Expected Reachability-Time Games

Vojtěch Forejt[a], Marta Kwiatkowska[a], Gethin Norman[b], Ashutosh Trivedi[c]

[a]*Department of Computer Science, University of Oxford, UK.*
[b]*School of Computing Science, University of Glasgow, UK.*
[c]*Department of Computer Science and Engineering, IIT Bombay, India.*

Abstract

Probabilistic timed automata are a suitable formalism to model systems with real-time, nondeterministic and probabilistic behaviour. We study two-player zero-sum games on such automata where the objective of the game is specified as the expected time to reach a target. The two players—called player Min and player Max—compete by proposing timed moves simultaneously and the move with a shorter delay is performed. The first player attempts to minimise the given objective while the second tries to maximise the objective. We observe that these games are not determined, and study decision problems related to computing the upper and lower values, showing that the problems are decidable and lie in the complexity class NEXPTIME ∩ co-NEXPTIME.

Keywords: Probabilistic Timed Automata, Two-Player Games, Competitive Optimisation, Controller Synthesis

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 controller synthesis corresponds to finding a winning (or optimal) strategy of the controller for some given performance objective. Timed automata [2] extend finite automata by providing a mechanism to model real-time behaviour, while priced timed automata are timed automata with (time-dependent) prices attached to the locations of the automata. If the game structure or objectives are dependent on time or price, e.g. when the objective corresponds to completing a given set of tasks within some deadline or within some cost, then games on timed automata are a well-established approach for controller synthesis, see e.g. [3, 4, 5, 6, 7].

In this paper we extend the above approach to a setting that is quantitative in terms of both 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 an extension of probabilistic time automata (PTA) [8, 9, 10], a model for real-time systems exhibiting nondeterministic and probabilistic behaviour.

In our model, called probabilistic timed game arena (PTGA), 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 exactly one of the players). Once the players have made their choices, the timed move with the shorter delay is performed and the token is moved according to the probabilistic transition function of the PTA. Intuitively, players 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.

Players Min and Max choose their moves in order to minimise and maximise, respectively, the objective function. The *upper value* of a game is the minimum expected time that Min can ensure, while the *lower value* of a game is the maximum expected value 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.

The objectives frequently studied include reachability, which asks for certain locations to be eventually reached, safety, which asks for a given target set to be avoided, or more complex properties, expressed using a formula of a linear temporal logic. The objective function is then an indicator function saying whether the property is satisfied on a play, and the expected value then corresponds to the *probability* of the property being true. In our paper we are interested in a more complex setting and study *reachability-time* time objectives, which express the *expected time* to reach a given target set. These objectives have 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. [11, 7]. The reachability-time objectives are a special case of *weight* or *price* objectives in which different numbers are assigned to locations, and the value of the objective function depends on the respective numbers and the time spent in the locations; in our setting, the numbers are fixed to be 1 and the objective function simply sums the times spent in for each location. Computing properties related to price functions often leads to undecidability, even in non-probabilistic setting [12, 13]. Studying simpler properties is thus motivated by the desire to obtain decidable properties while still being able to study sufficiently complex class of properties.

1.1. Contribution

We demonstrate the decidability of the problem of whether the upper (lower, or the optimal when it exists) value of a game with reachability-time objectives is at most a given bound. Our proofs immediately yield a NEXPTIME $\cap$ co-NEXPTIME complexity bound. To our best knowledge, this is the first decidability result for stochastic games on timed automata in which the objective concerns a random variable that takes non-binary values.

Our approach is based on extending the boundary region graph construction for timed automata [14] to PTGAs and demonstrating that the reachability-time problem can be reduced to the same problem on the boundary region graph. In particular, our proof aims to show that the limit of the step-bounded value functions in the timed automata and boundary region graph also coincide.

Generic results exist that allow one to prove that step-bounded values converge to the step-unbounded value, but to the best of our knowledge none are readily applicable in our setting where the state space is uncountable and little is known a priori about the value functions. For example, Banach fixpoint theorem requires the value iteration function (that takes a n-step value function and returns the $n+1$-step value function) to be a contraction on an underlying metric space, and it appears difficult to devise the metric space so that the contraction property is easily obtained. Another possible proof direction is Kleene fixpoint theorem, which requires Scott-continuity on the value functions, which again is a property that is difficult to establish in our setting. We are able to partly rely on the Knaster-Tarski fixpoint theorem which characterises the set of fixpoints, but it is not strong enough to prove the convergence itself, for reasons similar to the ones above. Several other theorems such as Brouwer fixpoint theorem or Kakutani fixpoint theorem are generally not suitable for proving properties that we require in turn-based stochastic games.

Hence, to prove that the limit of the step-bounded value functions is the desired value, we need to take a tailor-made approach. We first inductively show that, when the number of steps is bounded, then the value functions in timed automata and boundary region graph coincide and are non-expansive within a region. Here we make use of *quasi-simple functions* which generalise simple functions, previously used by Asarin and Maler in the study of games over non-probabilistic timed automata [3]. Then, using the non-expansiveness property, we show that the limit of the step-bounded value functions in the timed automata and boundary region graph also coincide. In this part we use Knaster-Tarski fixpoint theorem.

The definition of quasi-simple functions is a central component of our proof, as it is strong enough to enable us to utilise an approach used in proofs of fixpoint theorems, but on the other hand general enough to capture the values of reachability-time objectives. We believe that it can serve as a step from simple functions towards functions describing even more complex but still decidable objectives.

1.2. Related Work

Hoffman and Wong-Toi [15] were the first to define and solve the optimal controller synthesis problem for timed automata. For a detailed introduction to the topic of qualitative games on timed automata, see e.g. [16]. Asarin and Maler [3] initiated the study of quantitative games on timed automata by providing a symbolic algorithm to solve reachability-time objectives. The works of [17] and [14] show that the decision problem for such games over timed automata with at least two clocks is EXPTIME-complete. The tool UPPAAL Tiga [6] is capable of solving reachability and safety objectives for games on timed automata. Jurdziński and Trivedi [18] show the EXPTIME-completeness for average-time games on automata with two or more clocks.

A natural extension of games with reachability-time objectives are games on priced timed automata where the objective concerns the cumulated price of reaching a target. Both [4] and [5] present semi-algorithms for computing the value of such games for linear prices. In [12] the problem of checking the existence of optimal strategies is shown to be undecidable, with [13] showing undecidability holds even for three clocks and stopwatch prices.

As for 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 [19]. In [20], continuous-time games are verified against time-automata objectives, giving rise to systems whose semantics is related to the ones of [19]. The work of [21] studies probability of satisfying Büchi objectives in a timed game where perturbations of probabilities can take place, and [22] studies games on interactive Markov chains which are modelled as a game extension of timed automata.

Regarding one-player games on PTAs, in [23] 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. The work of [24] shows that the problem becomes decidable when the price functions are of a restricted form. In [25], simple functions are extended to devise a symbolic algorithm for computing minimum expected time to reach a target in one-player games on PTAs; the extension differs fundamentally from our quasi-simple functions. We also mention the approaches for analysing unpriced probabilistic timed automata against temporal logic specifications based on the region graph [8, 9] and either forwards [8] or backwards [26] reachability. The complexity of performing such verification is studied in [27] for almost-sure reachability, and in [28] for PCTL properties and a restricted number of clocks. Finally, [29] deals with a model similar to PTAs in which time evolves continuously and controllable "fixed delay" events are introduced.

A preliminary version of the work was published in conference proceedings [30]. The result presented in [30] required an assumption on the structure of the PTAs that enforced a terminal state to be reached almost surely under any pair of strategies. In this paper we *lift this restriction* and consider arbitrary PTAs. Further, the proofs in [30] contain a significant flaw which required major changes to be made to the proof, also for the restricted case. Thus, although the high-level idea behind the proof (bounding the difference of values for two configurations whose clock values are close to each other) stays the same, the actual steps of the proof changed significantly. Note that, although [30] also introduces quasi-simple functions, the definition used here is different (and not equivalent). Most notably, our proofs here use a much more "constructive" approach when defining value functions.

1.3. Outline

The structure of the paper is the following. In Section 2 we introduce Stochastic Games Arenas, which serve as semantics for games on PTAs. Games on PTAs are then introduced in Section 3, and Section 4 defines boundary region abstraction, which plays a fundamental role in our proofs. Section 5 provides the proofs for the main result.

2. Stochastic Game Arena

We now introduce a general notion of stochastic game arenas (SGAs), which will later serve as semantics for the model we study. The reader may notice that our definition of a stochastic game arena differs from the standard concurrent stochastic game arena [31, 32]. However, as we shall demonstrate later, it captures precisely the semantics of probabilistic timed game arena. In addition, presenting the basic concepts relating to

values in the general setting of SGAs allows us to use these concepts in the context of both probabilistic timed game arenas and their abstractions.

2.1. Stochastic Game Arena: Syntax and Semantics

We write $\mathbb{N}$ for the set of non-negative integers, $\mathbb{Q}$ for the rational numbers, $\mathbb{R}_{\geqslant 0}$ for the non-negative reals, and $\mathbb{R}^{\infty}_{\geqslant 0}$ for the reals with the maximum element ∞. A function $f : (\mathbb{R}^{\infty}_{\geqslant 0})^n \to \mathbb{R}^{\infty}_{\geqslant 0}$ is *non-expansive* if for any $\mathbf{x}, \mathbf{y} \in (\mathbb{R}^{\infty}_{\geqslant 0})^n$ we have $|f(\mathbf{x}) - f(\mathbf{y})| \leqslant |\mathbf{x} - \mathbf{y}|$ where $|\cdot|$ is the max norm, i.e. $|(x_1, \ldots, x_n)| = \max_{1 \leqslant i \leqslant n} |x_i|$. A *discrete probability distribution*, or just distribution, over a (possibly uncountable) set Q is a function $d : Q \to [0, 1]$ such that $\sum_{q \in Q} d(q) = 1$ and $\mathit{supp}(d) \stackrel{\text{def}}{=} \{q \in Q \mid d(q) > 0\}$ is at most countable. 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 $q \in Q$. Given a set Q and two functions $f, f' : Q \to \mathbb{R}^{\infty}_{\geqslant 0}$, we state $f \leqslant f'$ when $f(q) \leqslant f'(q)$ for all $q \in Q$. A function $f : Q \to \mathbb{R}_{\geqslant 0}$ is a *convex combination* of functions $f_1, \ldots, f_n : Q \to \mathbb{R}_{\geqslant 0}$ if there are non-negative coefficients $p_1, \ldots, p_n$ such that $\sum_{i=1}^{n} p_i = 1$ and we have $f(q) = \sum_{i=1}^{n} p_i \cdot f_i(q)$ for all $q \in Q$.

Definition 1 (Stochastic Game Arena (SGA)). *A stochastic game arena is given by a tuple* $\mathsf{G} = (S, A_{\mathrm{Min}}, A_{\mathrm{Max}}, p_{\mathrm{Min}}, p_{\mathrm{Max}}, \mathsf{win}, \tau_{\mathrm{Min}}, \tau_{\mathrm{Max}})$ *where:*

- S *is a possibly uncountable set of* states*;*
- A_{Min} *and* A_{Max} *are possibly uncountable sets of* actions controlled *by players* Min *and* Max *respectively, and* $\bot$ *is a distinguished action such that* $A_{\mathrm{Min}} \cap A_{\mathrm{Max}} = \{\bot\}$*;*
- $p_{\mathrm{Min}} : (S \times A_{\mathrm{Min}}) \rightharpoonup \mathcal{D}(S)$ *and* $p_{\mathrm{Max}} : (S \times A_{\mathrm{Max}}) \rightharpoonup \mathcal{D}(S)$ *are the probabilistic transition (partial) functions for players* Min *and* Max *respectively, such that* $p_{\mathrm{Min}}(s, \bot)$ *and* $p_{\mathrm{Max}}(s, \bot)$ *are undefined for all* $s \in S$*, and for any* $s \in S$ *either there exists* $\alpha \in A_{\mathrm{Min}}$ *such that* $p_{\mathrm{Min}}(s, \alpha)$ *is defined or there exists* $\alpha \in A_{\mathrm{Max}}$ *such that* $p_{\mathrm{Max}}(s, \alpha)$ *is defined;*
- $\mathsf{win} : (A_{\mathrm{Min}} \times A_{\mathrm{Max}}) \to (A_{\mathrm{Min}} \cup A_{\mathrm{Max}})$ *is a function specifying which of the actions chosen by the players takes place, requiring that for any* $(\alpha, \beta) \in A_{\mathrm{Min}} \times A_{\mathrm{Max}}$ *we have* $\mathsf{win}(\alpha, \beta) \in \{\alpha, \beta\}$*, and moreover* $\mathsf{win}(\alpha, \beta)$ *is never equal to* $\bot$ *unless* $\alpha = \beta = \bot$*;*
- $\tau_{\mathrm{Min}} : (S \times A_{\mathrm{Min}}) \rightharpoonup \mathbb{R}_{\geqslant 0}$ *and* $\tau_{\mathrm{Max}} : (S \times A_{\mathrm{Max}}) \rightharpoonup \mathbb{R}_{\geqslant 0}$ *are the time delay (partial) functions for players* Min *and* Max *respectively, specifying the* delay *associated with performing an action in a state.*

Note that SGAs introduced above are more general than classical stochastic games, in particular SGAs contain information about the time delays of actions.

We say that an SGA is *finite* if S, A_{Min} and A_{Max} are finite. For any state $s \in S$, we let $A_{\mathrm{Min}}(s)$ denote the set of actions available to player Min in s, i.e., the actions $\alpha \in A_{\mathrm{Min}}$ for which $p_{\mathrm{Min}}(s, \alpha)$ is defined, letting $A_{\mathrm{Min}}(s) = \{\bot\}$ if no such action exists. Similarly, $A_{\mathrm{Max}}(s)$ denotes the actions available to player Max in s and we let

$A(s)=A_{\mathrm{Min}}(s)\times A_{\mathrm{Max}}(s)$. From the conditions required of the probabilistic transition functions of the players, we have $(\bot,\bot)\notin A(s)$ for all $s\in S$.

A game on SGA G starts with a token in an *initial state* $s\in S$ 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 the probabilistic transition function of the player proposing the action that is favoured by the win function. Formally, we introduce the following auxiliary definition for an SGA.

Definition 2 (Probabilistic transition function of an SGA). *For any stochastic game arena* $\mathsf{G}=(S,A_{\mathrm{Min}},A_{\mathrm{Max}},p_{\mathrm{Min}},p_{\mathrm{Max}},\tau_{\mathrm{Min}},\tau_{\mathrm{Max}})$ *the* probabilistic transition function *of* G *is given by the partial function* $p:(S\times A_{\mathrm{Min}}\times A_{\mathrm{Max}})\rightharpoonup\mathcal{D}(S)$ *where for any* $s\in S$, $\alpha\in A_{\mathrm{Min}}$ *and* $\beta\in A_{\mathrm{Max}}$*:*

$$p(s,\alpha,\beta)=\begin{cases}\textit{undefined} & \textit{if } (\alpha,\beta)\notin A(s)\\ p_{\mathrm{Min}}(s,\alpha) & \textit{if } \mathsf{win}(\alpha,\beta)=\alpha\\ p_{\mathrm{Max}}(s,\beta) & \textit{otherwise.}\end{cases}$$

Using this definition, if we are in state s and the action pair $(\alpha,\beta)\in A(s)$ is chosen by the players, then the probability of making a transition to s' equals $p(s,\alpha,\beta)(s')$. We similarly define the time delay function τ of the SGA G by

$$\tau(s,\alpha,\beta)\stackrel{\text{def}}{=}\begin{cases}\text{undefined} & \text{if } (\alpha,\beta)\notin A(s)\\ \tau_{\mathrm{Min}}(s,\alpha) & \text{if } \mathsf{win}(\alpha,\beta)=\alpha\\ \tau_{\mathrm{Max}}(s,\beta) & \text{otherwise.}\end{cases}$$

A transition of G is a tuple $(s,(\alpha,\beta),s')$ such that $p(s,\alpha,\beta)(s')>0$ and a play of G is a finite or infinite sequence

$$\langle s_0,(\alpha_1,\beta_1),s_1,(\alpha_2,\beta_2),\ldots,s_i,(\alpha_{i+1},\beta_{i+1}),s_{i+1},\ldots\rangle$$

such that $(s_i,(\alpha_{i+1},\beta_{i+1}),s_{i+1})$ is a transition for all $i\geqslant 0$. The length of a play ρ, denoted $len(\rho)$, is defined as the number of transitions appearing in the play. For a finite play $\rho=\langle s_0,(\alpha_1,\beta_1),s_1,\ldots,(\alpha_k,\beta_k),s_k\rangle$, let $last(\rho)$ denote the last state s_k of the play. We write Ω (Ω_f) for the sets of infinite (finite) plays in G and $\Omega(s)$ ($\Omega_f(s)$) for the sets of infinite (finite) plays starting from $s\in S$.

