

Relating Average and Discounted Costs for Quantitative Analysis of Timed Systems

Rajeev Alur
University of Pennsylvania, Philadelphia, USA
alur@cis.upenn.edu

Ashutosh Trivedi
University of Pennsylvania, Philadelphia, USA
ashut@cis.upenn.edu

ABSTRACT

Quantitative analysis and controller synthesis problems for reactive real-time systems can be formalized as optimization problems on timed automata, timed games, and their probabilistic extensions. The limiting average cost and the discounted cost are two standard criteria for such optimization problems. In theory of finite-state probabilistic systems, a number of interesting results are available relating the optimal values according to these two different performance objectives. These results, however, do not directly apply to timed systems due to the infinite state-space of clock valuations. In this paper, we present some conditions under which the existence of the limit of optimal discounted cost objective implies the existence of limiting average cost to the same value. Using these results we answer an open question posed by Fahrenberg and Larsen, and give simpler proofs of some known decidability results on (probabilistic) timed automata. We also show the determinacy and decidability of average-time games on timed automata, and expected average-time games on probabilistic timed automata.

Categories and Subject Descriptors

D.4.7 [Organization and Design]: Real-time systems and embedded systems; B.5.2 [Design Aids]: Optimization, Verification

General Terms

Theory, Verification

Keywords

Timed Automata, Quantitative Analysis, Blackwell optimality, and Tauberian Theorems

1. INTRODUCTION

A number of recent case-studies [8, 15, 14, 4] demonstrate the applicability of timed automata [1] for model based design and analysis of embedded real-time systems. Probabilistic extensions of timed automata, such as probabilistic timed automata [21], and duration probabilistic automata [23] are useful in modeling uncertainties (component failure, probability distribution on delays etc.),

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EMSOFT'11, October 9–14, 2011, Taipei, Taiwan.

Copyright 2011 ACM 978-1-4503-0714-7/11/10 ...\$10.00.

while two-player games [3, 7] on timed automata are natural candidates for modeling open real-time systems. The success of timed automata as a model of real-time systems can be partly attributed to its theoretical properties, such as decidable model checking problems, and partly to the availability of an excellent set of tools such as UPPAAL [29], Kronos [20], RED [27], and, PRISM [25].

Performance analysis and optimal controller synthesis questions for real-time embedded systems can be reduced to optimization problems on priced extensions of (probabilistic) timed automata for a given performance metric [4, 2]. The limiting average cost and the discounted cost are two standard criteria for associating a performance metric with an (infinite) execution of a reactive system. For a given execution r let $\pi(r) = \langle p_i \in \mathbb{R} \rangle_{i=0}^{\infty}$ be the sequence of numbers such that p_i is the price at step i . Then, the limiting average cost of the execution r is the limit, if exists, of finite length averages:

$$\text{Avg}(\pi(r)) \stackrel{\text{def}}{=} \lim_{N \rightarrow \infty} \frac{1}{N+1} \sum_{i=0}^N p_i.$$

On the other hand, the discounted cost of the execution r for a discount factor $\lambda \in (0, 1]$ is the total discounted sum of prices, where the price at step i is discounted by $(1 - \lambda)^i$:

$$\text{Disct}^\lambda(\pi(r)) \stackrel{\text{def}}{=} \lambda \sum_{i=0}^{\infty} (1 - \lambda)^i p_i.$$

The notions of limiting average cost and discounted cost are, loosely speaking, quite opposite to each other: while the limiting average cost criterion is insensitive to prices appearing in the initial prefixes, the discounted cost criterion gives more weight to prices appearing early in a sequence than the later ones. Hence it is perhaps surprising at first to learn that these two criteria are related to each-other. Hardy and Littlewood [11] showed that for a bounded (or positive) sequence π , if the limit $\lim_{\lambda \rightarrow 0^+} \text{Disct}^\lambda(\pi)$ exists then the average cost $\text{Avg}(\pi)$ exists as well, and it is equal to $\lim_{\lambda \rightarrow 0^+} \text{Disct}^\lambda(\pi)$. In this paper we study optimization problems on timed automata (and its generalizations) with discounted cost and average cost criteria, and show the relation between their optimal values. In addition to their purely theoretical interest, such relations yield new proofs of decidability of certain average cost games on timed automata.

Connections between optimal discounted cost and average cost criteria are well-studied for probabilistic systems. Mertens and Neyman [24] showed that for two-player zero-sum stochastic games (and hence for Markov decision processes (MDPs)) with finitely many states and actions, the value of the game with average cost objective is equal to the limit of values of games with discounted cost objectives. Lehrer and Sorin [22] showed a more general result in the context of MDPs with infinite state spaces. They proved that

uniform convergence of the limit of optimal discounted cost objective implies uniform convergence of the average reward objective to the same value. The concept of Blackwell-optimal strategy is tightly coupled with the relation between the limit of optimal discounted cost and the optimal average cost. Blackwell [5] showed the existence of strategies in finite Markov decision processes that are optimal for all discount factors sufficiently close to 0. It is well known [11, 26] that every positional Blackwell-optimal strategy is also average optimal, while the opposite is not true. The result of Mertens and Neyman show the existence of Blackwell-optimal strategies on finite stochastic games. To the best of our knowledge no analog of such results is known for timed automata, and the results from finite state probabilistic systems do not automatically generalize to timed systems due to infinite state space of clock valuations.

The main contribution of this paper are the theorems relating the limit of optimal discounted cost and the optimal average cost criteria on timed automata. We also show the existence of ε -Blackwell optimal (Blackwell optimal up to a given precision $\varepsilon > 0$) strategies in priced timed automata. We generalize these results for probabilistic timed automata, and games on probabilistic timed automata. Our results yield a simpler proof of decidability and determinacy of optimal average-time games on timed automata, and the first proof of decidability and determinacy for expected average-time games on probabilistic timed automata. We also prove a conjecture of Fahrenberg and Larsen [9] on the connection between continuously-discounted cost and average cost (per time-unit) optimization problem on timed automata.

The rest of the paper is organized in the following manner. We begin the technical part of the paper by introducing the concepts and the notations in a general setting of MDPs and stochastic games with infinite states and actions. In Section 3 we introduce priced timed automata and present our main result. In Section 4 and 5 we discuss the extensions of our results for probabilistic timed automata and two player games, respectively.

2. PRELIMINARIES

We write \mathbb{N} for the set of natural numbers, \mathbb{R} for the set of real numbers, and $\mathbb{R}_{\geq 0}$ for the set of non-negative reals. For sets X and Y , we write $[X \rightarrow Y]$ for the set of functions $F : X \rightarrow Y$ and $[X \dashrightarrow Y]$ for the set of partial functions $F : X \dashrightarrow Y$.

