward for state s is equal to $\mathbb{E}_s^{\mu, \chi} \left\{ \sum_{i=0}^{\infty} \lambda^i \cdot \pi(X_i, Y_{i+1}) \right\}$.

10. AlAttili, I., Houben, F., Igna, G., Michels, S., Zhu, F., Vaandrager, F.: Adaptive scheduling of data paths using Uppaal Tiga. CoRR **abs/0912.1897** (2009)
11. de Alfaro, L., Faella, M., Henzinger, T.A., Majumdar, R., Stoelinga, M.: The element of surprise in timed games. In: Proc. CONCUR'03. Volume 2761 of LNCS., Springer (2003)
12. Courcoubetis, C., Yannakakis, M.: Minimum and maximum delay problems in real-time systems. In: FMSD 1992. Volume 1., Dordrecht, Kluwer (1992)
13. Jurdziński, M., Trivedi, A.: Reachability-time games on timed automata. In: Proc. ICALP'07. Volume 4596 of LNCS., Springer (2007)
14. Jurdziński, M., Trivedi, A.: Average-time games. In: Proc. FSTTCS'08. Volume 2 of Leibniz International Proceedings in Informatics., Schloss Dagstuhl (2008)
15. Jurdziński, M., Trivedi, A.: Concavely-priced timed automata. In: Proc. FORMATS'08. Volume 5215 of LNCS., Springer (2008)
16. Jurdziński, M., Kwiatkowska, M., Norman, G., Trivedi, A.: Concavely-priced probabilistic timed automata. In: Proc. CONCUR'09. Volume 5710 of LNCS., Springer (2009)
17. Chadha, R., Legay, A., Prabhakar, P., Viswanathan, M.: Complexity bounds for the verification of real-time software. In: Proc. VMCAI'10. Volume 5944 of LNCS. (2010)
18. Hoffmann, G., Wong-Toi, H.: The input-output control of real-time discrete event systems. In: Proc. RTS'92, IEEE Press (1992)
19. Asarin, E., Maler, O., Pnueli, A.: Symbolic controller synthesis for discrete and timed systems. In: Hybrid Systems II. Volume 999 of LNCS., Springer (1995)
20. Brihaye, T., Henzinger, T.A., Prabhu, V.S., Raskin, J.: Minimum-time reachability in timed games. In: Proc. ICALP'07. Volume 4596 of LNCS., Springer (2007)
21. Brihaye, T., Bruyère, V., Raskin, J.: On optimal timed strategies. In: Proc. FORMATS'05. Volume 3829 of LNCS., Springer (2005)
22. Bouyer, P., Brihaye, T., Markey, N.: Improved undecidability results on weighted timed automata. Information Processing Letters **98** (2006)
23. Bouyer, P., Forejt, V.: Reachability in stochastic timed games. In: Proc. ICALP'09. Volume 5556 of LNCS., Springer (2009)
24. Kwiatkowska, M., Norman, G., Parker, D., Sproston, J.: Performance analysis of probabilistic timed automata using digital clocks. FMSD **29** (2006)
25. Berendsen, J., Chen, T., Jansen, D.: Undecidability of cost-bounded reachability in priced probabilistic timed automata. In: Proc. TAMC'09. Volume 5532 of LNCS., Springer (2009)
26. Forejt, V., Kwiatkowska, M., Norman, G., Trivedi, A.: Expected reachability-time games. Technical Report RR-10-07, Oxford University Computing Laboratory (2010)
27. Kwiatkowska, M., Norman, G., Trivedi, A.: Quantitative games on probabilistic timed automata. CoRR **abs/1001.1933** (2010)
28. Neyman, A., Sorin, S., eds.: Stochastic Games and Applications. Volume 570 of NATO Science Series C. Kluwer Academic Publishers (2004)
29. Condon, A.: The complexity of stochastic games. Information and Computation **96** (1992)
30. Filar, J., Vrieze, K.: Competitive Markov Decision Processes. Springer (1997)
31. Bouyer, P., Brihaye, T., Bruyère, V., Raskin, J.: On the optimal reachability problem on weighted timed automata. FMSD **31**(2) (2007)
32. Alur, R., Dill, D.: A theory of timed automata. Theoretical Computer Science **126** (1994)
33. Jurdziński, M., Sproston, J., Laroussinie, F.: Model checking probabilistic timed automata with one or two clocks. Logical Methods in Computer Science **4**(3) (2008)
34. Bouyer, P., Brinksma, E., Larsen, K.G.: Staying alive as cheaply as possible. In: Proc. HSCC'04. Volume 2993 of LNCS., Springer (2004)