Definition 3 (SGA Strategy). *Let* $\mathsf{G}=(S,A_{\mathrm{Min}},A_{\mathrm{Max}},p_{\mathrm{Min}},p_{\mathrm{Max}},\tau_{\mathrm{Min}},\tau_{\mathrm{Max}})$ *be a SGA. A* strategy *of* Min *is a function* $\mu:\Omega_f\to A_{\mathrm{Min}}$ *such that* $\mu(\rho)\in A_{\mathrm{Min}}(last(\rho))$ *for all finite plays* $\rho\in\Omega_f$*. 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.*

For any finite play, a strategy of Min (Max) returns an action available to Min (Max) in the last state of the play.

For a SGA G, state s of G and strategy pair $(\mu,\chi)\in\Sigma_{\mathrm{Min}}\times\Sigma_{\mathrm{Max}}$, let $\Omega^{\mu,\chi}(s)$ ($\Omega_f^{\mu,\chi}(s)$) denote the set of infinite (finite) plays in which Min and Max play according to μ and χ, respectively. Given a finite play $\rho\in\Omega_f^{\mu,\chi}(s)$, a basic cylinder set $Cyl(\rho)$ is the set of infinite plays in $\Omega^{\mu,\chi}(s)$ for which ρ is a prefix. Using standard results from probability theory [33] we can construct a probability space

$(\Omega^{\mu,\chi}(s), \mathcal{F}^{\mu,\chi}(s), Pr_s^{\mu,\chi})$ where $\mathcal{F}^{\mu,\chi}(s)$ is the smallest σ-algebra generated by the basic cylinder sets and $Pr_s^{\mu,\chi} : \mathcal{F} \to [0,1]$ is the unique probability measure such that for any finite play $\rho = \langle s_0, (\alpha_1, \beta_1), s_1, \ldots, s_{k-1}, (\alpha_k, \beta_k), s_k \rangle \in \Omega_f^{\mu,\chi}(s)$:

$$Pr_s^{\mu,\chi}(Cyl(\rho)) = \prod_{i=1}^{k} p(s_{i-1}, \alpha_i, \beta_i)(s_i)$$

where $pre(\rho, i) = \langle s_0, (\alpha_1, \beta_1), s_1, \ldots, s_{i-1}, (\alpha_i, \beta_i), s_i \rangle$ for all $i < k$.

Given a *real-valued random variable* $f : \Omega \to \mathbb{R}_{\geqslant 0}^{\infty}$, the expression $\mathbb{E}_s^{\mu,\chi}(f)$ denotes the expected value of f with respect to the probability measure $Pr_s^{\mu,\chi}$.

We extend $Pr_s^{\mu,\chi}(f)$ also to the cases where the game is assumed to start from a finite play $\rho = \langle s_0, (\alpha_1, \beta_1), s_1, \ldots, (\alpha_k, \beta_k), s_k \rangle$ as opposed to a state s, and we let $Pr_\rho^{\mu,\chi} = Pr_{last(\rho)}^{\mu_\rho, \chi_\rho}$, where the strategy μ_ρ is defined from μ by

$$\mu_\rho(\rho') = \mu(\langle s_0, (\alpha_1, \beta_1), s_1, \ldots, (\alpha_k, \beta_k), s_k, (\alpha'_{k+1}, \beta'_{k+1}), s'_{k+1}, \ldots (\alpha_\ell, \beta_\ell), s_\ell \rangle)$$

for all $\rho' = \langle s'_k, \ldots (\alpha'_k, \beta'_k), s'_{k+1}, \ldots (\alpha_\ell, \beta_\ell), s_\ell \rangle$ such that $s_k = s'_k$, and $\mu_\rho(\rho')$ is defined arbitrarily otherwise; the strategy χ_ρ is defined analogously. We then also use $\mathbb{E}_\rho^{\mu,\chi}$ defined with respect to $Pr_\rho^{\mu,\chi}$.

2.2. Reachability-time objective in Stochastic Game Arena

We now define the reachability-time objective for plays of SGAs.

Definition 4. *For an SGA* G *and target set of states* F *of* G*, the* (finite-horizon) n-step reachability-time objective *associated with an infinite play* $\rho = \langle s_0, (\alpha_1, \beta_1), s_1, \ldots \rangle$ *is given by:*

$$\mathsf{Reach}_F^n(\rho) = \textstyle\sum_{i=1}^{k} \tau(s_{i-1}, \alpha_i, \beta_i)$$

where $k = \min\{i \in \mathbb{N} \mid s_i \in F\}$ *if* $s_j \in F$ *for some* $j < n \in \mathbb{N}$ *and* $k = n$ *otherwise. Furthermore, the* (infinite-horizon) reachability-time objective *(with target set* $F \subseteq S$*) associated with an infinite play* ρ *is given by:*

$$\mathsf{Reach}_F(\rho) = \lim_{n \to \infty} \mathsf{Reach}_F^n(\rho)\,.$$

In the definition of the infinite horizon objective the limit always exists, but it can be infinite. To simplify notation, we often omit the target set F when it is clear from the context.

In our games on an SGAs players Min and Max move a token along the edges in order to minimise and maximise, respectively, the (n-step) reachability-time objective function. Formally, for an SGA G and an objective Reach^n we define lower and upper value with respect to Reach^n for G in state $s \in S$ by

$$\underline{\mathsf{Val}}_{\mathsf{G}}^n(s) \stackrel{\text{def}}{=} \sup_{\chi \in \Sigma_{\mathrm{Max}}} \inf_{\mu \in \Sigma_{\mathrm{Min}}} \mathbb{E}_s^{\mu,\chi}(\mathsf{Reach}^n)$$
$$\overline{\mathsf{Val}}_{\mathsf{G}}^n(s) \stackrel{\text{def}}{=} \inf_{\mu \in \Sigma_{\mathrm{Min}}} \sup_{\chi \in \Sigma_{\mathrm{Max}}} \mathbb{E}_s^{\mu,\chi}(\mathsf{Reach}^n)$$

respectively. Similarly, for an objective Reach we define the lower and upper values:

$$\underline{\mathsf{Val}}_{\mathsf{G}}(s) \stackrel{\text{def}}{=} \sup_{\chi \in \Sigma_{\mathrm{Max}}} \inf_{\mu \in \Sigma_{\mathrm{Min}}} \mathbb{E}_s^{\mu,\chi}(\mathsf{Reach})$$

$$\overline{\mathsf{Val}}_{\mathsf{G}}(s) \stackrel{\text{def}}{=} \inf_{\mu\in\Sigma_{\mathrm{Min}}} \sup_{\chi\in\Sigma_{\mathrm{Max}}} \mathbb{E}_s^{\mu,\chi}(\mathsf{Reach}).$$

In the cases when the lower and upper values coincide, we denote this value simply as $\mathsf{Val}^n_{\mathsf{G}}(s)$ or $\mathsf{Val}_{\mathsf{G}}(s)$ and say that the corresponding game is *determined*. We omit G if it is clear from the context, e.g. we write simply $\overline{\mathsf{Val}}$ instead of $\overline{\mathsf{Val}}_{\mathsf{G}}$.

For $\mu \in \Sigma_{\mathrm{Min}}$, $\chi \in \Sigma_{\mathrm{Max}}$ and $s \in S$, let

$$\mathsf{Val}_{\mathsf{G}}(s,\mu) \stackrel{\text{def}}{=} \sup_{\chi'\in\Sigma_{\mathrm{Max}}} \mathbb{E}_s^{\mu,\chi'}(\mathsf{Reach}) \;\text{ and }\; \mathsf{Val}_{\mathsf{G}}(s,\chi) \stackrel{\text{def}}{=} \inf_{\mu'\in\Sigma_{\mathrm{Min}}} \mathbb{E}_s^{\mu',\chi}(\mathsf{Reach})\,.$$

We say μ is *optimal* (or ε-*optimal*), if $\mathsf{Val}_{\mathsf{G}}(s,\mu){=}\overline{\mathsf{Val}}_{\mathsf{G}}(s)$ (or $\mathsf{Val}_{\mathsf{G}}(s,\mu) \leqslant \overline{\mathsf{Val}}_{\mathsf{G}}(s){-}\varepsilon$) for all $s \in S$. Furthermore, χ is *optimal* (or ε-*optimal*), if $\mathsf{Val}_{\mathsf{G}}(s,\chi){=}\overline{\mathsf{Val}}_{\mathsf{G}}(s)$ (or $\mathsf{Val}_{\mathsf{G}}(s,\chi) \geqslant \underline{\mathsf{Val}}_{\mathsf{G}}(s){-}\varepsilon$) for all $s \in S$. If G is determined, then each player has an ε-optimal strategy for all $\varepsilon{>}0$.

Since we will consider two-player games on SGAs that are not determined, we are interested in the following problem with respect to the upper value of a game.

Definition 5. *Given an SGA* G*, initial state* $s \in S$*, reachability-time objective and value* $B \in \mathbb{Q}$*, the corresponding game* reachability-time problem *is to decide whether* $\overline{\mathsf{Val}}(s) \leqslant B$.

All results presented in the paper are still valid when replacing the upper value with the lower value. The following is a well-known result.

Theorem 6 ([34, 35]). *The reachability-time problem for infinite-horizon objectives over finite SGAs is in NP ∩ co-NP.*

Efficient algorithms exist to solve the problem over finite SGAs, e.g. using value iteration [36, 37].

2.3. Optimality Equations for SGAs

We now introduce optimality equations for reachability objectives over SGAs. For the remainder of this section we fix an SGA $\mathsf{G}{=}(S, A_{\mathrm{Min}}, A_{\mathrm{Max}}, p_{\mathrm{Min}}, p_{\mathrm{Max}}, \tau_{\mathrm{Min}}, \tau_{\mathrm{Max}})$ and a target set $F \subseteq S$.

Definition 7. *The Bellman-style equations for* n*-step reachability time objective are given as follows:* $\overline{\mathsf{Val}}^n(s){=}0$ *whenever* $n{=}0$ *or* $s \in F$*, and for* $n{\geqslant}0$ *and* $s \notin F$*:*

$$\overline{\mathsf{Val}}^{n+1}(s) = \inf_{\alpha\in A_{\mathrm{Min}}(s)} \sup_{\beta\in A_{\mathrm{Max}}(s)} \left\{\tau(s,\alpha,\beta) + \textstyle\sum_{s'\in S} p(s,\alpha,\beta)(s')\cdot\overline{\mathsf{Val}}^n(s')\right\}$$

The correctness of these equations can be easily obtained from the fact that for any $n{\geqslant}0$, $s \notin F$, path ρ with $last(\rho){=}s$ and strategies μ and χ, where $\alpha{=}\mathsf{win}(\mu(\rho),\chi(\rho))$, by definition of $\mathbb{E}_\rho^{\mu,\chi}$ we have:

$$\begin{aligned}
&\mathbb{E}_\rho^{\mu,\chi}(\mathsf{Reach}_F^{n+1}) = \mathbb{E}_s^{\mu_\rho,\chi_\rho}(\mathsf{Reach}_F^{n+1}) \\
&= \int_{\rho'\in\Omega^{\mu_\rho,\chi_\rho}(s)} \mathsf{Reach}_F^{n+1}(\rho')\, dPr_s^{\mu_\rho,\chi_\rho} \qquad \text{(by definition of expectation)}
\end{aligned}$$

$$= \sum_{s' \in S} \left(\int_{\bar{\rho} \in \Omega^{\mu_{\rho\alpha s'}, \chi_{\rho\alpha s'}}(s')} p(s, \mu_\rho(s), \chi_\rho(s))(s') \right.$$
$$\left. \cdot \left(\tau(s, \mu_\rho(s), \chi_\rho(s)) + \mathsf{Reach}_F^n(\bar{\rho}) \right) dPr_{s'}^{\mu_{\rho\alpha s'}, \chi_{\rho\alpha s'}}(\bar{\rho}) \right)$$
(by definition of $\Omega^{\mu_\rho, \chi_\rho}(s)$, $Pr_s^{\mu_\rho, \chi_\rho}$ and Reach_F^{n+1})

$$= \tau(s, \mu_\rho(s), \chi_\rho(s))$$
$$+ \sum_{s' \in S} p(s, \mu_\rho(s), \chi_\rho(s))(s') \cdot \left(\int_{\bar{\rho} \in \Omega^{\mu_{\rho\alpha s'}, \chi_{\rho\alpha s'}}(s')} \mathsf{Reach}_F^n(\bar{\rho})\, dPr_{s'}^{\mu_{\rho\alpha s'}, \chi_{\rho\alpha s'}}(\bar{\rho}) \right)$$
(rearranging)

$$= \tau(s, \mu(\rho), \chi(\rho)) + \sum_{s' \in S} p(s, \mu(\rho), \chi(\rho))(s') \cdot \mathbb{E}^{\mu,\chi}_{\rho\,\alpha\,s'}(\mathsf{Reach}_F^n)$$
(by properties of μ_ρ, χ_ρ and definition of expectation)

Let us now turn to the equations for infinite-horizon objectives.

Definition 8. *A function $P : S \to \mathbb{R}^\infty_{\geqslant 0}$ is a solution of the optimality equations $\overline{\mathsf{Opt}}_\mathsf{G}$, written $P \models \overline{\mathsf{Opt}}_\mathsf{G}$, if for any $s \in S$:*

$$P(s) = \begin{cases} 0 & \text{if } s \in F \\ \inf\limits_{\alpha \in A_{\mathrm{Min}}(s)} \sup\limits_{\beta \in A_{\mathrm{Max}}(s)} \left\{ \tau(s,\alpha,\beta) + \sum\limits_{s' \in S} p(s,\alpha,\beta)(s') \cdot P(s') \right\} & \text{if } s \notin F \end{cases}$$

and is a solution of the optimality equations $\underline{\mathsf{Opt}}_\mathsf{G}$, written $P \models \underline{\mathsf{Opt}}_\mathsf{G}$, if for any $s \in S$:

$$P(s) = \begin{cases} 0 & \text{if } s \in F \\ \sup\limits_{\beta \in A_{\mathrm{Max}}(s)} \inf\limits_{\alpha \in A_{\mathrm{Min}}(s)} \left\{ \tau(s,\alpha,\beta) + \sum\limits_{s' \in S} p(s,\alpha,\beta)(s') \cdot P(s') \right\} & \text{if } s \notin F. \end{cases}$$

To simplify the presentation, from now we will only concentrate on upper value $\overline{\mathsf{Val}}$. Analogous results for the lower value follow in a straightforward manner.