A (discrete) *probability distribution* over a (possibly uncountable) set X is a function $F : X \rightarrow [0, 1]$ such that support set of F , i.e., $\text{supp}(F) = \{x \in X : F(x) > 0\}$, is a countable set and $\sum_{x \in X} F(x) = 1$. Let $\mathcal{P}(X)$ denote the set of all probability distributions over X . We say that a probability distribution $F \in \mathcal{P}(X)$ is a *point distribution* if $F(x)=1$ for some $x \in X$.

2.1 Optimization Problems

The semantics of priced timed automata are weighted state transition systems with uncountably infinite states and actions, while the semantics of probabilistic timed automata are Markov decision processes with uncountably infinite states and actions. Before we introduce priced timed automata, and priced probabilistic timed automata, let us first introduce various concepts and definitions in this general framework.

DEFINITION 1 (MARKOV DECISION PROCESSES). A Markov decision process (MDP) is a tuple $\mathcal{M} = (S, A, T, \pi)$ where:

- S is a (possibly uncountable) set of states;
- A is a (possibly uncountable) set of actions;

– $T : S \times A \rightarrow \mathcal{P}(S)$ is a *probabilistic transition function* such that the set $A(s) = \{a \in A : T(s, a) \text{ is defined}\}$ is nonempty for every state $s \in S$;

– $\pi : S \times A \rightarrow \mathbb{R}_{\geq 0}$ is a *bounded and measurable price function*.

A weighted state-transition system is an MDP $\mathcal{M}=(S, A, T, \pi)$ such that $T(s, a)$ is a point distribution for every state $s \in S$ and action $a \in A(s)$. We say that an MDP (or a weighted state-transition system) \mathcal{M} is *finite* if both S and A are finite.

We say that $(s, a, s') \in S \times A \times S$ is a transition of the MDP \mathcal{M} if $T(s, a)(s')$ is positive. A finite run of \mathcal{M} is a finite sequence $\langle s_0, a_1, s_1, \dots, a_n, s_n \rangle$ such that (s_i, a_{i+1}, s_{i+1}) , for each $i < n$, is a transition of \mathcal{M} . Similarly, an infinite run is an infinite sequence $\langle s_0, a_1, s_1, a_2, \dots \rangle$ such that (s_i, a_{i+1}, s_{i+1}) , for each $i \in \mathbb{N}$, is a transition of \mathcal{M} . For a finite run r we write $\text{last}(r)$ and $\text{len}(r)$ for its last state and number of transitions, resp. For example if $r = \langle s_0, a_1, \dots, a_n, s_n \rangle$ then $\text{last}(r) = s_n$ and $\text{len}(r) = n$. We write $\text{first}(r)$ for the first state of a (finite or infinite) run $r = \langle s_0, a_1, \dots \rangle$, here $\text{first}(r) = s_0$. Let FRuns and Runs be the set of finite and infinite runs of \mathcal{M} , and let $\text{FRuns}(s)$ and $\text{Runs}(s)$ be the set of finite and infinite runs r such that $\text{first}(r) = s$.

A *strategy* $\mu : \text{FRuns} \rightarrow \mathcal{P}(A)$ is a function mapping finite runs to probability distributions on the set of enabled actions, i.e. for all finite runs $r \in \text{FRuns}$ we have $\text{supp}(\mu(r)) \subseteq A(\text{last}(r))$. We say that a strategy μ is *pure* if $\mu(r)$ is a point distribution for all $r \in \text{Runs}$, while we say that it is *stationary* if $\text{last}(r) = \text{last}(r')$ implies $\mu(r) = \mu(r')$ for all $r, r' \in \text{Runs}$. We also say that a strategy is *positional* if it is both pure and stationary. We write Σ for the set of all strategies and Π for the set of all positional strategies of the MDP \mathcal{M} .

For a strategy $\mu \in \Sigma$ and a starting state $s \in S$ we write $\text{Run}(s, \mu)$ for the set of all runs consistent with the strategy μ . Notice that if \mathcal{M} is a weighted state-transition system and μ is a pure strategy then $\text{Run}(s, \mu)$ is a singleton set. Given a finite run $r \in \text{FRuns}$ a basic cylinder set $\text{Cyl}(r)$ is defined as the set of all infinite runs with prefix r . Let \mathcal{F} be the σ -algebra generated by all cylinder sets. To analyze an MDP \mathcal{M} under a strategy μ and a starting state s we define the probability space $(\text{Runs}(s, \mu), \mathcal{F}, \text{Prob}(s, \mu))$ over the set of infinite runs $\text{Runs}(s, \mu)$ in the standard manner [26]. Note that $\text{Prob}(s, \mu)$ is the unique probability measure satisfying

$$\text{Prob}(s, \mu)(\text{Cyl}(r)) = \prod_{i=1}^n T(s_{i-1}, a_i)(s_i) \cdot \mu(r_{i-1})(a_i),$$

for all finite runs $r = \langle s_0, a_1, \dots, s_n \rangle$, where $r_i = \langle s_0, a_1, \dots, s_i \rangle$.

Given a *real-valued random variable* $f : \text{Runs} \rightarrow \mathbb{R}$ over the set of infinite runs, using standard techniques from probability theory, we define the expectation $\mathbb{E}(s, \mu)\{f\}$ of this variable with respect to the strategy $\mu \in \Sigma$ when starting in the state $s \in S$.

To compare the performance of an MDP under different strategies, we define optimality criteria (also known as performance criteria, or payoffs) associated with a starting state and a strategy. Discounted cost and average cost are two well-studied criteria in performance analysis of MDPs. For a run $r = \langle s_0, a_1, s_1, \dots \rangle$ we define its λ -discounted cost \mathcal{D}_λ (for $\lambda \in (0, 1]$) and N -average cost \mathcal{A}_N (for $N \geq 1$) as:

$$\mathcal{D}_\lambda(r) = \lambda \sum_{i=0}^{\infty} (1-\lambda)^i \pi(s_i, a_{i+1}) \text{ and } \mathcal{A}_N(r) = \frac{1}{N} \sum_{i=0}^{N-1} \pi(s_i, a_{i+1}).$$

Observe that $\mathcal{D}_\lambda(r)$ is always defined for every run r since π is a bounded function. Due to non-negativity of price functions the following result follows from a result [13] of Hardy and Littlewood.

THEOREM 1. *For every run r of an MDP \mathcal{M} we have that*

$$\begin{aligned} \liminf_{N \rightarrow \infty} \mathcal{A}_N(r) &\leq \liminf_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(r) \\ &\leq \limsup_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(r) \leq \limsup_{N \rightarrow \infty} \mathcal{A}_N(r). \end{aligned}$$

Moreover, if the middle inequality is an equality, then all inequalities are equalities.

The notion of λ -discounted cost \mathcal{D}_λ (for $\lambda \in (0, 1]$) and N -average cost \mathcal{A}_N (for $N \geq 1$) of a run can be generalized to λ -discounted cost $\mathcal{D}_\lambda(s, \mu)$ and N -average cost $\mathcal{A}_N(s, \mu)$ of a strategy μ of \mathcal{M} with starting state s as:

$$\mathcal{D}_\lambda(s, \mu) = \mathbb{E}(s, \mu)\{\mathcal{D}_\lambda\} \text{ and } \mathcal{A}_N(s, \mu) = \mathbb{E}(s, \mu)\{\mathcal{A}_N\}.$$

Finally, we define optimal λ -discounted cost $\mathcal{D}_\lambda(s)$ and optimal N -average cost $\mathcal{A}_N(s)$ of \mathcal{M} with starting state s as:

$$\mathcal{D}_\lambda(s) = \inf_{\mu \in \Sigma} \mathcal{D}_\lambda(s, \mu) \text{ and } \mathcal{A}_N(s) = \inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu).$$

It is natural to ask whether $\mathcal{D}_\lambda(s)$ and $\mathcal{A}_N(s)$ are also related in a manner similar to Theorem 1. Lehrer and Sorin [22] gave a partial answer to this question in the case of uniform convergence of \mathcal{D}_λ and \mathcal{A}_N . We say that a sequence $\langle f_n : S \rightarrow \mathbb{R} \rangle$ of functions converges uniformly to a limiting function $f : S \rightarrow \mathbb{R}$ if for every $\varepsilon > 0$ there exists $N_\varepsilon \in \mathbb{N}$ such that for all $s \in S$ and all $N \geq N_\varepsilon$ we have $|f_N(s) - f(s)| \leq \varepsilon$. On the other hand, the ordinary (point-wise) convergence only require that for every $\varepsilon > 0$ and $s \in S$ there exists $N_{\varepsilon, s} \in \mathbb{N}$ such that for all $N \geq N_{\varepsilon, s}$ we have $|f_N(s) - f(s)| \leq \varepsilon$.

THEOREM 2 ([22]). *For a Markov decision process \mathcal{M} and a function $F : S \rightarrow \mathbb{R}$ we have that $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda = F$ uniformly on S if and only if $\lim_{N \rightarrow \infty} \mathcal{A}_N = F$ uniformly on S .*

Lehrer and Sorin [22] showed that in the absence of uniform convergence this relation does not hold by giving an example of an MDP with countably infinite states and finite action sets where both $\lim_{n \rightarrow \infty} \mathcal{A}_N(s)$ and $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s)$ exist and differ.

A strategy μ is \mathcal{D}_λ -optimal if for all $s \in S$ we have that $\mathcal{D}_\lambda(s) = \mathcal{D}_\lambda(s, \mu)$. Observe that since the optimal λ -discounted cost \mathcal{D}_λ is defined as infimum over λ -discounted costs for an infinite set of strategies, no strategy may achieve the optimal cost. Hence, for a given precision $\varepsilon > 0$, we define the concept of ε -optimal strategies. For $\varepsilon > 0$ we say that a strategy μ is ε - \mathcal{D}_λ -optimal if for all $s \in S$ we have that $\mathcal{D}_\lambda(s) + \varepsilon \geq \mathcal{D}_\lambda(s, \mu)$. The concepts of \mathcal{A}_N -optimal and ε - \mathcal{A}_N -optimal strategies are defined analogously.

For a given $\varepsilon > 0$ we say that a strategy is ε -Blackwell-optimal if for every starting state $s \in S$ there exists a $\lambda_0 \in (0, 1]$ such that it is ε - \mathcal{D}_λ -optimal for all $\lambda \in (0, \lambda_0]$. The concept of Blackwell-optimal strategy is defined in straightforward manner. For $\varepsilon > 0$ we say that a strategy μ is ε -limit-average optimal if for every starting state $s \in S$ there exists $N_0 \in \mathbb{N}$ such that μ is ε - \mathcal{A}_N -optimal for all $N \geq N_0$.

For a given state $s \in S$ we also define the long-run average $\mathcal{A}(s, \mu)$ of a strategy $\mu \in \Sigma$ and the optimal average cost $\mathcal{A}(s)$ as:

$$\mathcal{A}(s, \mu) = \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu) \text{ and } \mathcal{A}(s) = \inf_{\mu \in \Sigma} \mathcal{A}(s, \mu).$$

The concepts of average optimal and ε -average-optimal strategies are defined in straightforward manner. The following result relates these average and discounted cost criteria for finite Markov decision processes.

THEOREM 3 ([24],[26]). *For a finite Markov Decision Process \mathcal{M} and starting state s we have that*

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s).$$

Moreover, there exist positional Blackwell-optimal, positional ε -limit-average optimal, and positional average-optimal strategies.

2.2 Two-Player Games

Two player zero-sum games are a well-established paradigm to study competitive optimization problems arising from performance evaluation of open systems. In these games two players correspond to the controller and the environment, and the goal of the controller is to choose the controllable actions in such a way that optimizes some global objective. We are interested in games played on the infinite graph of configurations of timed automata, where the two players, player Min (typically the controller) and player Max (the environment), aim to minimize and maximize, resp., the objective function over the infinite runs of the system. We are interested in showing the relation between the optimal value of discounted cost objective and average cost objective. In this subsection, we introduce various concepts in the setting of stochastic games with infinite state and action spaces, and review some classical results.

DEFINITION 2 (GAME ARENA). *A (stochastic) game arena is a tuple $\Gamma = (\mathcal{M} = (S, A, T, \pi), S_{\text{Min}}, S_{\text{Max}})$ where*

- \mathcal{M} is a Markov decision process,
- S_{Min} and S_{Max} are the sets of states controlled by player Min and player Max, respectively, such that S_{Min} and S_{Max} form a partition of S , i.e., $S_{\text{Min}} \cup S_{\text{Max}} = S$ and $S_{\text{Min}} \cap S_{\text{Max}} = \emptyset$.

We say that a game arena $\Gamma = (\mathcal{M}, S_{\text{Min}}, S_{\text{Max}})$ is deterministic if \mathcal{M} is a weighted state-transition system. Also, we say that a game arena is *finite* if \mathcal{M} is finite.

A strategy of player Min is a (partial) function $\mu : \text{FRuns} \rightarrow \mathcal{P}(A)$ that is defined for a finite run ending in states controlled by player Min, and returns a distribution over actions available in the last state of the run. Formally, a strategy μ is defined for finite runs $r \in \text{FRuns}$ if $\text{last}(r) \in S_{\text{Min}}$ and it is such that $\text{supp}(\mu(r)) \subseteq A(\text{last}(r))$. A strategy χ of player Max is defined analogously. The concepts of pure, stationary, and positional strategies of players are defined in a manner analogous to MDPs. We write Σ_{Min} and Σ_{Max} for the sets of strategies of players Min and Max, respectively, and Π_{Min} and Π_{Max} for the set of positional strategies of players Min and Max, resp.