Our aim is to utilise the optimality equations for $\overline{\mathsf{Opt}}_\mathsf{G}$ and prove that $\overline{\mathsf{Val}}$ and $\lim_{n\to\infty} \overline{\mathsf{Val}}^n$ are equal, as an initial step towards computing or approximating $\overline{\mathsf{Val}}$. Although this equivalence can seem obvious, it is not at all trivial and, due to the uncountable nature of SGAs, it is not possible to use results such as Kleene fixpoint theorem out of the box. In fact, in this paper we will only prove the equivalence for a special case of SGAs (sufficient for our purpose). Nevertheless, the following two lemmas can be established for SGAs in general.

Lemma 9. *For any solution $V \models \overline{\mathsf{Opt}}_\mathsf{G}$ we have $\overline{\mathsf{Val}} \leqslant V$.*

PROOF. Consider any $\varepsilon{>}0$ and let μ be a strategy for player Min that, for any finite play ρ, selects an $\varepsilon{\cdot}2^{-(len(\rho)+1)}$ optimal action. For an *initial state* $s \in S$ and a finite play ρ such that $last(\rho){=}s$, it follows that:

$$V(s) + \varepsilon{\cdot}2^{-(len(\rho)+1)} \geqslant \sup_{\beta \in A_{\mathrm{Max}}(s)} \left\{ \tau(s, \mu(\rho), \beta) + \sum_{s' \in S} p(\mu(\rho), \beta)(s') \cdot V(s') \right\}. \tag{1}$$

We will now show that for any path ρ, counter-strategy χ for Max and $n \in \mathbb{N}$ we have:

$$\mathbb{E}_{\rho}^{\mu,\chi}(\mathsf{Reach}_F^n) \leqslant V(last(\rho)) + \textstyle\sum_{m=len(\rho)+1}^{n+len(\rho)} \varepsilon{\cdot}2^{-m}\,. \tag{2}$$

We prove (2) by induction on $n \in \mathbb{N}$. The case for $n{=}0$ follows from Definition 4 and Definition 8.

Now suppose (2) holds for some $n \in \mathbb{N}$. Consider any finite path ρ where $last(\rho) = s$ and counter-strategy χ for Max. Now, if $s \in F$, then by Definition 4 we have:

$$\mathbb{E}_{\rho}^{\mu,\chi}(\mathsf{Reach}_F^{n+1}) = 1 = V(last(\rho)) \leqslant V(last(\rho)) + \textstyle\sum_{m=len(\rho)+1}^{(n+1)+len(\rho)} \varepsilon{\cdot}2^{-m}\,.$$

On the other hand, if $s \notin F$ and letting $a{=}\mathsf{win}(\mu(\rho),\chi(\rho))$, then by Definition 4 and Definition 8:

$$\begin{aligned}
&\mathbb{E}_{\rho}^{\mu,\chi}(\mathsf{Reach}_F^{n+1}) = \tau(s,\mu(\rho),\chi(\rho)) + \sum_{s'\in S} p(s,\mu(\rho),\chi(\rho))(s'){\cdot}\mathbb{E}_{\rho as'}^{\mu,\chi}(\mathsf{Reach}_F^{n}) \\
&\leqslant\ \tau(s,\mu(\rho),\chi(\rho)) + \sum_{s'\in S} p(s,\mu(\rho),\chi(\rho))(s'){\cdot}\left(V(s') + \sum_{m=len(\rho as')+1}^{n+len(\rho as')} \varepsilon{\cdot}2^{-m}\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by induction)} \\
&\leqslant\ \tau(s,\mu(\rho),\chi(\rho)) + \sum_{s'\in S} p(s,\mu(\rho),\chi(\rho))(s'){\cdot}\left(V(s') + \sum_{m=len(\rho)+2}^{n+len(\rho)+1} \varepsilon{\cdot}2^{-m}\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by definition of } len(\cdot)) \\
&=\ \left(\tau(s,\mu(\rho),\chi(\rho)) + \sum_{s'\in S} p(s,\mu(\rho),\chi(\rho))(s'){\cdot}V(s')\right) + \sum_{m=len(\rho)+2}^{(n+1)+len(\rho)} \varepsilon{\cdot}2^{-m} \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(rearranging)} \\
&\leqslant\ V(s) + \varepsilon{\cdot}2^{-(len(\rho)+1)} + \sum_{m=len(\rho)+2}^{(n+1)+len(\rho)} \varepsilon{\cdot}2^{-m} \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(by (1))} \\
&=\ V(last(\rho)) + \sum_{m=len(\rho)+1}^{(n+1)+len(\rho)} \varepsilon{\cdot}2^{-m} \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{(rearranging.)}
\end{aligned}$$

Since these are all the cases to consider, it follows that (2) holds by induction on n.

Letting $\rho = s$ and taking the limit of n in (2), we have $\mathbb{E}_{s}^{\mu,\chi}(\mathsf{Reach}_F) \leqslant V(s) + \varepsilon$ and, since ε and χ were arbitrary, it follows that $\overline{\mathsf{Val}}(s) \leqslant V(s)$ as required. □

Lemma 10. $\overline{\mathsf{Val}} \geqslant \lim_{n\to\infty} \overline{\mathsf{Val}}^n$.

PROOF. The proof follows straightforwardly from the fact that for any $n \in \mathbb{N}$ and finite play ρ we have that $\mathsf{Reach}_F(\rho) \geqslant \mathsf{Reach}_F^n(\rho)$. □

3. Probabilistic Timed Game Arenas

In this section we introduce Probabilistic Timed Game Arenas (PTGAs) which extend classical timed automata [2] with discrete distributions and a partition of the actions between two players Min and Max. However, before we present syntax and semantics of PTGAs, we need to introduce the concept of clock variables and related notions.

3.1. Clocks, Constraints, Regions, and Zones

Clocks. Let $\mathcal{C}$ be a finite set of *clocks*. A *clock valuation* on $\mathcal{C}$ is a function $\nu : \mathcal{C} \to \mathbb{R}_{\geqslant 0}$ and we write $V(\mathcal{C})$ (or just V when $\mathcal{C}$ is clear from the context) for the set of clock valuations. Abusing notation, we also treat a valuation ν as a point in $(\mathbb{R}_{\geqslant 0})^{|\mathcal{C}|}$. Let $\mathbf{0}$ denote the clock valuation that assigns 0 to all clocks. If $\nu \in V$ and $t \in \mathbb{R}_{\geqslant 0}$ then we write $\nu{+}t$ for the clock valuation defined by $(\nu{+}t)(c) \stackrel{\text{def}}{=} \nu(c){+}t$ for all $c \in \mathcal{C}$. For $C \subseteq \mathcal{C}$, we write ν_C for the valuation where $\nu_C(c)$ equals 0 if $c \in C$ and $\nu(c)$ otherwise. For $X \subseteq V(\mathcal{C})$, we write $\overline{X}$ for the smallest closed set in V containing X. Although clocks are usually allowed to take arbitrary non-negative values, for notational convenience we assume that there is an upper bound $K \in \mathbb{N}$ such that for every clock $c \in \mathcal{C}$ we have that $\nu(c) \leqslant K$.

Clock constraints. A *clock constraint* over $\mathcal{C}$ with upper bound $K \in \mathbb{N}$ is a conjunction of *simple constraints* of the form $c \bowtie i$ or $c{-}c' \bowtie i$, where $c, c' \in \mathcal{C}$, $i \in \mathbb{N}$, $i{\leqslant}K$, and $\bowtie \in \{<, >, =, \leqslant, \geqslant\}$. For $\nu \in V(\mathcal{C})$ and $K \in \mathbb{N}$, let $\mathrm{SCC}(\nu, K)$ be the set of clock constraints with upper bound K which hold in ν, i.e. those constraints that resolve to `true` after substituting each occurrence of a clock x with $\nu(x)$.

Clock regions. Every clock region is an equivalence class of the indistinguishability-by-clock-constraints relation, and vice versa. For a given set of clocks $\mathcal{C}$ and upper bound $K \in \mathbb{N}$ on clock constraints, a *clock region* is a maximal set $\zeta{\subseteq}V(\mathcal{C})$ such that $\mathrm{SCC}(\nu, K){=}\mathrm{SCC}(\nu', K)$ for all $\nu, \nu' \in \zeta$. For the set of clocks $\mathcal{C}$ and upper bound K we write $\mathcal{R}(\mathcal{C}, K)$ for the corresponding finite set of clock regions. We write $[\nu]$ for the clock region of ν. If $\zeta{=}[\nu]$, write ζ_C for $[\nu_C]$; this definition is well-defined, since for any clock valuations ν and ν' if $[\nu]{=}[\nu']$ then $[\nu_C]{=}[\nu'_C]$.

Clock zones. A *clock zone* is a convex set of clock valuations, which is a union of a set of clock regions. We write $\mathcal{Z}(\mathcal{C}, K)$ for the set of clock zones over the set of clocks $\mathcal{C}$ and upper bound K. Observe that a set of clock valuations is a clock zone if and only if it is definable by a clock constraint. Although more than one clock constraint can represent the same zone, for any clock zone ζ, there exists an $O(|\mathcal{C}|^3)$ algorithm to compute the (unique) canonical clock constraint of ζ [38]. We therefore interchange the semantic and syntactic interpretation of clock zones.

When the set of clocks and upper bound is clear from the context we write $\mathcal{R}$ and $\mathcal{Z}$ for the set of regions and zones respectively.

3.2. Probabilistic Timed Game Arena: Syntax

For the remainder of the paper we fix a positive integer K, and work with K-bounded clocks and clock constraints.

Definition 11 (Probabilistic Timed Game Arena (PTGA)). *A probabilistic timed game arena is a tuple* $\mathsf{T}{=}(L, \mathcal{C}, \mathit{Inv}, \mathit{Act}_{\mathrm{Min}}, \mathit{Act}_{\mathrm{Max}}, E, \delta)$ *where*

- L *is a finite set of* locations*;*
- $\mathcal{C}$ *is a finite set of* clocks*;*

- $Inv : L \to \mathcal{Z}$ *is an* invariant condition*;*
- Act_{Min} *and* Act_{Max} *are disjoint finite sets of* actions*, and we use* Act *for the set* $Act_{\mathrm{Min}} \cup Act_{\mathrm{Max}}$
- $E : L \times Act \to \mathcal{Z}$ *is an* action enabling condition*;*
- $\delta : L \times Act \to \mathcal{D}(2^{\mathcal{C}} \times L)$ *is a* probabilistic transition function*.*

When we consider a PTGA 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 [28], we assume that probabilities are expressed as ratios of two natural numbers, each written in binary, and zones in the definition of Inv and E are expressed as clock constraints.

A standard *probabilistic timed automaton* (PTA) is a PTGA where one of Act_{Min} and Act_{Max} is empty. On the other hand, the standard (non-probabilistic) *timed game arena* (*timed automaton*) is a PTGA (PTA) such that $\delta(\ell, a)$ is a point distribution for all $\ell \in L$ and $a \in Act$.

3.3. Probabilistic Timed Game Arena: Semantics

Let $\mathsf{T}{=}(L, \mathcal{C}, Inv, Act_{\mathrm{Min}}, Act_{\mathrm{Max}}, E, \delta)$ be a probabilistic timed game arena. A *configuration* of a PTGA is a pair (ℓ, ν), where ℓ is a location and ν a clock valuation such that $\nu \in Inv(\ell)$. For any $t \in \mathbb{R}_{\geqslant 0}$, 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'.

A game on a PTGA starts in an *initial configuration* $(\ell, \nu) \in L \times V$ and Min and Max construct an infinite play by repeatedly choosing available timed actions $(t_a, a) \in \mathbb{R}_{\geqslant 0} \times Act_{\mathrm{Min}}$ and $(t_b, b) \in \mathbb{R}_{\geqslant 0} \times Act_{\mathrm{Max}}$ proposing $\bot$ 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 $\bot$, and otherwise Max is responsible. We assume the players cannot simultaneously choose $\bot$, i.e. that in any configuration there is at least one timed action available.

Definition 12 (PTGA Semantics). *Let* $\mathsf{T} = (L, \mathcal{C}, Inv, Act_{\mathrm{Min}}, Act_{\mathrm{Max}}, E, \delta)$ *be a PTGA. The semantics of* T *is given by the SGA*

$$[\![\mathsf{T}]\!]{=}(S, A_{\mathrm{Min}}, A_{\mathrm{Max}}, p_{\mathrm{Min}}, p_{\mathrm{Max}}, \mathsf{win}, \tau_{\mathrm{Min}}, \tau_{\mathrm{Max}})$$

where

- $S \subseteq L \times V$ *is the (possibly uncountable) set of states such that* $(\ell, \nu) \in S$ *if and only if* $\nu \in Inv(\ell)$*;*
- $A_{\mathrm{Min}} = (\mathbb{R}_{\geqslant 0} \times Act_{\mathrm{Min}}) \cup \{\bot\}$ *and* $A_{\mathrm{Max}} = (\mathbb{R}_{\geqslant 0} \times Act_{\mathrm{Max}}) \cup \{\bot\}$ *are the sets of timed actions of players* Min *and* Max*;*

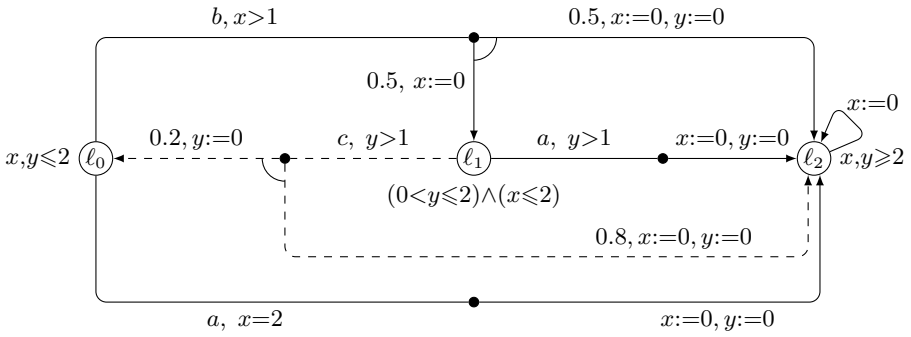

Figure 1: Example of a probabilistic timed game arena.

- *for* $\star \in \{\mathrm{Min}, \mathrm{Max}\}$, $(\ell, \nu) \in S$ *and* $(t, a) \in A_\star$ *the probabilistic transition function* $p_\star$ *is defined when* $\nu{+}t' \in \mathit{Inv}(\ell)$ *for all* $0{\leqslant}t'{\leqslant}t$, $\nu{+}t \in E(\ell, a)$ *and for any* (ℓ', ν')*:*

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

- *for* $(t_a, a) \in \mathbb{R}_{\geqslant 0} {\times} \mathit{Act}_{\mathrm{Min}}$ *and* $(t_b, b) \in \mathbb{R}_{\geqslant 0} {\times} \mathit{Act}_{\mathrm{Max}}$*, we define*

$$\mathsf{win}((t_a, a), (t_b, b)) = \begin{cases} (t_a, a) & \text{if } t_a {<} t_b \\ (t_b, b) & \text{otherwise.} \end{cases}$$

  If one of the arguments to win *is* $\bot$*, we define the returning value to be the other argument.*

- *the time delay function is given by* $\tau_\star(s, (t, a)) = t$ *for all* $\star \in \{\mathrm{Min}, \mathrm{Max}\}$, $s \in S$ *and* $(t, a) \in A_\star$ *such that* $p_\star(s, (t, a))$ *is defined.*

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 time delay function of the SGA corresponds to the elapsed time of each move.