We write $\text{Runs}(s, \mu, \chi)$ for all the runs of Γ starting from state s where player Min plays according to the strategy $\mu \in \Sigma_{\text{Min}}$ and player Max plays according to strategy $\chi \in \Sigma_{\text{Max}}$. The probability space $(\text{Runs}(s, \mu, \chi), \mathcal{F}, \text{Prob}(s, \mu, \chi))$ and the expectation $\mathbb{E}(s, \mu, \chi)\{f\}$ of the real-valued random variable $f : \text{Runs} \rightarrow \mathbb{R}$ are defined in a straightforward manner.

A game on a game arena is characterized by the optimality criterion players associate with a run. In a zero-sum game it is sufficient to define the performance criterion of player Min, and the criterion of player Max is the opposite. Similar to the optimization problems discussed in Section 2.1, we consider \mathcal{D}_λ -game, \mathcal{A}_N -game, \mathcal{A} -game corresponding to λ -discounted cost, N -average cost, and the average cost criteria, respectively.

We define the concept of the \mathcal{D}_λ -value of a strategy μ of player Min in a game Γ with starting state s as:

$$\mathcal{D}_\lambda(s, \mu) = \sup_{\chi \in \Sigma_{\text{Max}}} \mathbb{E}(s, \mu, \chi)\{\mathcal{D}_\lambda\},$$

and \mathcal{D}_λ -value of a strategy χ of player Max in a game Γ with starting state s as:

$$\mathcal{D}_\lambda(s, \chi) = \inf_{\mu \in \Sigma_{\text{Min}}} \mathbb{E}(s, \mu, \chi) \{ \mathcal{D}_\lambda \}.$$

The upper value $uval(\mathcal{D}_\lambda)$ of a \mathcal{D}_λ -game is then defined as the upper limit on the value of optimization criterion that player Min can ensure irrespective of the strategy used by player Max, and is equal to

$$uval(\mathcal{D}_\lambda)(s) = \inf_{\mu \in \Sigma_{\text{Min}}} \mathcal{D}_\lambda(s, \mu).$$

The concept of the lower value $lval(\mathcal{D}_\lambda)$ of the \mathcal{D}_λ -game is analogous:

$$lval(\mathcal{D}_\lambda)(s) = \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{D}_\lambda(s, \chi).$$

It is easy to verify that $lval(\mathcal{D}_\lambda)(s) \leq uval(\mathcal{D}_\lambda)(s)$. If for every state $s \in S$ we have that $lval(\mathcal{D}_\lambda)(s) = uval(\mathcal{D}_\lambda)(s)$ then we say that \mathcal{D}_λ -game is determined. In such a case, we say that the value $\mathcal{D}_\lambda(s)$ exists, where $\mathcal{D}_\lambda(s) = lval(\mathcal{D}_\lambda)(s) = uval(\mathcal{D}_\lambda)(s)$.

For a given $\varepsilon > 0$ we say that a strategy χ is ε - \mathcal{D}_λ -optimal if

$$\mathcal{D}_\lambda(s, \chi) \geq \mathcal{D}_\lambda(s) - \varepsilon.$$

and a strategy χ is \mathcal{D}_λ -optimal if $\mathcal{D}_\lambda(s, \chi) = \mathcal{D}_\lambda(s)$. The \mathcal{D}_λ -optimality of the strategies of player Min is defined analogously. Notice that if the \mathcal{D}_λ -game is determined, then each player has an ε -optimal strategy for all $\varepsilon > 0$. We say that \mathcal{D}_λ -game is *positionally determined* if

$$\begin{aligned} \mathcal{D}_\lambda(s) &= \inf_{\mu \in \Pi_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathbb{E}(s, \mu, \chi) \{ \mathcal{D}_\lambda \} \\ &= \sup_{\chi \in \Pi_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \mathbb{E}(s, \mu, \chi) \{ \mathcal{D}_\lambda \} \end{aligned}$$

for all $s \in S$. It is straightforward to see that if \mathcal{D}_λ -game is positionally determined, then both players have *positional* ε -optimal strategies for all $\varepsilon > 0$. The concepts of upper value, lower value, optimal and ε -strategies, etc., for N -average cost games are defined in a similar fashion. For average cost games the only noteworthy difference is:

$$\begin{aligned} \mathcal{A}(s, \mu) &= \sup_{\chi \in \Sigma_{\text{Max}}} \mathbb{E}(s, \mu, \chi) \{ \mathcal{A}^* \} \\ \mathcal{A}(s, \chi) &= \inf_{\mu \in \Sigma_{\text{Min}}} \mathbb{E}(s, \mu, \chi) \{ \mathcal{A}_* \}, \end{aligned}$$

where $\mathcal{A}^*(\text{Run}(s, \mu, \chi)) = \limsup_{N \rightarrow \infty} \mathcal{A}_N(\text{Run}(s, \mu, \chi))$, while $\mathcal{A}_*(\text{Run}(s, \mu, \chi)) = \liminf_{N \rightarrow \infty} \mathcal{A}_N(\text{Run}(s, \mu, \chi))$.

For a given $\varepsilon > 0$ we say that a strategy of player Min (Max) is ε -Blackwell-optimal if for every starting state $s \in S$ there exists a $\lambda_0 \in (0, 1]$ such that it is ε - \mathcal{D}_λ -optimal for all $\lambda \in (0, \lambda_0]$. The concept of a Blackwell-optimal strategy of each player is defined in straightforward manner. We also define ε -limit-average optimal strategies for both players. It is well known that for finite game arenas ε -limit-average optimality is a stronger notion than ε -Blackwell optimality.

PROPOSITION 4 ([24]). *For every finite stochastic game arena Γ , $\varepsilon > 0$, and for both players we have that every ε -limit-average optimal strategy is also ε' -Blackwell optimal for $\varepsilon' = 2\varepsilon$.*

The following result relates the value of discounted cost game to the value of average cost game in a game arena.

THEOREM 5 ([24],[11]). *For a finite game arena Γ , all \mathcal{D}_λ -games, \mathcal{A}_N -games and \mathcal{A} -games are determined. For every starting state s we have that*

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s).$$

Moreover, both players have positional Blackwell-optimal, positional ε -limit-average optimal, and positional average-optimal strategies.

2.3 Operator Based Approach

Rosenberg and Sorin [28] studied the functional operators $\Phi : [0, 1] \times [S \rightarrow \mathbb{R}_{\geq 0}] \rightarrow [S \rightarrow \mathbb{R}_{\geq 0}]$ and $\Psi : [S \rightarrow \mathbb{R}_{\geq 0}] \rightarrow [S \rightarrow \mathbb{R}_{\geq 0}]$ and used them to show the connection between discounted cost and average cost criteria for absorbing games and incomplete information repeated games. We review these operators and some of their properties that we later exploit to prove our results for timed systems. For a game arena Γ we define the operators Φ and Ψ in the following manner: for $\lambda \in (0, 1]$ and $f : S \rightarrow \mathbb{R}_{\geq 0}$ we have

$$\begin{aligned} \Phi(\lambda, f)(s) &\stackrel{\text{def}}{=} \text{opt}_{a \in A(s)} \left\{ \lambda \cdot \pi(s, a) + (1-\lambda) \sum_{s' \in S} T(s, a)(s') \cdot f(s') \right\}, \\ \Psi(f)(s) &\stackrel{\text{def}}{=} \text{opt}_{a \in A(s)} \left\{ \pi(s, a) + \sum_{s' \in S} T(s, a)(s') \cdot f(s') \right\}, \end{aligned}$$

where opt is \inf if $s \in S_{\text{Min}}$, and is \sup if $s \in S_{\text{Max}}$.

The following proposition follows from the fact that Φ is monotonic and for every function $f : S \rightarrow \mathbb{R}$ and $d \in \mathbb{R}$ we have $\Phi(\lambda, f+d) = \Phi(\lambda, f) + (1-\lambda) \cdot d$.

PROPOSITION 6 ([28]). *The operator $\Phi(\lambda, \cdot)$ is a contraction with coefficient $(1-\lambda)$ on the set $[S \rightarrow \mathbb{R}_{\geq 0}]$ with respect to ∞ -norm.*

Let $\mathbf{0} \in [S \rightarrow \mathbb{R}_{\geq 0}]$ be such that $\mathbf{0}(s) = 0$ for all $s \in S$. The following equalities involving operators are easy to verify.

PROPOSITION 7 ([28]). *1. $\mathcal{A}_N = \Phi\left(\frac{1}{N}, \mathcal{A}_{N-1}\right)$.*

2. $\mathcal{A}_N = \frac{1}{N} \Psi^N(\mathbf{0})$.

3. $\mathcal{D}_\lambda = \Phi(\lambda, \mathcal{D}_\lambda)$, and

4. $\mathcal{D}_\lambda = \lim_{N \rightarrow \infty} (\Phi(\lambda, \cdot))^N(f)$, for any $f \in [S \rightarrow \mathbb{R}_{\geq 0}]$.

Proposition 7 and Theorem 5 imply the following proposition (that can also be proved directly by extending the operator approach of [28]).

PROPOSITION 8. *For a finite game arena Γ we have that*

$$\lim_{\lambda \rightarrow 0^+} \lim_{N \rightarrow \infty} \Phi(\lambda, \cdot)^N(\mathbf{0})(s) = \lim_{N \rightarrow \infty} \frac{1}{N} \Psi^N(\mathbf{0})(s),$$

for every $s \in S$.

To emphasize some set's correspondence to its game arena (or MDP) Γ , specially when it is not clear from the context, we affix Γ in the superscript. For example, we write Runs^Γ and $\Sigma_{\text{Min}}^\Gamma$ to denote the set of runs and the set of strategies for player Min, resp., in a stochastic game on Γ .

3. PRICED TIMED AUTOMATA

To present the syntax and the semantics of priced timed automata we need to introduce clock variables and related concepts.

Let \mathcal{C} be a finite set of nonnegative real variables which we refer to as *clocks*. A *clock valuation* on \mathcal{C} is a function $\nu : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$ and we write V for the set of clock valuations. Abusing notation, we treat a valuation ν as a point in $\mathbb{R}_{\geq 0}^{|\mathcal{C}|}$. A *clock constraint* over \mathcal{C} is a conjunction of *simple constraints* of the form $c \bowtie i$ or $c - c' \bowtie i$, where $c, c' \in \mathcal{C}$, $i \in \mathbb{N}$ and $\bowtie \in \{<, >, =, \leq, \geq\}$. A clock zone is a set of clock valuations satisfied by some clock constraint. We write \mathcal{Z} for the set of clock zones over \mathcal{C} .

DEFINITION 3 (PRICED TIMED AUTOMATA). A priced timed automaton is a tuple $\mathcal{T} = (L, \mathcal{C}, Inv, Act, Enb, Rst, \delta, pr)$ where:

- L is a finite set of locations;
- \mathcal{C} is a finite set of clocks;
- $Inv : L \rightarrow \mathcal{Z}$ is an invariant condition;
- Act is a finite set of actions;
- $Enb : L \times Act \rightarrow \mathcal{Z}$ is an action guard function;
- $Rst : Act \rightarrow 2^{\mathcal{C}}$ is a clock reset function;
- $\delta : L \times Act \rightarrow L$ is a transition function; and
- $pr : L \cup Act \rightarrow \mathbb{N}$ is a price function assigning price-rates to locations and fixed prices to actions.

We say that a timed automaton is *bounded* if there exists a constant $K \in \mathbb{N}$ such that $Inv(\ell) \subseteq \{\nu \in V : \nu(c) \leq K \text{ for all } c \in \mathcal{C}\}$ for all $\ell \in L$. For technical convenience we restrict our attention to bounded priced timed automata.

If $\nu \in V$ and $t \in \mathbb{R}_{\geq 0}$ then we write $\nu+t$ for the clock valuation defined by $(\nu+t)(c) = \nu(c)+t$ for all $c \in \mathcal{C}$. For $C \subseteq \mathcal{C}$, we write $\nu[C:=0]$ for the valuation where $\nu[C:=0](c)$ equals 0 if $c \in C$ and $\nu(c)$ otherwise. A *configuration* of a timed automaton is a pair (ℓ, ν) , where ℓ is a location and ν a clock valuation satisfying the invariant of the location ℓ , i.e., $\nu \in Inv(\ell)$. For any delay $t \in \mathbb{R}_{\geq 0}$ 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 $Enb(\ell, a)$ is satisfied) after t time units have elapsed. Furthermore, if the timed action (t, a) is performed from (ℓ, ν) then the next location is equal to $\delta(\ell, a)$, while the next valuation ν' is obtained by resetting clocks in $Rst(a)$ after waiting for t time-units from ν , i.e. $\nu' = (\nu + t)[Rst(a):=0]$. The price associated with this timed action is $pr(\ell) \cdot t + pr(a)$.

DEFINITION 4 (PTA: SEMANTICS). The semantics of a priced timed automaton $\mathcal{T} = (L, \mathcal{C}, Inv, Act, Enb, Rst, \delta, pr)$ is given by a weighted state-transition system $\llbracket \mathcal{T} \rrbracket = (S, A, T, \pi)$ where

- $S \subseteq L \times V$ is the set of states such that $(\ell, \nu) \in S$ if and only if $\nu \in Inv(\ell)$;
- $A = \mathbb{R}_{\geq 0} \times Act$ is the set of timed actions;
- $T : S \times A \rightarrow S$ is the transition function such that for a state $(\ell, \nu) \in S$ and a timed action $(t, a) \in A$ we have that $(t, a) \in A((\ell, \nu))$ if and only if
 1. $\nu+t' \in Inv(\ell)$ for all $t' \leq t$ and
 2. $\nu+t \in Enb(\ell, a)$.