Time Divergence. When modelling real-time systems it is important to restrict attention to time divergent (or non-Zeno) behaviour. More precisely, one should not consider strategies which lead to behaviour in which time does not advance beyond a certain point, as this cannot occur in a real system. We achieve this by restricting attention to *structurally non-Zeno* PGTAs, these are PGTA where all strategies will yield time-divergent behaviour by construction. We use the syntactic conditions given in [39] for PTAs and are derived from those for timed automata [40, 41].

Example 13. *Consider the PTGA in Figure 1; we use solid and dashed lines to indicate actions controlled by* Min *and* Max *respectively. Considering location* ℓ_1*, the invariant condition is* $(0{<}y{\leqslant}2){\wedge}(x{\leqslant}2)$*, actions* a *and* c *are enabled when* $y{>}1$ *and,*

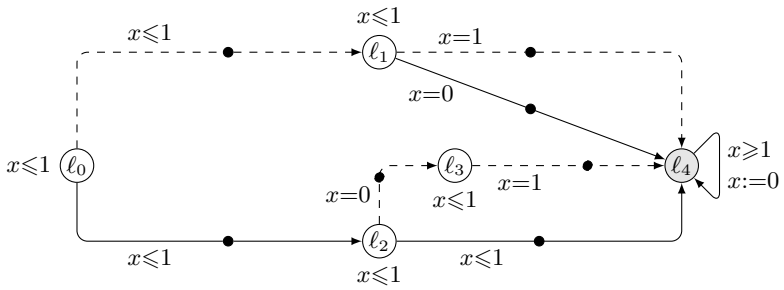

Figure 2: Example demonstrating that PTGAs are not determined. Action names are omitted for brevity.

if a is taken, we move to ℓ_2, 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_2.

Let us denote clock valuations by tuples where the first (second) coordinate correspond to the clock x (y). Starting in the configuration $(\ell_0, (0, 0))$ 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*

$$\langle(\ell_0,(0,0)),\ ((1.1,b),\bot),\ (\ell_1,(0,1.1)),\ ((0.5,a),\ (0.2,c)),\ (\ell_0,(0.2,0)),\ ((1.1,b),\bot),\ (\ell_2,(0,0))\rangle$$

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 location ℓ_2.

3.4. Reachability-time problem over PTGA

We are interested in the reachability-time problem for games over the semantics of a PTGA T. We assume that the target set is given as a set L_F of locations (the corresponding target of the SGA $[\![\mathsf{T}]\!]$, with state space S, is given by $F{=}\{(\ell,\nu) \in S \mid \ell \in L_F\}$). However, the results presented can be easily generalised to target sets of location-zone pairs.

3.5. Non-determinacy of PTGA with reachability-time objectives

Before proceeding with the definitions that we need to prove the main decidability result of the paper, we show, through the following counter-example, that PTGAs are not determined, even when the game contains only non-strict inequalities.

Example 14. *Considering the PTGA given in Figure 2 with target set $L_F{=}\{\ell_4\}$, recall that we use solid and dashed lines to indicate actions controlled by* Min *and* Max *respectively. Constructing the optimality equations $\overline{\mathsf{Opt}}_{\mathsf{G}}$ for the SGA semantics of this PTGA, we have, after some simplifications:*

$$P(\ell_4, x) = 0$$
$$P(\ell_3, x) = 1{-}x$$

$$P(\ell_2, x) = \begin{cases} 1 & \textit{if } x{=}0 \\ 0 & \textit{otherwise} \end{cases}$$
$$P(\ell_1, x) = \begin{cases} 0 & \textit{if } x{=}0 \\ 1{-}x & \textit{otherwise} \end{cases}$$

and $P(\ell_0, 0)$ *is equal to the minimum of:*

$$\max\{0{+}P(\ell_2, 0), 0{+}P(\ell_1, 0)\} \tag{3}$$

and

$$\inf_{0<t\leqslant 1}\left\{\max\left\{\sup_{t<t'\leqslant 1}\big(t{+}P(\ell_2,t)\big), \sup_{0\leqslant t'\leqslant t}\big(t'{+}P(\ell_1,t')\big)\right\}\right\}. \tag{4}$$

The expression (3) is equal to 1 *and corresponds to player* Min *leaving* ℓ_0 *immediately (when the clock* x *equals 0). The expression (4) corresponds to the infimum over leaving* ℓ_0 *after a non-zero delay (when the clock* x *is greater than 0) and is also equal to* 1*. Combining these results we have that* $P(\ell_0, 0){=}1$.

On the other hand, considering the optimality equations $\mathsf{Opt}_{\mathsf{G}}$*, the values for the locations* $\ell_1, \ldots, \ell_4$ *are as above, while the value for* $P(\ell_0, 0)$ *equals the maximum of:*

$$0{+}P(\ell_1, 0) \;\; \textit{and} \;\; \sup_{0<t\leqslant 1}\left\{\min\left\{\inf_{t\leqslant t'\leqslant 1}\big(t{+}P(\ell_1,t)\big), \inf_{0\leqslant t'<t}\big(t'{+}P(\ell_2,t')\big)\right\}\right\}. \tag{5}$$

The first expression in (5) equals 0 *and corresponds to player* Max *leaving* ℓ_0 *immediately. The second expression in (5) corresponds to the supremum over leaving* ℓ_0 *after a non-zero delay, and is also equal to* 0*, and therefore it follows that* $P(\ell_0, \mathbf{0}){=}0$*. Hence the game is not determined as the upper and lower values of the game differ in the state* $(\ell_0, 0)$.

4. Boundary region abstraction

The region graph [2] is useful for solving time-abstract optimisation problems on timed automata. The region graph, however, is not suitable for solving competitive optimisation problems and games on timed automata as it abstracts away the timing information. The corner-point abstraction [42], which captures digital clock semantics [43] of timed automata, is an abstraction of timed automata where the configurations of the system are restricted to $L{\times}\mathbb{N}^{|\mathcal{C}|}$, i.e. transitions are allowed only when all clocks have non-negative integer values. Although this abstraction retains some timing information, it is not convenient for proof techniques based on dynamic programming, used in this paper. The boundary region abstraction (BRA) [14], a generalisation of the corner-point abstraction, is better suited for such proof techniques. More precisely, we need to prove certain properties of values in a PTGA, which we can do only when reasoning about all the states of the PTGA. 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 [14] to handle PTGAs. First, we require a number of preliminary concepts.

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 $\zeta \to^* \zeta'$, if there exist $\nu \in \zeta$, $\nu' \in \zeta'$ and $t \in \mathbb{R}_{\geqslant 0}$ such that $\nu' = \nu{+}t$ and say ζ' is the *time successor* of ζ if $\zeta \neq \zeta'$ and $\nu{+}t' \in \zeta \cup \zeta'$ for all $t' \leqslant t$ and write $\zeta \to \zeta'$ to denote this fact. We also use $\zeta \to^+ \zeta'$ if there is ζ'' such that $\zeta \to \zeta'' \to^* \zeta'$. For regions $\zeta, \zeta' \in \mathcal{R}$ such that $\zeta \to^* \zeta'$ we write $[\zeta, \zeta']$ for the zone $\cup\{\zeta'' \mid \zeta \to^* \zeta'' \wedge \zeta'' \to^* \zeta'\}$.

Intuition for the Boundary Region Abstraction. In our definition of the boundary region abstraction (BRA) we capture the intuition that, when studying the "optimal" behaviour of the players, it is sufficient to consider moves that take place near the start or end of the regions. This allows us to abstract from moves that specify the precise time, but instead allow the players to say which regions they wish to enter, and then either say that they want to take the move at the start of the region ($\inf$), or at its end ($\sup$).

Based on this intuition we define the boundary region abstraction of a probabilistic game arena as follows.

Definition 15 (Boundary region abstraction (BRA)). *For a probabilistic timed game arena $\mathsf{T}{=}(L, \mathcal{C}, Inv, Act_{\mathrm{Min}}, Act_{\mathrm{Max}}, E, \delta)$, the boundary region abstraction of T is given by the SGA $\widehat{\mathsf{T}}{=}(\widehat{S}, \widehat{A}_{\mathrm{Min}}, \widehat{A}_{\mathrm{Max}}, \widehat{p}_{\mathrm{Min}}, \widehat{p}_{\mathrm{Max}}, \widehat{\mathsf{win}}, \widehat{\tau}_{\mathrm{Min}}, \widehat{\tau}_{\mathrm{Max}})$ 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 $\zeta \subseteq Inv(\ell)$ and $\nu \in \overline{\zeta}$ (recall that $\overline{\zeta}$ denotes the closure of ζ);*
- *$\widehat{A}_{\mathrm{Min}} = (Act_{\mathrm{Min}}{\times}\mathcal{R}{\times}\{\inf, \sup\}) \cup \{\bot\}$ is the set of actions of player Min;*
- *$\widehat{A}_{\mathrm{Max}} = (Act_{\mathrm{Max}}{\times}\mathcal{R}{\times}\{\inf, \sup\}) \cup \{\bot\}$ is the set of actions of player Max;*
- *for $\star \in \{\mathrm{Min}, \mathrm{Max}\}$, $\hat{s} = (\ell, \nu, \zeta) \in \widehat{S}$ and $\alpha = (a, \zeta'', \mathrm{opt}) \in \widehat{A}_\star$ such that $\zeta \to^* \zeta''$, the probabilistic transition function $\widehat{p}_\star$ is defined if $[\zeta, \zeta''] \subseteq Inv(\ell)$ and $\zeta'' \subseteq E(\ell, a)$ and for any $(\ell', \nu', \zeta') \in \widehat{S}$:*

  $$\widehat{p}_\star(\hat{s}, \alpha)((\ell', \nu', \zeta')) = \textstyle\sum_{C \subseteq \mathcal{C} \wedge \nu''_C = \nu' \wedge \zeta''_C = \zeta'} \delta[\ell, a](C, \ell')$$

  where $\nu'' = \mathrm{opt}_{\nu+t \in \zeta'', t \geqslant 0}\, \nu{+}t$;
- *$\widehat{\mathsf{win}}((a, \zeta_a, \mathrm{opt}_a), (b, \zeta_b, \mathrm{opt}_b))$ is equal to $(a, \zeta_a, \mathrm{opt}_a)$ if (i) $\zeta_a \to^+ \zeta_b$ or (ii) $\zeta_a = \zeta_b$, $\mathrm{opt}_a = \inf$ and $\mathrm{opt}_b = \sup$; it is equal to $(b, \zeta_b, \mathrm{opt}_b)$ otherwise;*
- *for $\star \in \{\mathrm{Min}, \mathrm{Max}\}$, $(\ell, \nu, \zeta) \in \widehat{S}$ and $(a_\alpha, \zeta_\alpha, \mathrm{opt}) \in \widehat{A}_\star$ such that $\widehat{p}_\star$ is defined the time delay function is given by $\widehat{\tau}_\star((\ell, \nu, \zeta), (a_\alpha, \zeta_\alpha, \mathrm{opt})) = \mathrm{opt}_{\nu+t \in \zeta_\alpha}\, t$.*

Given a target set of locations L_F of T, the corresponding target set of the BRA is given by $\widehat{F}{=}\{(\ell, \nu, \zeta) \in \widehat{S} \mid \ell \in L_F\}$.

To simplify notation, for two elements $a \in \widehat{A}_{\mathrm{Min}}$ and $b \in \widehat{A}_{\mathrm{Max}}$ we write $a \leqslant b$ to denote that $\widehat{\mathsf{win}}(a, b){=}a$. We use analogous notation also for other SGAs. For an element $s{=}(\ell, \nu) \in L{\times}V$, we use $\widehat{s}$ to denote the element $(\ell, \nu, [\nu]) \in \widehat{S}$.

Although the boundary region abstraction is not a finite SGA, for a fixed initial state we can restrict attention to a finite SGA, adapting an approach from [44] as follows.

Proposition 16. *Let* T *be a PTGA and* $\widehat{\mathsf{T}}$ *the corresponding BRA. For any state of* $\widehat{\mathsf{T}}$*, its reachable sub-graph is finite and constructible in time exponential in the size of* T*.*

PROOF. The most demanding part of the proof is to show that there is a set V of valuations that has exponential size and contains ν for any state (ℓ, ν, ζ) reachable in the sub-graph of $\widehat{\mathsf{T}}$.

For $r \in \mathbb{R}_{\geqslant 0}$ we write $\langle r \rangle$ for the fractional part of r, i.e. $r - \lfloor r \rfloor$. For a clock valuation ν we define its fractional signature $\wr \nu \wr$ to be the sequence $(f_0, f_1, \ldots, f_m)$ such that $f_0{=}0$, $f_i{<}f_j$ if $i{<}j$, for all $i, j \leqslant m$, and $f_1, f_2, \ldots, f_m$ are all the non-zero fractional parts of clock values in the clock valuation ν. In other words, for every $i{\geqslant}1$ there is a clock c such that $\langle \nu(c) \rangle{=}f_i$, and for every clock $c \in \mathcal{C}$ there is $i{\leqslant}m$ such that $\langle \nu(c) \rangle{=}f_i$.

Let $\oplus$ denote addition modulo m. For $0{\leqslant}k{\leqslant}m$ we define the k-shift of a fractional signature $(f_0, f_1, \ldots, f_m)$ as the fractional signature $(f'_0, f'_1, \ldots, f'_m)$ such that for all $0 \leqslant i \leqslant m$ we have $f'_i = \langle f_{i \oplus k} + 1 - f_k \rangle$. Note that a k''-shift $(f''_0, \ldots, f''_m)$ of a k'-shift $(f'_0, \ldots, f'_m)$ of $(f_0, f_1, \ldots, f_m)$ is an $(k' \oplus k'')$-shift of $(f_0, f_1, \ldots, f_m)$ because for any i we have:

$$\begin{aligned} f''_i &= \langle f'_{i \oplus k''} {+} 1 {-} f'_{k''} \rangle \\ &= \langle \langle f_{i \oplus k'' \oplus k'} {+} 1 {-} f_{k'} \rangle {+} 1 {-} \langle f_{k'' \oplus k'} {+} 1 {-} f_{k'} \rangle \rangle \\ &= \langle f_{i \oplus (k'' \oplus k')} {+} 1 {-} f_{k'' \oplus k'} \rangle \,. \end{aligned}$$

This means that, by successive application of shifts, only m different fractional signatures can be obtained. We further say that a fractional signature $(f'_0, f'_1, \ldots, f'_n)$ is a subsequence of another fractional signature $(f_0, f_1, \ldots, f_m)$ if $n{\leqslant}m$ and for all $i{\leqslant}n$ there exists $j{\leqslant}m$ such that $f'_i{=}f_j$.

For any state (ℓ, ν, ζ) of the BRA $\widehat{\mathsf{T}}$, we claim that it is only possible to transition to states (ℓ', ν', ζ') such that $\langle \nu' \rangle$ is a subsequence of a k-shift of $\langle \nu \rangle$, for some k. To see that, notice that the ν_α in the definition of $\widehat{p}_\star$ (Definition 15) satisfies that $\langle \nu_\alpha \rangle$ is a k-shift of $\langle \nu \rangle = (f_0, \ldots f_m)$ for k chosen so that f_m is the fractional part of clocks that have integer value in ν_α. Subsequently resetting clocks gives rise to a subsequence of a fractional signature, and so $\langle \nu' \rangle$ (for ν' from the defining sum of $\widehat{p}_\star$) is a subsequence of $\langle \nu_\alpha \rangle$. □

Example 17. *Returning to Example 13 (see Figure 1), a sub-graph of BRA reachable from* $(\ell_0, (0.3, 0.1), 0{<}y{<}x{<}1)$ *for the PTGA of Figure 1 is shown in Figure 3. The names of the regions correspond to the regions depicted in the bottom right corner. Edges are labelled* (a, ζ, opt) *and the intuitive meaning is to wait until we reach the lower or upper (depending on* opt*) boundary of the region. For some regions, for example* ζ_4*, the boundaries coincide and we keep this redundancy to simplify the notation. Considering the region* ζ_1*, we see that it is determined by the constraints* $(1{<}x{<}2){\wedge}(0{<}y{<}1){\wedge}(y{<}x{-}1)$*. The bold numbers on edges correspond to the time delay before the action labelling the edge is taken. Figure 3 includes the actions available in the initial state and one of the action pairs that are available in the state* $(\ell_1, (0, 1), (x{=}0){\wedge}(1{<}y{<}2))$*. To simplify the figure, the probabilities that are equal to* 0.5 *are omitted.*

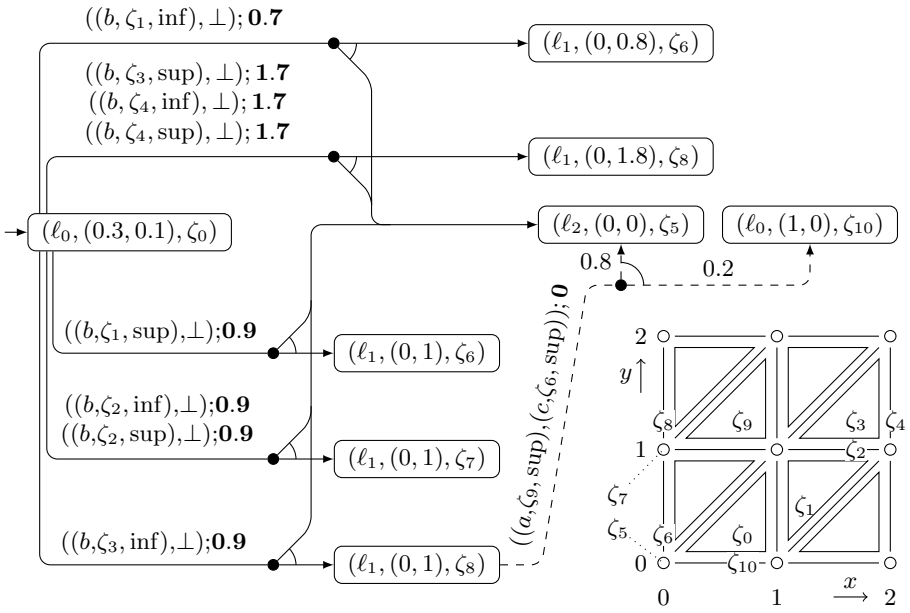

Figure 3: Sub-graph of the boundary region abstraction for the PTGA of Figure 1, with the region names as depicted in the bottom right corner.

5. Decidability of the Reachability-Time Problem

In this section we show decidability of the reachability-time problem, which is the main result of the paper. The result is formalised in the following theorem.

Theorem 18. *Let* T *be a PTGA. The reachability-time problem for infinite-horizon objectives in* T *is in NEXPTIME∩co-NEXPTIME.*

The crucial, and most demanding, step of the proof of Theorem 18 is proving that the problems on PTGAs can be reduced to problems on BRAs. This fact is formalised in Theorem 19. Theorem 18 then follows straightforwardly from Theorem 19, Proposition 16 and Theorem 6.

Theorem 19. *Let* T *be a PTGA and* $\widehat{\mathsf{T}}$ *the corresponding BRA. The answers to the reachability-time problems for* T *and* $\widehat{\mathsf{T}}$ *are the same.*

The remainder of this paper is devoted to the proof of Theorem 19. First, in Section 5.1 we introduce quasi-simple functions and prove some of their properties. Then, in Section 5.2 we show that values in the games we study can be characterised using quasi-simple functions, and that this allows us to establish the correspondence between PTGA and its boundary region abstraction.

For the remainder of this section, unless otherwise specified, we fix a PTGA $\mathsf{T} = (L, \mathcal{C}, Inv, Act_{\min}, Act_{\max}, E, \delta)$, set of target locations F_L, suppose the semantics of

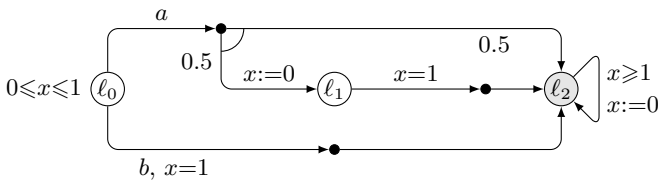

Figure 4: Example demonstrating optimal strategies are not regionally positional.

T is given by:

$$[\![\mathsf{T}]\!] = (S, A_{\mathrm{Min}}, A_{\mathrm{Max}}, p_{\mathrm{Min}}, p_{\mathrm{Max}}, \mathsf{win}, \tau_{\mathrm{Min}}, \tau_{\mathrm{Max}})$$

with corresponding target set $F{=}\{(\ell,\nu) \in S \mid \ell \in F_L\}$ and the boundary region abstraction of T is given by

$$\widehat{T} = (\widehat{S}, \widehat{A}_{\mathrm{Min}}, \widehat{A}_{\mathrm{Max}}, \widehat{p}_{\mathrm{Min}}, \widehat{p}_{\mathrm{Max}}, \widehat{\mathsf{win}}, \widehat{\tau}_{\mathrm{Min}}, \widehat{\tau}_{\mathrm{Max}})$$

with corresponding target set $\widehat{F}{=}\{(\ell,\nu,\zeta) \in \widehat{S} \mid \ell \in F_L\}$.

5.1. Quasi-simple Functions

To prove properties of controllers for (non-probabilistic) timed systems, Asarin and Maler [3] introduced simple functions, a finitely representable class of functions with the property that every decreasing sequence is finite. We define these functions here and show that they are not sufficient for our purpose.

Definition 20 (Simple Functions). *Given a set of valuations $X{\subseteq}V$, a function $f : X{\to}\mathbb{R}_{\geqslant 0}$ is* simple *if there exists $e \in \mathbb{N}$ and either $f(\nu){=}e$ for all $\nu \in X$, or there exists a clock $c \in C$ such that $f(\nu){=}e{-}\nu(c)$ for all $\nu \in X$. Furthermore, a function $f : \widehat{S}{\to}\mathbb{R}_{\geqslant 0}$ 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 upper values for n-step reachability-time objectives are regionally simple, and because the fixpoint is reached for some n the upper value for reachability-time objective is regionally simple. Also, using the properties of simple functions, [14] shows that, for a non-probabilistic game reachability-time objectives, 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 PTGAs, applying the value improvement function does not necessarily preserve regional-simplicity. Moreover, as the example below demonstrates, neither is the value of the game necessarily regionally simple nor optimal strategies regionally positional.

Example 21. *Consider the one-player PTGA shown in Figure 4. Observe that, for every state (ℓ_0,ν) in the region $(\ell_0, 0{<}x{<}1)$, the optimal expected time to reach ℓ_2 equals*

$$\min\{\inf_{t\geqslant 0}\{t + 0.5{\cdot}1 + 0.5{\cdot}0\}, 1{-}\nu(x)\} = \min\{0.5, 1{-}\nu(x)\}.$$

Hence, the values in PTGA with reachability-time objectives may not be regionally simple. Moreover, the optimal strategy is not regionally positional, since if $\nu(x) \leqslant 0.5$ then the optimal strategy is to take action a immediately, while otherwise the optimal strategy is to wait until $\nu(x)=1$ and then take action b.

Due to these results it is not possible to work with simple functions. Our proof instead relies on regional non-expansiveness of value functions. Given $X \subseteq V$, a function $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$ is *non-expansive* if for all $x, y \in X$ we have $|f(x)-f(y)| \leqslant |x-y|$. A function $f : \widehat{S} \to \mathbb{R}^{\infty}_{\geqslant 0}$ is *regionally non-expansive* if $f(\ell, \cdot, \zeta)$ is non-expansive for any ℓ and ζ, and similarly any $f : S \to \mathbb{R}^{\infty}_{\geqslant 0}$ is regionally non-expansive if $f(\ell, \cdot)$ is non-expansive when its domain is restricted to a single region.

The proof direction that we take requires us to establish that $\lim_{n\to\infty} \overline{\mathsf{Val}}^n$ is non-expansive. To do this, we will show that for each $n \in \mathbb{N}$ the function $\overline{\mathsf{Val}}^n$ is non-expansive. However, a direct proof by induction would fail and instead we are required to prove a stronger claim about the functions $\overline{\mathsf{Val}}^n$. To do this, we first introduce *quasi-simple functions*.

Definition 22 (Quasi-Simple Functions). *Let $X \subseteq V$ be a set of clock valuations. The class of quasi-simple functions is built by first defining every simple function to be quasi-simple, and then inductively by stipulating that convex combination, maximum and minimum of finitely many quasi-simple functions are quasi-simple.*

A function $f : \widehat{S} \to \mathbb{R}^{\infty}_{\geqslant 0}$ is regionally quasi-simple if $f(\ell, \cdot, \zeta)$ is quasi-simple for all $\ell \in L$ and $\zeta \in \mathcal{R}$, and any $f : S \to \mathbb{R}^{\infty}_{\geqslant 0}$ is regionally quasi-simple if $f(\ell, \cdot)$ is quasi-simple when its domain is restricted to a single region.

We will later show that functions $\overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]} : S \to \mathbb{R}_{\geqslant 0}$ and $\overline{\mathsf{Val}}^n_{\widehat{T}} : \widehat{S} \to \mathbb{R}_{\geqslant 0}$ for $n \in \mathbb{N}$ are regionally quasi-simple. From this using the lemma below we can then demonstrate that these functions are non-expansive.

Lemma 23. *Every quasi-simple function is non-expansive.*

PROOF. Consider any quasi-simple function $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$. We will prove by induction on the structure of f (see Definition 22) that for any $\nu_1, \nu_2 \in X$ we have $|f(\nu_1) - f(\nu_2)| \leqslant |\nu_1 - \nu_2|$.

- If f is a simple function, then either f is a constant, and hence:

  $$|f(\nu_1) - f(\nu_2)| = 0 \leqslant |\nu_1 - \nu_2|$$

  or $f = e - \nu(c)$ for some clock c, in which case:

  $$|f(\nu_1) - f(\nu_2)| = |\nu_2(c) - \nu_1(c)| \leqslant |\nu_2 - \nu_1|$$

  as required.

- If f is a convex combination $p_1, \ldots, p_n$ of quasi-simple functions $f_1, \ldots, f_n$, then:

  $$|\textstyle\sum_{i=1}^{n} p_i \cdot f_i(\nu_1) - \sum_{i=1}^{n} p_i \cdot f_i(\nu_2)| \leqslant |\sum_{i=1}^{n} p_i \cdot (f_i(\nu_1) - f_i(\nu_2))|$$

$\leqslant \sum_{i=1}^{n} p_i \cdot |\nu_1 - \nu_2|$ (by induction)
$= |\nu_1 - \nu_2|$ (since we are considering a convex combination)

as required.

- If f is the maximum of two quasi-simple functions f_1 and f_2, then without loss of generality we suppose $f_1(\nu_1) \geqslant f_2(\nu_1)$. In the case when $f_1(\nu_2) \geqslant f_2(\nu_2)$ we have

  $$|\max\{f_1(\nu_1), f_2(\nu_1)\} - \max\{f_1(\nu_2), f_2(\nu_2)\}| \leqslant |f_1(\nu_1) - f_1(\nu_2)| \leqslant |\nu_1 - \nu_2|$$

  since f_1 is non-expansive. On the other hand, in the case when $f_1(\nu_2) < f_2(\nu_2)$:

  $$|\max\{f_1(\nu_1), f_2(\nu_1)\} - \max\{f_1(\nu_2), f_2(\nu_2)\}| \leqslant |f_1(\nu_1) - f_2(\nu_2)| \ .$$

  Now either $f_1(\nu_1) \geqslant f_2(\nu_2)$, and therefore we have:

  $$|f_1(\nu_1) - f_2(\nu_2)| \leqslant |f_1(\nu_1) - f_1(\nu_2)| \leqslant |\nu_1 - \nu_2|$$

  since f_1 is non-expansive, or $f_1(\nu_1) < f_2(\nu_2)$ in which case:

  $$|f_1(\nu_1) - f_2(\nu_2)| \leqslant |f_2(\nu_1) - f_2(\nu_2)| \leqslant |\nu_1 - \nu_2|$$

  since f_2 is non-expansive. Since these are all the cases to consider we have f is non-expansive as required.

- If f is minimum of two quasi-simple functions the proof follows similarly to the case when f is the maximum of two quasi-simple functions.

Since these are the only cases to consider the proof is complete. □

In the proofs below we will make use of several technical properties of quasi-simple functions. First, however, we require an alternative representation of quasi-simple functions in terms of parse trees.

Let Υ be the set of all parse trees whose leaves are simple functions and whose nodes are the operations: min, max and convex combination. Clearly, every tree $\Delta \in \Upsilon$ corresponds to a unique quasi-simple function which we will call $qs(\Delta)$. Conversely, every quasi-simple function corresponds to infinitely many trees from Υ. The definition below gives us a unique representative.

Definition 24. *Let the rank of a quasi-simple function $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$, denoted $rank(f)$, be the smallest k such that there is a tree $\Delta \in \Upsilon$ of height k such that $qs(\Delta) = f$. For any quasi-simple function $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$ we define a unique representative parse tree Δ_f by induction on the rank of f.*

- *If $rank(f) = 0$, then let Δ_f to be any tree with height 0 such that $qs(\Delta_f) = f$.*
- *If $rank(f) = k+1$ for some $k \in \mathbb{N}$, there must be an operation op (either min, max or convex combination) and integer n such that f is obtained by taking the op of the quasi-simple functions $f_1, \ldots, f_n$, each of which has rank at most k. Therefore, by induction we have representatives $\Delta_{f_1}, \ldots, \Delta_{f_n}$ for $f_1, \ldots, f_n$. Let Δ_f be the tree with root op and subtrees $\Delta_{f_1}, \ldots, \Delta_{f_n}$. Clearly, by construction we have $qs(\Delta_f) = f$.*

The first technical property of quasi-simple functions will allow us to establish that when we take a delay so that a boundary of a region is reached, quasi-simplicity is preserved.

Lemma 25. *Let $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$ be a quasi-simple function, c a clock and i an integer such that $\nu(c) \geqslant i$ for all $\nu \in X$. If $f^{elapse} : X \to \mathbb{R}^{\infty}_{\geqslant 0}$ is the function where for any $\nu \in X$ we have $f^{elapse}(\nu) = t_\nu + f(\nu + t_\nu)$ and $t_\nu = \nu(c) - i$, then f^{elapse} is quasi-simple.*

PROOF. Consider any quasi-simple function $f : X \to \mathbb{R}^{\infty}_{\geqslant 0}$ and let Δ_f be its representative parse tree constructed using Definition 24. Let Δ_f^{mod} be the modified parse tree where any leaf labeled with a constant simple function e is replaced by the non-constant simple function $e' - \nu(c)$, where $e' = e + i$.

We will prove that $f^{elapse} = qs(\Delta_f^{mod})$, which demonstrates that f^{elapse} is quasi-simple as required. The proof is by induction on the rank of f. If $rank(f) = 0$, then there are two cases to consider.

- If Δ_f is a leaf labelled with a constant simple function which for any $\nu \in X$ returns e for some $e \in \mathbb{N}$, then for any $\nu \in X$:

$$\begin{aligned} f^{elapse}(\nu) &= t_\nu + f(\nu + t_\nu) \\ &= t_\nu + e && \text{(by definition of } \Delta_f) \\ &= i - \nu(c) + e && \text{(by definition of } t_\nu) \\ &= e' - \nu(c) && \text{(by definition of } e') \end{aligned}$$

  which equals $qs(\Delta_g^{mod})(\nu)$ as required.