For all $(\ell, \nu) \in S$, and $(t, a) \in A((\ell, \nu))$ we have that

$$T((\ell, \nu), (t, a)) = (\delta(\ell, a), (\nu + t)[Rst(a):=0]).$$

- $\pi : S \times A \rightarrow \mathbb{R}_{\geq 0}$ is the price function and it such that for $(\ell, \nu) \in S$ and $(t, a) \in A((\ell, \nu))$ we have that

$$\pi((\ell, \nu), (t, a)) = pr(\ell) \cdot t + pr(a).$$

We are interested in the relation between the optimal λ -discounted cost \mathcal{D}_λ , optimal N -average cost \mathcal{A}_N , and optimal average cost \mathcal{A} on the weighted state-transition system $\llbracket \mathcal{T} \rrbracket$. The following theorem is our main result.

THEOREM 9. For every bounded priced timed automaton \mathcal{T} and every starting state s we have

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s).$$

Moreover, there exist ε -Blackwell optimal and ε -limit-average optimal strategies for every $\varepsilon > 0$.

In order to prove this theorem we present an abstraction of priced timed automata in the next subsection, and using that we prove the correctness of the reduction of optimal λ -discounted cost and optimal N -average cost problems for a priced timed automaton to corresponding problems on a finite graph. In Section 3.2 we show the existence of ε -Blackwell optimal strategies in a timed automaton \mathcal{T} , and in Section 3.3 we show the proof of Theorem 9 by showing a crucial Lemma 14. Finally, in Section 3.4 we prove a similar theorem for optimal continuous discounting cost problem studied by Fahrenberg and Larsen.

3.1 Region Abstractions

The decidability of the qualitative reachability problem on timed automata was shown by Alur and Dill [1] by showing a reduction to the so-called region graph. The region graph, however, is not suitable for qualitative optimization problems as it does not preserve any timing/price information. Several generalizations of region abstractions, e.g. corner-point abstraction [6, 9], digital clocks [21], and boundary region abstractions [18], have been proposed to solve quantitative optimization problems on timed automata. We present the boundary region graph abstraction, and prove the correctness of the reduction of discounted cost optimization problem on priced timed automata to that on its boundary region abstraction. The boundary region graph has the property that for a fixed starting state the reachable subgraph (a state-transition system) from that state is finite. This subgraph coincides with the corner-point abstraction if the valuation ν of the starting state is a *grid point*, i.e. $\nu \in \mathbb{N}^{|\mathcal{C}|}$.

Before we present the boundary region abstraction let us introduce some concepts and notation related to the region abstraction.

For a bounded timed automaton \mathcal{T} , let $SCC(\nu)$ be the finite set of simple constraints which hold in $\nu \in V$. A *clock region* is a maximal set $R \subseteq V$ such that $SCC(\nu) = SCC(\nu')$ for all $\nu, \nu' \in R$. We write \mathcal{R} for the set of clock regions of \mathcal{T} . Observe that every clock region is an equivalence class of the indistinguishability-by-clock-constraints relation, and vice versa. For a clock valuation ν we write $[\nu]$ for its clock region and, if $R = [\nu]$, we write $R[C:=0]$ for the region $[\nu[C:=0]]$.

For regions $R, R' \in \mathcal{R}$, we say that the clock region R' is in the future of clock region R , or R is in the past of R' , if there are $\nu \in R, \nu' \in R'$ and delay $d \in \mathbb{R}_{\geq 0}$ such that $\nu' = \nu+d$; we then write $R \rightarrow_* R'$. For regions $R, R' \in \mathcal{R}$ such that $R \rightarrow_* R'$ we write $[R, R']$ for the union of the set of regions that are in the future of R and in the past of R' . We say that a clock region R is *thin* if $[\nu] \neq [\nu+\varepsilon]$ for every $\nu \in R$ and $\varepsilon > 0$, and *thick* otherwise. We write $\mathcal{R}_{\text{Thin}}$ and $\mathcal{R}_{\text{Thick}}$ for the sets of thin and thick regions, respectively. For a set $X \subseteq V$ of valuations we write $\text{clos}(X)$ and $\text{bd}(X)$ for the closure and the boundary, respectively, of the set X with respect to Euclidean topology on $\mathbb{R}_{\geq 0}^{|\mathcal{C}|}$.

The main idea of the boundary region abstraction is that in this abstraction from every configuration we permit only those time delays that let the system reach a configuration either in a thin region or close to a boundary of a thick region. There is an additional overhead of storing regions as part of configuration and timed actions, because thick regions are open sets of configurations and the configurations lying on the boundary of a thick region do not belong to the thick region itself.

DEFINITION 5 (BOUNDARY REGION GRAPH). *The boundary region graph of $\mathcal{T} = (L, \mathcal{C}, \text{Inv}, \text{Act}, \text{Enb}, \text{Rst}, \delta, \text{pr})$ with semantics $\llbracket \mathcal{T} \rrbracket = (S, \hat{A}, T, \pi)$ is given by a weighted state-transition system $\hat{\mathcal{T}} = (\hat{S}, \hat{A}, \hat{T}, \hat{\pi})$ where*

– $\hat{S} \subseteq L \times V \times \mathcal{R}$ is a set of states such that $(\ell, \nu, R) \in \hat{S}$ if and only if $(\ell, \nu) \in S$ and $\nu \in \text{clos}(R)$;

– $\hat{A} = \mathbb{R}_{\geq 0} \times \mathcal{R} \times \text{Act}$ is a set of boundary timed actions;

– $\hat{T} : \hat{S} \times \hat{A} \rightarrow \hat{S}$ is a transition function such that for a state $(\ell, \nu, R) \in \hat{S}$ and $(t, R_a, a) \in \hat{A}$ we have that $(t, R_a, a) \in \hat{A}((\ell, \nu, R))$ if and only if $[R, R_a] \subseteq \text{Inv}(\ell)$, $R_a \subseteq \text{Enb}(\ell, a)$, and $\nu + t \in \text{bd}(R_a)$. For all $(\ell, \nu, R) \in \hat{S}$ and $(t, R_a, a) \in \hat{A}((\ell, \nu, R))$ we have

$$\hat{T}((\ell, \nu, R), (t, R_a, a)) = (\delta(\ell, a), \nu + t[Rst(a) := 0], R_a[Rst(a) := 0]).$$

– $\hat{\pi} : \hat{S} \times \hat{A} \rightarrow \mathbb{R}$ is a price function such that for $s = (\ell, \nu, R) \in \hat{S}$ and $(t, R_a, a) \in \hat{A}((\ell, \nu, R))$ we have that

$$\hat{\pi}((\ell, \nu, R), (t, R_a, a)) = \text{pr}(\ell) \cdot t + \text{pr}(a).$$

We write $\hat{\mathcal{T}}_s$ for the weighted state-transition graph obtained by restricting the set of states to the set of reachable states from $s \in \hat{S}$.

By the definition of a boundary region graph it follows that for every $(\ell, \nu, R) \in \hat{S}$ if $(t, R_a, a) \in \hat{A}((\ell, \nu, R))$ then

- either R_a is thin, and $\nu + t \in \text{clos}(R_a)$, or
- R_a is thick, and $t = \inf \{t' \mid \nu + t' \in \text{clos}(R_a)\}$ or $t = \sup \{t' \mid \nu + t' \in \text{clos}(R_a)\}$.

Moreover, such t is a natural number if ν is a grid point, otherwise there is a clock $c \in \mathcal{C}$ and number $b \in \mathbb{N}$ such that $t = b - \nu(c)$. The following property is now immediate.

PROPOSITION 10 ([18]). *If \mathcal{T} is a bounded priced timed automaton then for every $s \in \hat{S}$ the set $\hat{A}(s)$ of enabled boundary actions is finite. Moreover, for every $s \in \hat{S}$, the weighted state-transition system $\hat{\mathcal{T}}_s$ is finite and its size is exponential in the size of \mathcal{T} .*

Now let us review the functional operators Φ and Ψ for a timed automaton \mathcal{T} and the related operators $\hat{\Phi}$ and $\hat{\Psi}$ for its boundary region graph $\hat{\mathcal{T}}$. For every state $s = (\ell, \nu) \in S$ of \mathcal{T} , a function $f : S \rightarrow \mathbb{R}$, and $\lambda \in (0, 1]$ we define

$$\Phi(\lambda, f)(s) \stackrel{\text{def}}{=} \inf_{\tau \in A(s)} \{ \lambda \cdot \pi(s, \tau) + (1 - \lambda) \cdot f(T(s, \tau)) \},$$

$$\Psi(f)(s) \stackrel{\text{def}}{=} \inf_{\tau \in A(s)} \{ \pi(s, \tau) + f(T(s, \tau)) \}.$$

Similarly, for every state $s = (\ell, \nu, R) \in \hat{S}$ of $\hat{\mathcal{T}}$, a function $f : \hat{S} \rightarrow \mathbb{R}$, and $\lambda \in (0, 1]$ we define

$$\hat{\Phi}(\lambda, f)(s) \stackrel{\text{def}}{=} \min_{\tau \in \hat{A}(s)} \{ \lambda \cdot \hat{\pi}(s, \tau) + (1 - \lambda) \cdot f(\hat{T}(s, \tau)) \},$$

$$\hat{\Psi}(f)(s) \stackrel{\text{def}}{=} \min_{\tau \in \hat{A}(s)} \{ \hat{\pi}(s, \tau) + f(\hat{T}(s, \tau)) \}.$$

We say that a function $f : \hat{S} \rightarrow \mathbb{R}$ is regionally concave and C -continuous, if for every location $\ell \in L$, and region $R \in \mathcal{R}$ the function $f(\ell, \cdot, R) : \nu \in \text{clos}(R) \mapsto f(\ell, \nu, R)$ is concave and Lipschitz continuous with constant C . The following proposition is useful in showing the connection between operators on \mathcal{T} and $\hat{\mathcal{T}}$.

PROPOSITION 11. *The function $\hat{\Psi}^N(\mathbf{0})$ and $\hat{\Phi}(\lambda, \cdot)^N(\mathbf{0})$ are regionally concave and Lipschitz continuous with constants $(W \cdot N)$ and W , respectively, where W is the largest location price-rate.*

The proof of this proposition follows from the fact that if a function $f : \hat{S} \rightarrow \mathbb{R}$ is regionally concave and C -continuous then $\hat{\Phi}(\lambda, f)$ and $\hat{\Psi}(f)$ are regionally concave and Lipschitz continuous with constants $\max\{C, W\}$ and $C + W$ respectively. It is shown using the following properties of concave functions:

- If $f : S \rightarrow \mathbb{R}$ is concave, and $g : S \times \mathbb{R}_{\geq 0} \rightarrow S$ is affine, then $f(g(\cdot)) : S \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is concave.
- If $f_1, f_2 : S \rightarrow \mathbb{R}$ are concave and $w_1, w_2 \in \mathbb{R}_{\geq 0}$ then $w_1 \cdot f_1 + w_2 \cdot f_2$ is concave.
- If $f_1, f_2, \dots, f_n : S \rightarrow \mathbb{R}$ are concave functions, then their point-wise minimum is also concave.

The proof for Lipschitz continuity follows from the following properties:

- If $f : S \rightarrow \mathbb{R}$ is k_1 -continuous and $g : S \times \mathbb{R}_{\geq 0} \rightarrow S$ is k_2 -continuous, then $f(g(\cdot)) : S \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is $k_1 \cdot k_2$ -continuous.
- If $f_1, f_2 : S \rightarrow \mathbb{R}$ are k_1 -continuous and k_2 -continuous, respectively, and $w_1, w_2 \in \mathbb{R}_{\geq 0}$ then $w_1 \cdot f_1 + w_2 \cdot f_2$ is $w_1 k_1 + w_2 k_2$ continuous.
- If $f_1, f_2, \dots, f_n : S \rightarrow \mathbb{R}$ are continuous with constants k_1, k_2, \dots, k_n , resp., then their point-wise minimum is also Lipschitz continuous with constant $\max\{k_1, k_2, \dots, k_n\}$.

For a state $s = (\ell, \nu)$ of \mathcal{T} we write $[s] = (\ell, \nu, [\nu])$ for the corresponding state in $\hat{\mathcal{T}}$. The following relation shows the correctness of the reduction of the optimal discounted cost problem from a priced timed automaton to its boundary region graph.

PROPOSITION 12. *For every bounded priced timed automaton \mathcal{T} we have*

$$\begin{aligned} \Psi^N(\mathbf{0})(s) &= \hat{\Psi}^N(\mathbf{0})([s]), \text{ and} \\ \lim_{N \rightarrow \infty} \Phi(\lambda, \cdot)^N(\mathbf{0})(s) &= \lim_{N \rightarrow \infty} \hat{\Phi}(\lambda, \cdot)^N(\mathbf{0})([s]), \end{aligned}$$

for every $s \in S$ and every $N \in \mathbb{N}$.