- If Δ_f is a leaf labelled with a simple function which for any $\nu \in X$ returns $e - \nu(c')$ for some $e \in \mathbb{N}$ and clock c', then we have for any $\nu \in X$:

$$\begin{aligned} f^{elapse}(\nu) &= t_\nu + f(\nu + t_\nu) \\ &= t_\nu + e - (\nu(c') + t_\nu) && \text{(by definition of } \Delta_f) \\ &= e - \nu(c') && \text{(rearranging)} \end{aligned}$$

  which again equals $qs(\Delta_g^{mod})(\nu)$ as required.

For the inductive step, suppose $rank(f) = k+1$ for some $k \in \mathbb{N}$ and for any quasi-simple function of rank less than or equal to k the result holds. Since $rank(f) = k+1$ there must be an operation op (either min, max or convex combination) and integer n such that f is obtained by taking the op of some quasi-simple functions $f_1, \ldots, f_n$, each of which has rank at most k. Now, for any $\nu \in X$:

$$\begin{aligned} f^{elapse}(\nu) &= t_\nu + f(\nu + t_\nu) \\ &= t_\nu + op(f_1(\nu + t_\nu), \ldots, f_n(\nu + t_\nu)) && \text{(by definition of } f) \\ &= op(t_\nu + f_1(\nu + t_\nu), \ldots, t_\nu + f_n(\nu + t_\nu)) && \text{(rearranging)} \\ &= op(f_1^{elapse}(\nu), \ldots, f_n^{elapse}(\nu)) && \text{(by definition of } f_i^{elapse} \text{ for } 1 \leqslant i \leqslant n) \end{aligned}$$

$$= op(qs(\Delta_{f_1}^{mod})(\nu), \ldots, qs(\Delta_{f_n}^{mod})(\nu)) \quad \text{(by the inductive hypothesis)}$$
$$= qs(\Delta_f^{mod})(\nu) \quad \text{(by definition of } \Delta_f^{mod})$$

This completes the induction step, and hence the lemma holds. □

The next lemma states that resetting clocks preserves quasi-simplicity.

Lemma 26. *For any region ζ and quasi-simple function $g : \zeta_C \to \mathbb{R}^{\infty}_{\geqslant 0}$, the function $g^{reset} : \zeta \to \mathbb{R}^{\infty}_{\geqslant 0}$ defined by $g^{reset}(\nu) = g(\nu_C)$ is quasi-simple.*

PROOF. For a quasi-simple function f, let Δ_f^{mod} be the modified parse tree of Δ_f where a leaf labelled with a non-constant simple function which for any $\nu \in \zeta_C$ returns $e - \nu(c)$ for some integer e and clock $c \in C$ is replaced with a leaf labelled by the constant function e. The proof follows by showing $f^{reset} = qs(\Delta_f^{mod})$ for all quasi-simple functions f. This proof is by induction on the rank of f.

If $rank(f) = 0$, then Δ_f is a leaf and there are three cases to consider.

- If Δ_f is a leaf labelled with a constant simple function which for any $\nu \in \zeta_C$ returns e for some $e \in \mathbb{N}$, then for any $\nu \in \zeta$ by construction:

$$\begin{aligned} f^{reset}(\nu) &= f(\nu_C) \\ &= e && \text{(by definition of } \Delta_f) \\ &= qs(\Delta_f^{mod}) && \text{(by construction)} \end{aligned}$$

- If Δ_f is a leaf labelled with a simple function which for any $\nu \in \zeta_C$ returns $e - \nu(c')$ for some $e \in \mathbb{N}$ and clock $c' \notin C$, then for any $\nu \in \zeta$:

$$\begin{aligned} f^{reset}(\nu) &= f(\nu_C) \\ &= e - \nu_C(c) && \text{(by definition of } \Delta_f) \\ &= e - \nu(c) && \text{(since } c \notin C) \\ &= qs(\Delta_f^{mod}) && \text{(by construction)} \end{aligned}$$

- If Δ_f is a leaf labelled with a simple function which for any $\nu \in \zeta_C$ returns $e - \nu(c)$ for some $e \in \mathbb{N}$ and clock $c \in C$, then we have for any $\nu \in \zeta$:

$$\begin{aligned} f^{reset}(\nu) &= f(\nu_C) \\ &= e - \nu_C(c) && \text{(by definition of } \Delta_f) \\ &= e && \text{(since } c \in C) \\ &= qs(\Delta_f^{mod}) && \text{(by construction)} \end{aligned}$$

For the inductive step, suppose $rank(f) = k+1$ for some $k \in \mathbb{N}$ and for any quasi-simple function of rank less than or equal to k the result holds. Since $rank(f) = k+1$ there must be an operation op (either min, max and convex combination) and integer

n such that f is obtained by taking the op of some quasi-simple functions $f_1, \ldots, f_n$, each of which has rank at most k. Therefore for any $\nu \in \zeta$ by construction:

$$\begin{aligned} f^{reset}(\nu) &= f(\nu_C) \\ &= op(f_1(\nu_C), \ldots, f_n(\nu_C)) && \text{(definition of } f\text{)} \\ &= op(f_1^{reset}(\nu), \ldots, f_n^{reset}(\nu)) && \text{(by definition of } f_i^{reset} \text{ for } 1 \leqslant i \leqslant n\text{)} \\ &= op(qs(\Delta_{f_1}^{mod})(\nu), \ldots, qs(\Delta_{f_n}^{mod})(\nu)) && \text{(by the inductive hypothesis)} \\ &= qs(\Delta_f^{mod})(\nu) && \text{(by definition of } \Delta_g^{mod}\text{)} \end{aligned}$$

which completes the proof. □

The following technical lemma will allow us to establish that, assuming quasi-simplicity in successor states, the players' optimal behaviour is to pick delays so that boundaries of regions are reached.

Lemma 27. *Let $f : X \to \mathbb{R}_{\geqslant 0}^{\infty}$ be a quasi-simple function. For any $x \in X$ and $t \in \mathbb{R}_{\geqslant 0}$ such that $x+t \in X$:*

- $\sup_{t' \leqslant t \wedge x+t' \in X} \{t'+f(x+t')\} = t+f(x+t)$;
- $\inf_{t' \geqslant t \wedge x+t' \in X} \{t'+f(x+t')\} = t+f(x+t)$.

PROOF. Consider any quasi-simple function $f : X \to \mathbb{R}_{\geqslant 0}^{\infty}$ and clock x. It suffices to show that the function $t \mapsto t+f(x+t)$ is increasing. Now for any $t_1, t_2 \in \mathbb{R}_{\geqslant 0}$ such that $t_1 \leqslant t_2$ and $x+t_1, x+t_2 \in X$, we have:

$$\begin{aligned} t_2+f(x+t_2) &= t_1+f(x+t_1)+((t_2-t_1)+(f(x+t_2)-f(x+t_1))) \\ &\geqslant t_1+f(x+t_1) \end{aligned}$$

where the inequality follows since the term $(t_2-t_1)+(f(x+t_2)-f(x+t_1))$ is non-negative by the non-expansiveness of f (see Lemma 23). □

5.2. Establishing correspondence of PTGA and boundary region abstraction

Having introduced quasi-simple functions and their properties, we will now show how they relate to PTGAs and how they can be utilised to finish the proof of Theorem 19. The proof is notationally heavy, and to alleviate some of the technical notation we first introduce a number of functions (and properties of these functions) that will allow us to abbreviate some of the notation. Intuitively, these functions are counterparts to $\overline{\mathsf{Val}}$ functions that in addition to an initial state also take the first action to be taken.

Definition 28. *Let $n \in \mathbb{N}$, $\ell \in L$ and $\zeta \subseteq Inv(\ell)$. For $\star \in \{\mathrm{Min}, \mathrm{Max}\}$ and $(a, \zeta', \mathrm{opt}) \in \widehat{A}_\star(\ell, \zeta)$, let $\overline{\mathsf{Val}}_{\widehat{\mathsf{T}}}^{n+1}((\ell, \cdot, \zeta), (a, \zeta', \mathrm{opt})) : \overline{\zeta} \to \mathbb{R}_{\geqslant 0}$ be the function where for any $\nu \in \overline{\zeta}$ we have $\overline{\mathsf{Val}}_{\widehat{\mathsf{T}}}^{n+1}((\ell, \nu, \zeta), (a, \zeta', \mathrm{opt}))$ equal*

$$\mathrm{opt}_{\nu+t \in \zeta'} \left\{ t + \sum_{(\tilde{\ell}, \tilde{\nu}, \tilde{\zeta}) \in \widehat{S}} \widehat{p}_\star(\ell, \nu, \zeta), (a, \zeta', \mathrm{opt}))(\tilde{\ell}, \tilde{\nu}, \tilde{\zeta}) \cdot \overline{\mathsf{Val}}_{\widehat{\mathsf{T}}}^{n}(\tilde{\ell}, \tilde{\nu}, \tilde{\zeta}) \right\}.$$

Furthermore, for $\nu \in Inv(\ell)$ and $(t,a) \in A_{\mathrm{Min}} \cup A_{\mathrm{Max}}$ such that $\nu{+}t \in \bar{\zeta}$ let:

$$\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}((\ell,\nu),(t,a)) = t + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,a](C,\ell') \cdot \overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(\ell',(\nu{+}t)_C)$$

$$\overline{\mathsf{Val}}^{n+1}_{mix}((\ell,\nu),(t,a),\zeta) = t + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,a](C,\ell') \cdot \overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\ell',(\nu{+}t)_C,\zeta_C)\,.$$

Intuitively, $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}((\ell,\nu,\zeta),(a,\zeta',\mathrm{opt}))$ corresponds to the optimal value in (ℓ,ν,ζ) when the length of the horizon is $n{+}1$ and the first action performed is fixed to be (a,ζ',opt). Similarly, $\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}((\ell,\nu),(t,a))$ corresponds to the optimal value in (ℓ,ν) for the length of the horizon $n{+}1$ when the first action performed is (t,a). The definition of $\overline{\mathsf{Val}}^{n+1}_{mix}((\ell,\nu),(t,a),\zeta)$ gives an auxiliary function which combines the values of $[\![\mathsf{T}]\!]$ and $\widehat{\mathsf{T}}$. Intuitively, it corresponds to taking a fixed action in $[\![\mathsf{T}]\!]$, and then transferring to $\widehat{\mathsf{T}}$ for n more steps.

Next we show that, within a region, the values in the BRA $\widehat{\mathsf{T}}$ are quasi-simple when we restrict to a finite horizon reachability objectives. To simplify the notation, we assume that in any state player Min can pick at least one action, and that, for each action a player Min can select, there exists an action b player Max can select that is preferred, i.e. $b{=}\mathsf{win}(a,b)$, and also an action b that is not preferred, i.e. $a{=}\mathsf{win}(a,b)$ (in addition, there can be actions b not satisfying any of the two conditions). We refer to this assumption as *choice freedom*.

Lemma 29. *Assume T is choice-free. For any $n \in \mathbb{N}$ and $s \in S$:*

$$\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s) \;=\; \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(s,(t,a)), \sup_{(t',b)\in A_{\mathrm{Max}}(s)\wedge t'\leqslant t} \overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(s,(t',b)) \right\}$$

Furthermore, for any $n \in \mathbb{N}$ and $\widehat{s} \in \widehat{S}$ we have that $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s})$ equals[1]*:*

$$\min_{(a,\zeta,\mathrm{inf})\in \widehat{A}_{\mathrm{Min}}(\widehat{s})} \left\{ \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\mathrm{inf})), \right.$$

$$\left. \max_{\substack{(b,\zeta',\mathrm{sup})\in \widehat{A}_{\mathrm{Max}}(\widehat{s}) \\ (b,\zeta',\mathrm{sup})\leqslant(a,\zeta,\mathrm{inf})}} \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(b,\zeta',\mathrm{sup})), \max_{(b,\zeta,\mathrm{inf})\in \widehat{A}_{\mathrm{Max}}(\widehat{s})} \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(b,\zeta,\mathrm{inf})) \right\}$$

PROOF. For $\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s)$, the proof follows easily using Definition 8, choice-freedom, and properties of win. For $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s})$ we use the definition of $\widehat{\mathsf{win}}$ together with the fact that $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\mathrm{inf})) \leqslant \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\mathrm{sup}))$ for all $a \in Act$. The latter follows from Definition 15. □

From now on, we will assume T is choice-free. Note that this is purely a notational advantage, which will allow us to use Lemma 29. The proofs we give can be easily

[1] Recall that $a{\leqslant}b$ denotes the fact that $\widehat{\mathsf{win}}(a,b){=}a$.

extended to non-choice-free T by omitting an appropriate part of the equations. For example, if $Act_{\mathrm{Min}}(s)=\varnothing$, then the first equation in the Lemma 29 reduces to

$$\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s) = \sup_{(t',b)\in A_{\mathrm{Max}}(s)} \overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(s,(t',b))\,.$$

We now proceed with the following lemma which states that the n-step value functions on a BRA are regionally quasi-simple.

Lemma 30. *For any $n \in \mathbb{N}$, $\ell \in L$ and $\zeta \in \mathcal{R}$ such that $\zeta \subseteq Inv(\ell)$, the function $\overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\ell,\cdot,\zeta) : \overline{\zeta}\to\mathbb{R}_{\geqslant 0}$ is quasi-simple.*

PROOF. Consider any $\ell \in L$ and $\zeta \in \mathcal{R}$ such that $\zeta \subseteq Inv(\ell)$. We proceed by induction on $n \in \mathbb{N}$. For n=0, by Definition 7 we have that $\overline{\mathsf{Val}}^{0}_{\widehat{\mathsf{T}}}(\ell,\cdot,\zeta)$ is constant and equals 0, and hence quasi-simple.

Now suppose the claim holds for $n \in \mathbb{N}$. If $\ell \in L_F$, then $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\ell,\cdot,\zeta)$ is constant and equals 0, and hence quasi-simple. It therefore remains to consider the case when $\ell \notin L_F$. By Definition 15 for $\star \in \{\mathrm{Min},\mathrm{Max}\}$ we have $\widehat{A}_\star(\ell,\nu,\zeta) = \widehat{A}_\star(\ell,\nu',\zeta)$ for all $v, v' \in \overline{\zeta}$, hence we use $\widehat{A}_\star(\ell,\zeta)$ to denote $\widehat{A}_\star(\ell,\nu,\zeta)$ for any $v \in \overline{\zeta}$ and $\star \in \{\mathrm{Min},\mathrm{Max}\}$.

Using induction, Lemma 26, Lemma 25 and the quasi-simplicity of a convex combination of quasi-simple functions, it follows that the function $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}((\ell,\cdot,\zeta),\alpha) : \overline{\zeta}\to\mathbb{R}_{\geqslant 0}$ given in Definition 28 is quasi-simple for any $\star \in \{\mathrm{Min},\mathrm{Max}\}$ and $\alpha \in \widehat{A}_\star(\ell,\zeta)$.

Now, by Definition 7, for any $\nu \in \overline{\zeta}$:

$$\begin{aligned}\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\ell,\nu,\zeta) \;=\; & \min_{\alpha\in\widehat{A}_{\mathrm{Min}}(\ell,\zeta)} \max_{\beta\in\widehat{A}_{\mathrm{Max}}(\ell,\zeta)} \Bigg\{ \widehat{\tau}((\ell,\nu,\zeta),\alpha,\beta) \\ & \qquad + \sum_{(\tilde{\ell},\tilde{\nu},\tilde{\zeta})\in\widehat{S}} \widehat{p}((\ell,\nu,\zeta),\alpha,\beta)(\tilde{\ell},\tilde{\nu},\tilde{\zeta}) \cdot \overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\tilde{\ell},\tilde{\nu},\tilde{\zeta}) \Bigg\} \\ = \; & \min_{\alpha\in\widehat{A}_{\mathrm{Min}}(\ell,\zeta)} \max \left\{ \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}((\ell,\nu,\zeta),\alpha), \max_{\beta\in\widehat{A}_{\mathrm{Max}}(\ell,\zeta)\wedge\beta\leqslant\alpha} \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}((\ell,\nu,\zeta),\beta) \right\}\end{aligned}$$

by Definition 15 and Definition 28. Hence, $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\ell,\cdot,\zeta) : \overline{\zeta}\to\mathbb{R}_{\geqslant 0}$ equals an expression which takes the maxima and minima of quasi-simple functions, and therefore by Definition 22 is also quasi-simple. □

The following lemma demonstrates that, for finite-horizon reachability-time objective, the values in the BRA and PTGA coincide.

Lemma 31. *For any $n \in \mathbb{N}$ and $s \in S$ we have $\overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(s) = \overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\widehat{s})$, and hence $\overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(\ell,\cdot) : \overline{\zeta}\to\mathbb{R}_{\geqslant 0}$ is regionally quasi-simple for any $\ell \in L$ and $\zeta \in \mathcal{R}$ such that $\zeta \subseteq Inv(\ell)$.*

PROOF. Consider any $s{=}(\ell,\nu) \in S$. We proceed by induction on $n \in \mathbb{N}$. If $n{=}0$, then by Definition 7 both $\overline{\mathsf{Val}}^0_{[\![\mathsf{T}]\!]}(s)$ and $\overline{\mathsf{Val}}^0_{\widehat{\mathsf{T}}}(\widehat{s})$ equal 0, and hence the result holds.

Now assume that the lemma holds for some $n \in \mathbb{N}$. If $\ell \in L_F$ then

$$\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s) = \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s}) = 0$$

and the result follows. It therefore remains to consider the case when $\ell \notin L_F$. Using Lemma 29 we have:

$$\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s) = \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max \left\{ \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t,a)),\ \sup_{\substack{(t',b)\in A_{\mathrm{Max}}(s)\\ t'\leqslant t}} \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t',b)) \right\} . \quad (6)$$

Let $A^{\leqslant t}_{\mathrm{Max}}(s) = \{(t',b) \in A_{\mathrm{Max}}(s) \mid t'\leqslant t\}$ be the set of actions available to Player Max in s with delay up to t, let $\mathcal{R}(\nu,t) = [[v],[v{+}t]]$ be the regions obtainable from ν by delaying at most t time units, and $Act_{\mathrm{Max}}(\ell,\zeta) = \{b \in Act_{\mathrm{Max}} \mid \zeta \subseteq E(\ell,b)\}$ be the set of actions of T available when in location ℓ and region ζ. It follows by Definition 12 that:

$$A^{\leqslant t}_{\mathrm{Max}}(s) = \bigcup_{\zeta\in\mathcal{R}(\nu,t)} \{(t',b) \in \mathbb{R}_{\geqslant 0}{\times}Act(\ell,\zeta) \mid \nu{+}t' \in \zeta \wedge t'\leqslant t\} \quad (7)$$

Furthermore, letting $\mathcal{R}(\nu) = \{\zeta \in \mathcal{R} \mid \zeta \subseteq Inv(\ell) \wedge [\nu] \rightarrow^* \zeta\}$ be the set of regions obtainable from ν by some delay and $Act_{\mathrm{Min}}(\ell,\zeta) = \{a \in Act_{\mathrm{Min}} \mid \zeta \subseteq E(\ell,a)\}$ the set of actions of player Min available in location ℓ and region ζ, again by Definition 12 we have:

$$A_{\mathrm{Min}}(s) = \bigcup_{\zeta\in\mathcal{R}(\nu)} \{(t,a) \in \mathbb{R}_{\geqslant 0}{\times}Act_{\mathrm{Min}}(\ell,\zeta) \mid \nu{+}t \in \zeta\} \quad (8)$$

Now, by Definition 28, letting $t^+_{\nu,\zeta'} = \sup\{t' \mid \nu{+}t' \in \zeta'\}$ we have:

$$\begin{aligned}
&\sup_{(t',b)\in A^{\leqslant t}_{\mathrm{Max}}} \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t',b)) \\
&= \sup_{(t',b)\in A^{\leqslant t}_{\mathrm{Max}}} \left\{ t' + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,b](C,\ell') \cdot \overline{\mathsf{Val}}^{n}_{[\![\mathsf{T}]\!]}(\ell',(\nu{+}t')_C) \right\} \\
&= \sup_{(t',b)\in A^{\leqslant t}_{\mathrm{Max}}} \left\{ t' + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,b](C,\ell') \cdot \overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\ell',(\nu{+}t')_C,[(\nu{+}t')_C]) \right\} \\
&\qquad\qquad \text{(by induction)} \\
&= \max_{\substack{\zeta'\in\mathcal{R}(\nu,t)\\ b\in Act_{\mathrm{Max}}(\ell,\zeta')}} \sup_{t'\leqslant t\wedge\nu+t'\in\zeta'} \left\{ t' + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,b](C,\ell') \cdot \overline{\mathsf{Val}}^{n}_{\widehat{\mathsf{T}}}(\ell',(\nu{+}t')_C,\zeta'_C) \right\} \\
&\qquad\qquad \text{(by (7))}
\end{aligned}$$

$$= \max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \left\{ \min\{t, t^+_{\nu,\zeta'}\} + \sum_{(C,\ell') \in 2^{\mathcal{C}} \times L} \delta[\ell,b](C,\ell') \cdot \overline{\mathsf{Val}}^n_{\mathsf{T}}(\ell', (\nu + \min\{t, t^+_{\nu,\zeta'}\})_C, \zeta'_C) \right\}$$

(since $\overline{\mathsf{Val}}^n_{\mathsf{T}}$ is quasi-simple and by Lemma 27)

$$= \max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (\min\{t, t^+_{\nu,\zeta'}\}, b), \zeta') \qquad \text{(by Definition 28)}$$

Now substituting this into (6), it follows that $\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s)$ equals

$$= \inf_{(t,a) \in A_{\mathrm{Min}}(s)} \max \left\{ \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t,a)), \max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (\min\{t, t^+_{\nu,\zeta'}\}, b), \zeta') \right\}$$

$$= \min_{\substack{\zeta \in \mathcal{R}(\nu) \\ a \in A_{\mathrm{Min}}(\ell,\zeta)}} \inf_{\nu + t \in \zeta} \max \left\{ \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t,a)), \max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (\min\{t, t^+_{\nu,\zeta'}\}, b), \zeta') \right\}$$

(by (8))

For any $\zeta \in \mathcal{R}(\nu)$ and $(t,a) \in A_{\mathrm{Min}}(\ell,\zeta)$ the expression

$$\max_{\zeta' \in \mathcal{R}(\nu,t) \wedge b \in Act_{\mathrm{Max}}(\ell,\zeta')} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (\min\{t, t^+_{\nu,\zeta'}\}, b), \zeta')$$

equals the maximum of

$$\max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \setminus \zeta \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (t^+_{\nu,\zeta'}, b), \zeta') \quad \text{and} \quad \max_{b \in Act_{\mathrm{Max}}(\ell,\zeta)} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (\min\{t, t^+_{\nu,\zeta}\}, b), \zeta)$$

and both of these expressions decrease as t decreases. Moreover, letting $t^-_{\nu,\zeta} = \inf\{t' \mid \nu + t' \in \zeta\}$ and using Lemma 27, Definition 28, and induction, we have:

$$\inf_{\nu + t \in \zeta} \overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s,(t,a)) = \overline{\mathsf{Val}}^{n+1}_{mix}(s, (t^-_{\nu,\zeta}, a), \zeta)$$

Consequently, it follows that $\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s)$ equals:

$$\min_{\substack{\zeta \in \mathcal{R}(\nu) \\ a \in A_{\mathrm{Min}}(\ell,\zeta)}} \max \left\{ \overline{\mathsf{Val}}^{n+1}_{mix}(s, (t^-_{\nu,\zeta}, a), \zeta), \max_{\substack{\zeta' \in \mathcal{R}(\nu,t) \setminus \zeta \\ b \in Act_{\mathrm{Max}}(\ell,\zeta')}} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (t^+_{\nu,\zeta}, b), \zeta'), \max_{b \in Act_{\mathrm{Max}}(\ell,\zeta)} \overline{\mathsf{Val}}^{n+1}_{mix}(s, (t^-_{\nu,\zeta}, b), \zeta) \right\}$$

By definition of $t^-_{\nu,\zeta}$ and $t^+_{\nu,\zeta'}$ and Definition 28 we have:

$$\overline{\mathsf{Val}}^{n+1}_{mix}(s,(t^-_{\nu,\zeta},a),\zeta) = \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\inf))$$
$$\overline{\mathsf{Val}}^{n+1}_{mix}(s,(t^+_{\nu,\zeta},a),\zeta) = \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\sup))$$

and hence, using definition of $\mathcal{R}(\nu,t)$, $\overline{\mathsf{Val}}^{n+1}_{[\![\mathsf{T}]\!]}(s)$ equals:

$$\min_{(a,\zeta,\inf)\in\widehat{A}_{\mathrm{Min}}(\widehat{s})} \left\{ \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(a,\zeta,\inf)), \right.$$
$$\left. \max_{\substack{(b,\zeta',\sup)\in\widehat{A}_{\mathrm{Max}}(\widehat{s}) \\ (b,\zeta',\sup)\leqslant(a,\zeta,\inf)}} \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(b,\zeta',\sup)), \max_{(b,\zeta,\inf)\in\widehat{A}_{\mathrm{Max}}(\widehat{s})} \overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s},(b,\zeta,\inf)) \right\}$$

which from Lemma 29 equals $\overline{\mathsf{Val}}^{n+1}_{\widehat{\mathsf{T}}}(\widehat{s})$, completing the proof. □

In the rest of this subsection we use the lemmas to prove properties of the infinite-horizon setting.

Lemma 32. *The function* $\lim_{n\to\infty}\overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$ *is regionally non-expansive.*

PROOF. The proof follows from Lemma 31 and the fact that a limit of non-expansive functions is a non-expansive function. □

Lemma 33. *The function* $\lim_{n\to\infty}\overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$ *is a solution of the optimality equations* $\overline{\mathsf{Opt}}_{[\![\mathsf{T}]\!]}$.

PROOF. Let $\Gamma = \lim_{n\to\infty}\overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$ and for any $(\ell,\nu)\in S$ and $(t,a)\in A_{\mathrm{Min}}((\ell,\nu))$ let:

$$\Gamma((\ell,\nu),(t,a)) = t + \sum_{(C,\ell')\in 2^{\mathcal{C}}\times L} \delta[\ell,a](C,\ell')\cdot\Gamma(\ell',(\nu{+}t)_C)\,.$$

By Definition 8 and by using similar arguments as those from the proof of Lemma 29, to prove the result it is sufficient to demonstrate that for any $s\in S$:

$$\Gamma(s) = \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \Gamma(s,(t,a)), \sup_{(t',b)\in A_{\mathrm{Max}}(s)\wedge t'\leqslant t} \Gamma(s,(t',b)) \right\}. \quad (9)$$

Showing the left hand side is less than or equal to the right hand side follows easily from the monotonicity of the operator defining $\overline{\mathsf{Opt}}_{[\![\mathsf{T}]\!]}$, i.e. of the operator $\mathcal{F}$: $(S{\to}\mathbb{R}^\infty_{\geqslant 0}){\to}(S{\to}\mathbb{R}^\infty_{\geqslant 0})$ given by

$$\mathcal{F}(\gamma)(s) = \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \gamma(s,(t,a)), \sup_{(t',b)\in A_{\mathrm{Max}}(s)\wedge t'\leqslant t} \gamma(s,(t',b)) \right\}$$

and from the Knaster-Tarski fixpoint theorem which implies that for all ordinals $o_1 \leqslant o_2$ we have $\mathcal{F}^{o_1}(0) \leqslant \mathcal{F}^{o_2}(0)$ where 0 is the lowest element in the complete lattice of functions $S \to \mathbb{R}^{\infty}_{\geqslant 0}$ ordered with respect to $\leqslant$.

We complete the proof of (9) by showing the left hand side is greater than or equal to the right hand side. Consider any $s \in S$. If $\Gamma(s)$ is infinite, then the result follows. On the other hand, if $\Gamma(s)$ is finite, it is sufficient to show that for any $\varepsilon > 0$:

$$\Gamma(s) + \varepsilon \geqslant \inf_{(t,a) \in A_{\mathrm{Min}}(s)} \max \left\{ \Gamma(s,(t,a)), \sup_{(t',b) \in A_{\mathrm{Max}}(s) \wedge t' \leqslant t} \Gamma(s,(t',b)) \right\}.$$

We begin by selecting a finite sequence $t_1, \ldots, t_m$ of positive reals such that for any possible delay t in $s = (\ell, \nu)$ there exists t_i (denoted $nr(t)$) with $[\nu + t] = [\nu + nr(t)]$ and $|t - nr(t)| \leqslant \varepsilon/6$. Note that such a sequence $t_1, \ldots, t_m$ can always be selected as the clock values are bounded. By construction we have for any $t \in \mathbb{R}_{\geqslant 0}$ and $C \subseteq \mathcal{C}$:

$$|(\nu + t) - (\nu + nr(t))| \leqslant \varepsilon/6 \quad \text{and} \quad |(\nu + t)_C - (\nu + nr(t))_C| \leqslant \varepsilon/6 \tag{10}$$

Now for any $(t, a) \in A_{\mathrm{Min}} \cup A_{\mathrm{Max}}$ we have:

$$\begin{aligned}
&\left|\Gamma(s,(t,a)) - \Gamma(s,(nr(t),a))\right| \\
&\leqslant |t - nr(t)| + \sum_{(C,\ell') \in 2^{\mathcal{C}} \times L} \delta[\ell,a](C,\ell') \cdot |\Gamma(\ell',(\nu+t)_C) - \Gamma(\ell',(\nu+nr(t))_C)| \\
&\leqslant |t - nr(t)| + \sum_{(C,\ell') \in 2^{\mathcal{C}} \times L} \delta[\ell,a](C,\ell') \cdot |(\nu+t)_C - (\nu+nr(t))_C| \\
&\qquad\qquad \text{(since } \Gamma \text{ is regionally non-expansive (Lemma 32))} \\
&\leqslant |t - nr(t)| + \sum_{(C,\ell') \in 2^{\mathcal{C}} \times L} \delta[\ell,a](C,\ell') \cdot \varepsilon/6 \qquad \text{(by (10))} \\
&\leqslant |t - nr(t)| + \varepsilon/6 \qquad \text{(since } \delta[\ell,a](C,\ell') \text{ is a distribution)} \\
&\leqslant \varepsilon/6 + \varepsilon/6 \qquad \text{(by construction of } nr(t)) \\
&= \varepsilon/3\,. \qquad (11)
\end{aligned}$$

By similar arguments (using Lemma 23 and Lemma 31) we can show that for any $n \in \mathbb{N}$:

$$\left|\overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t,a)) - \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(nr(t),a))\right| \leqslant \varepsilon/3 \tag{12}$$

Since $\Gamma = \lim_{n \to \infty} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$, for any $1 \leqslant i \leqslant m$ there exists $N_i \in \mathbb{N}$ such that:

$$\left|\Gamma(\ell',(\nu+t_i)_C) - \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(\ell',(\nu+t_i)_C)\right| \leqslant \varepsilon/3$$

for all $\ell' \in L$, $C \subseteq \mathcal{C}$ and $n \geqslant N_i$. Setting $N = 1 + \max_{1 \leqslant i \leqslant m} N_i$, it follows that:

$$\left|\Gamma(s,(nr(t),a)) - \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(nr(t),a))\right| \leqslant \varepsilon/3 \tag{13}$$

for all $(t, a) \in A_{\mathrm{Min}} \cup A_{\mathrm{Max}}$ and $n \geqslant N$. Now using (11) we have for any $n \geqslant N$:

$$\Gamma(s,(t,a)) \leqslant \Gamma(s,(nr(t),a)) + \varepsilon/3$$

$$\leqslant \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(nr(t),a)) + \varepsilon/3 + \varepsilon/3 \qquad \text{(by (13))}$$
$$\leqslant \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t,a)) + \varepsilon/3 + 2\varepsilon/3 \qquad \text{(by (12))}$$
$$= \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t,a)) + \varepsilon \qquad (14)$$

Finally, for any $n \geqslant N$ we have:

$$\inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \Gamma(s,(t,a)), \sup_{(t',b)\in A_{\mathrm{Max}}(s)\wedge t'\leqslant t} \Gamma(s,(t',b)) \right\}$$
$$\leqslant \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t,a)) + \varepsilon, \sup_{\substack{(t',b)\in A_{\mathrm{Max}}(s)\\ t'\leqslant t}} \left\{ \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t',b)) + \varepsilon \right\} \right\} \qquad \text{(by (14))}$$
$$= \inf_{(t,a)\in A_{\mathrm{Min}}(s)} \max\left\{ \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t,a)), \sup_{(t',b)\in A_{\mathrm{Max}}(s)\wedge t'\leqslant t} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s,(t',b)) \right\} + \varepsilon \qquad \text{(rearranging)}$$
$$= \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}(s) + \varepsilon \qquad \text{(by Lemma 29)}$$
$$\leqslant \Gamma(s) + \varepsilon \qquad \text{(since } \overline{\mathsf{Val}}^m_{[\![\mathsf{T}]\!]}(s) \leqslant \overline{\mathsf{Val}}^{m+1}_{[\![\mathsf{T}]\!]} \text{ for all } m)$$

which completes the proof. □

We are now a few steps away from concluding the proof of the main result of the paper.

Theorem 34. $\overline{\mathsf{Val}}_{[\![\mathsf{T}]\!]} = \lim_{n\to\infty} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$.

PROOF. Using Lemma 10 it follows that $\overline{\mathsf{Val}}_{[\![\mathsf{T}]\!]} \geqslant \lim_{n\to\infty} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$. On the other hand, Lemma 33 states that $\lim_{n\to\infty} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$ is a solution of the equations $\overline{\mathsf{Opt}}_{[\![\mathsf{T}]\!]}$ and Lemma 9 states that $\overline{\mathsf{Val}}_{[\![\mathsf{T}]\!]} \leqslant V$ for any solution V of the equations $\overline{\mathsf{Opt}}_{[\![\mathsf{T}]\!]}$. Therefore, we have $\overline{\mathsf{Val}}_{[\![\mathsf{T}]\!]} \leqslant \lim_{n\to\infty} \overline{\mathsf{Val}}^n_{[\![\mathsf{T}]\!]}$, which completes the proof. □

The above theorem together with Lemma 31 tells us that, to compute $\overline{\mathsf{Val}}_{[\![\mathsf{T}]\!]}(s)$, it is sufficient to compute $\lim_{n\to\infty} \overline{\mathsf{Val}}^n_{\widehat{\mathsf{T}}}(s)$, which is equal to $\overline{\mathsf{Val}}_{\widehat{\mathsf{T}}}(s)$ using results similar to [36, 37]. This completes the proof of Theorem 19.

6. Conclusions

In this paper we introduced the reachability-time problem for PTGAs and showed that it is decidable and in NEXPTIME ∩ co-NEXPTIME. Our proof relies on an analysis of step-bounded value functions, showing that they are *quasi-simple* and non-expansive when infinite horizon is taken. This allows us to reduce the problem to the reachability-time problem on a finite abstraction. As opposed to the preliminary version of the work presented in [30], the reduction works for an unrestricted class of PTGAs.

Although the computational complexity of solving games on timed automata is high, UPPAAL Tiga [6] is able to solve practical reachability and safety properties for timed games by using efficient symbolic zone-based algorithms [7, 11]. A natural future direction is to investigate the possibility of devising similar algorithms for probabilistic timed games.

On the theoretical level, we plan to study if our approach can be utilised for extensions of reachability-time objectives by considering an appropriate class of reward-based properties.

Acknowledgements. This work is partly supported by ERC Advanced Grant VERIWARE and EPSRC grant EP/M023656/1. Vojtěch Forejt is also affiliated with FI MU, Brno, Czech Republic.

References

[1] P. Ramadge, W. Wonham, The control of discrete event systems, in: Proc. IEEE, Vol. 77(1), 1989, pp. 81–98.

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

[3] E. Asarin, O. Maler, As soon as possible: Time optimal control for timed automata, in: F. Vaandrager, J. van Schuppen (Eds.), Proc. 2nd Int. Workshop Hybrid Systems: Computation and Control (HSCC'99), Vol. 1569 of LNCS, Springer, 1999, pp. 19–30.

[4] R. Alur, M. Bernadsky, P. Madhusudan, Optimal reachability for weighted timed games, in: J. Díaz, J. Karhumäki, A. Lepistö, D. Sannella (Eds.), Proc. 31st Int. Coll. Automata, Languages and Programming (ICALP'04), Vol. 3142 of LNCS, Springer, 2004, pp. 122–133.

[5] P. Bouyer, F. Cassez, E. Fleury, K. G. Larsen, Optimal strategies in priced timed game automata, in: K. Lodaya, M. Mahajan (Eds.), Proc. 24th Int. Conf. Foundations of Software Technology and Theoretical Computer Science (FSTTCS'04), Vol. 3328 of LNCS, Springer, 2004, pp. 148–160.

[6] G. Behrmann, A. Cougnard, A. David, E. Fleury, K. Larsen, D. Lime, UPPAAL-Tiga: Time for playing games!, in: W. Damm, H. Hermanns (Eds.), Proc. 19th Int. Conf. Computer Aided Verification (CAV'07), Vol. 4590 of LNCS, Springer, 2007, pp. 121–125.

[7] F. Cassez, J. Jessen, K. Larsen, J. Raskin, P. Reynier, Automatic synthesis of robust and optimal controllers: an industrial case study, in: F. Vaandrager, J. van Schuppen (Eds.), Proc. 12th Int. Conf. Hybrid Systems: Computation and Control (HSCC 2009), Vol. 5469 of LNCS, Springer, 2009, pp. 90–104.

[8] M. Kwiatkowska, G. Norman, R. Segala, J. Sproston, Automatic verification of real-time systems with discrete probability distributions, Theoretical Computer Science 282 (2002) 101–150.

[9] H. Jensen, Model checking probabilistic real time systems, in: B. Bjerner, M. Larsson, B. Nordström (Eds.), Proc. 7th Nordic Workshop Programming Theory, Report 86:247–261, Chalmers University of Technology, 1996, pp. 247–261.

[10] D. Beauquier, Probabilistic timed automata, Theoretical Computer Science 292 (1) (2003) 65–84.

[11] I. AlAttili, F. Houben, G. Igna, S. Michels, F. Zhu, F. Vaandrager, Adaptive scheduling of data paths using Uppaal Tiga, in: Proc. 1st Workshop Quantitative Formal Methods: Theory and Applications (QFM'09), Vol. 13 of Electronic Proceedings in Theoretical Computer Science, Open Publishing Association, 2009, pp. 1–11.

[12] T. Brihaye, V. Bruyère, J. Raskin, On optimal timed strategies, in: P. Pettersson, W. Yi (Eds.), Proc. 3rd Int. Conf. Formal Modeling and Analysis of Timed Systems (FORMATS'05), Vol. 3829 of LNCS, Springer, 2005, pp. 49–64.

[13] P. Bouyer, T. Brihaye, N. Markey, Improved undecidability results on weighted timed automata, Information Processing Letters 98 (2006) 188–194.

[14] M. Jurdziński, A. Trivedi, Reachability-time games on timed automata, in: L. Arge, C. Cachin, T. Jurdziński, A. Tarlecki (Eds.), Proc. 34th Int. Coll. Automata, Languages and Programming (ICALP'07), Vol. 4596 of LNCS, Springer, 2007, pp. 838–849.

[15] G. Hoffmann, H. Wong-Toi, The input-output control of real-time discrete event systems, in: Proc. Real-Time Systems Symposium (RTSS'92), IEEE Press, 1992, pp. 256–265.

[16] E. Asarin, O. Maler, A. Pnueli, Symbolic controller synthesis for discrete and timed systems, in: P. Antsaklis, W. Kohn, A. Nerode, S. Sastry (Eds.), Hybrid Systems II, Vol. 999 of LNCS, Springer, 1995, pp. 1–20.

[17] T. Brihaye, T. Henzinger, V. Prabhu, J. Raskin, Minimum-time reachability in timed games, in: L. Arge, C. Cachin, T. Jurdziński, A. Tarlecki (Eds.), Proc. 34th Int. Coll. Automata, Languages and Programming (ICALP'07), Vol. 4596 of LNCS, Springer, 2007, pp. 825–837.

[18] M. Jurdziński, A. Trivedi, Average-time games, in: R. Hariharan, M. Mukund, V. Vinay (Eds.), Proc. 28th Int. Conf. Foundations of Software Technology and Theoretical Computer Science (FSTTCS'08), Vol. 2 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008, pp. 340–351.

[19] P. Bouyer, V. Forejt, Reachability in stochastic timed games, in: S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikoletseas, W. Thomas (Eds.), Proc. 36th Int. Coll. Automata, Languages and Programming (ICALP'09), Vol. 5556 of LNCS, Springer, 2009, pp. 103–114.

[20] T. Brázdil, J. Krcál, J. Kretínský, A. Kucera, V. Rehák, Stochastic real-time games with qualitative timed automata objectives, in: P. Gastin, F. Laroussinie (Eds.), Proc. 21st Int. Conf. Concurrency Theory (CONCUR'10), Vol. 6269 of LNCS, Springer, 2010, pp. 207–221.

[21] Y. Oualhadj, P.-A. Reynier, O. Sankur, Probabilistic robust timed games, in: P. Baldan, D. Gorla (Eds.), Proc. 25th Int. Conf. Concurrency Theory (CONCUR'14), Vol. 8704 of LNCS, Springer, 2014, pp. 203–217.

[22] T. Brázdil, H. Hermanns, J. Krcál, J. Kretínský, V. Rehák, Verification of open interactive Markov chains, in: D. D'Souza, T. Kavitha, J. Radhakrishnan (Eds.), IARCS Annual Conf. Foundations of Software Technology and Theoretical Computer Science, (FSTTCS'12), Vol. 18 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012, pp. 474–485.

[23] J. Berendsen, T. Chen, D. Jansen, Undecidability of cost-bounded reachability in priced probabilistic timed automata, in: J. Chen, S. Cooper (Eds.), Proc. 6th Int. Conf. Theory and Applications of Models of Computation (TAMC'09), Vol. 5532 of LNCS, Springer, 2009, pp. 128–137.

[24] M. Jurdziński, M. Kwiatkowska, G. Norman, A. Trivedi, Concavely-priced probabilistic timed automata, in: M. Bravetti, G. Zavattaro (Eds.), Proc. 20th Int. Conf. Concurrency Theory (CONCUR'09), Vol. 5710 of LNCS, Springer, 2009, pp. 415–430.

[25] A. Jovanovic, M. Kwiatkowska, G. Norman, Symbolic minimum expected time controller synthesis for probabilistic timed automata, in: S. Sankaranarayanan, E. Vicario (Eds.), Proc. 13th Int. Conf. Formal Modeling and Analysis of Timed Systems (FORMATS'15), Vol. 9268 of LNCS, Springer, 2015, pp. 140–155.

[26] M. Kwiatkowska, G. Norman, J. Sproston, F. Wang, Symbolic model checking for probabilistic timed automata, Information and Computation 205 (7) (2007) 1027–1077.

[27] F. Laroussinie, J. Sproston, State explosion in almost-sure probabilistic reachability, Information Processing Letters 102 (6) (2007) 236–241.

[28] M. Jurdziński, J. Sproston, F. Laroussinie, Model checking probabilistic timed automata with one or two clocks, Logical Methods in Computer Science 4 (3) (2008) 1–28.

[29] T. Brázdil, L. Korenciak, J. Krcál, P. Novotný, V. Rehák, Optimizing performance of continuous-time stochastic systems using timeout synthesis, in: J. Campos, B. R. Haverkort (Eds.), Proc. 12th Int. Conf. Quantitative Evaluation of Systems (QEST'15), Vol. 9259 of LNCS, Springer, 2015, pp. 141–159.

[30] V. Forejt, M. Kwiatkowska, G. Norman, A. Trivedi, Expected reachability-time games, in: K. Chatterjee, T. Henzinger (Eds.), Proc. 8th Int. Conf. Formal Modelling and Analysis of Timed Systems (FORMATS'10), LNCS, Springer, 2010, pp. 122–136.

[31] J. Filar, K. Vrieze, Competitive Markov Decision Processes, Springer, 1997.

[32] L. de Alfaro, R. Majumdar, Quantitative solution of omega-regular games, Journal of Computer and System Sciences 68 (2004) 374–397.

[33] A. Neyman, S. Sorin (Eds.), Stochastic Games and Applications, Vol. 570 of NATO Science Series C, Kluwer Academic Publishers, 2004.

[34] A. Condon, The complexity of stochastic games, Information and Computation 96 (1992) 203–224.

[35] U. Zwick, M. Paterson, The complexity of mean payoff games on graphs, Theoretical Computer Science 158 (1996) 343–359.

[36] K. Chatterjee, T. Henzinger, Value iteration, in: O. Grumberg, H. Veith (Eds.), 25 Years of Model Checking, Vol. 5000 of LNCS, Springer, 2008, pp. 107–138.

[37] A. Condon, On algorithms for simple stochastic games, in: Advances in Computational Complexity Theory, American Mathematical Society, 1993, pp. 51–73.

[38] D. Dill, Timing assumptions and verification of finite-state concurrent systems, in: J. Sifakis (Ed.), Proc. Automatic Verification Methods for Finite State Systems, Vol. 407 of LNCS, Springer, 1990, pp. 197–212.

[39] G. Norman, D. Parker, J. Sproston, Model checking for probabilistic timed automata, Formal Methods in System Design 43 (2) (2013) 164–190.

[40] S. Tripakis, Verifying progress in timed systems, in: Proc. ARTS'99, Vol. 1601 of LNCS, Springer, 1999, pp. 299–314.

[41] S. Tripakis, S. Yovine, A. Bouajjani, Checking timed Büchi automata emptiness efficiently, Formal Methods in System Design 26 (3) (2005) 267–292.

[42] P. Bouyer, E. Brinksma, K. G. Larsen, Staying alive as cheaply as possible, in: R. Alur, G. Pappas (Eds.), Proc. 7th Int. Workshop Hybrid Systems: Computation and Control (HSCC'04), Vol. 2993 of LNCS, Springer, 2004, pp. 203–218.

[43] T. Henzinger, Z. Manna, A. Pnueli, What good are digital clocks?, in: W. Kuich (Ed.), Proc. 19th Int. Coll. Automata, Languages and Programming (ICALP'92), Vol. 623 of LNCS, Springer, 1992, pp. 545–558.

[44] A. Trivedi, Competitive optimisation on timed automata, Ph.D. thesis, Department of Computer Science, The University of Warwick (2009).