PROOF. We first show by induction on N that $\Psi^N(\mathbf{0})(s) = \hat{\Psi}^N(\mathbf{0})([s])$. Let us first see the base case for $N = 1$. Let $s = (\ell, \nu) \in S$.

$$\begin{aligned} \Psi(\mathbf{0})(s) &= \inf_{(t, a) \in A(s)} \{ \pi(s, (t, a)) + \mathbf{0}(T(s, (t, a))) \} \\ &= \inf_{(t, a) \in A(\ell, \nu)} \{ \pi(s, (t, a)) \} \\ &= \min_{a \in \text{Act}} \min_{R \in \mathcal{R}} \inf_{\{t : \nu + t \in R\}} \{ \pi(s, (t, a)) \} \\ &= \min_{a \in \text{Act}} \min_{R \in \mathcal{R}} \inf_{\{t : \nu + t \in R\}} \{ \text{pr}(\ell) \cdot t + \text{pr}(a) \} \\ &= \min_{a \in \text{Act}} \min_{R \in \mathcal{R}} \min_{\{t : \nu + t \in \text{bd}(R)\}} \{ \text{pr}(\ell) \cdot t + \text{pr}(a) \} \quad (1) \\ &= \min_{(t, R, a) \in \hat{A}} \{ \pi([s], (t, R, a)) + \mathbf{0}(\hat{T}([s], (t, R, a))) \} \\ &= \hat{\Psi}(\mathbf{0})([s]). \end{aligned}$$

The only non-trivial equality is (1) which follows as for a fixed (ℓ, ν) the function $\text{pr}(\ell) \cdot t + \text{pr}(a)$ is concave (linear) in t , and for every concave function $f : (a, b) \rightarrow \mathbb{R}$ we have that

$$\inf_{\{t : a < t < b\}} f(t) = \{ \bar{f}(a), \bar{f}(b) \},$$

where \bar{f} is the unique continuous extension of f on $[a, b]$.

For the inductive step assume that $\Psi^N(\mathbf{0})(s) = \hat{\Psi}^N(\mathbf{0})([s])$. We show that the equality holds for $N + 1$. Following a reasoning similar to the base case, in order to prove $\Psi^{N+1}(\mathbf{0}) = \hat{\Psi}^{N+1}(\mathbf{0})$ we need to show that for every state $(\ell, \nu) \in S$, action $a \in Act$ and region R in future of ν , the function

$$\text{pr}(\ell) \cdot t + \text{pr}(a) + \hat{\Psi}^N(\mathbf{0})([T(s, (t, a))]) \quad (2)$$

is concave in t on the domain $\{t : \nu + t \in R\}$. It follows easily considering the following facts:

- $\hat{\Psi}^N(\mathbf{0})$ is regionally concave (Proposition 11).
- If $f : S \rightarrow \mathbb{R}$ is concave, and $g : \mathbb{R}_{\geq 0} \rightarrow S$ is affine, then $f(g(\cdot)) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is concave as well.
- If $f_1, f_2 : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ are concave and $w_1, w_2 \in \mathbb{R}_{\geq 0}$ then $w_1 \cdot f_1 + w_2 \cdot f_2$ is concave as well.

The proof for the equivalence of Φ and $\hat{\Phi}$ operators is along the similar lines and hence is omitted. \square

3.2 Blackwell Optimal Strategies

To see the difference between different optimality criteria, consider the timed automaton shown in Figure 1 with four location ℓ_0 – ℓ_3 and one clock x . The triplet on a transitions shows clock constraint of the action guard, action label, and the reset clock set. Let the price of a timed action be equal to the time delay, i.e. $\text{pr}(\ell) = 1$ and $\text{pr}(a) = 0$ for all $\ell \in L$ and $a \in A$. In this simple example there are only three strategies.

- (Strategy a): Choose action a in location ℓ_0 after 0 time units, and then take the only available action d in location ℓ_1 every 2 time unit. The \mathcal{D}_λ cost for this strategy is $2(1 - \lambda)$, \mathcal{A}_N -cost is $2 - 2/N$, and the average cost is 2.
- (Strategy b): Choose action b in location ℓ_0 after 10 time units, and then take the only available action d in location ℓ_2 every 1 time unit. The \mathcal{D}_λ cost for this strategy is $1 + 9\lambda$, \mathcal{A}_N -cost is $1 + 9/N$, and the average cost is 1.
- (Strategy c): Choose action c in location ℓ_0 after 1 time unit, and then take the only available action d in location ℓ_3 every 1 time unit. The \mathcal{D}_λ -cost, the \mathcal{A}_N -cost and the average-cost for this strategy is 1.

It is easy to see that the strategy a is \mathcal{D}_λ -optimal for all $\lambda \geq 0.5$, while for all other $\lambda \in (0, 0.5]$ the strategy c is optimal, hence the strategy c is Blackwell-optimal. Also, strategy c is limit-average-optimal for all $N \geq 2$. Notice that under average cost criteria both strategies b and c are optimal, while it is obvious that strategy c is always preferable to strategy b . To make this example more interesting, consider the case when the guard on the location ℓ_3 is




Figure 1: Discount optimal, average optimal and Blackwell-optimal strategies

strict, i.e. $x > 1$. Clearly there are no Blackwell-optimal or limit-average-optimal strategies in this case. However, for every $\varepsilon > 0$ the strategy that chooses $1 + \varepsilon$ delay is ε -limit-average optimal for all $N \geq (2 - \varepsilon)/(1 - \varepsilon)$, and ε -Blackwell optimal for all $\lambda \in (0, \lambda_0]$ for $\lambda_0 = 1/N_0$. The following theorem shows the existence of ε -Blackwell optimal and ε -limit-average optimal strategies for all bounded priced timed automaton.

THEOREM 13. *For every bounded priced timed automaton \mathcal{T} and every starting state $s \in S$ there exist ε -limit-average optimal and ε -Blackwell optimal strategies for every $\varepsilon > 0$.*

PROOF. We only sketch the existence of ε -limit-average optimal strategies. The proof for the existence of ε -Blackwell optimal strategies is along the similar lines and hence omitted.

Let us fix a state $s \in S$ of \mathcal{T} . Notice that $[s] \in \hat{S}$ is its corresponding state in the boundary region graph $\hat{\mathcal{T}}$ of \mathcal{T} . From Proposition 10 we know that for every $[s] \in \hat{S}$ the reachable weighted state-transition system of $\hat{\mathcal{T}}$ from $[s]$ is finite. The existence of an ε -limit-average optimal strategy in $\hat{\mathcal{T}}$ now follows from Theorem 3. In other words, for every $\varepsilon > 0$ there exists a strategy $\mu_\varepsilon \in \Sigma^{\hat{\mathcal{T}}}$ and $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$ we have

$$\mathcal{A}_N^{\hat{\mathcal{T}}}([s], \mu_\varepsilon) \leq \mathcal{A}_N^{\hat{\mathcal{T}}}([s]) + \varepsilon. \quad (3)$$

It was shown in [6, 18] that for every $s \in S$, $\varepsilon > 0$ and $\mu \in \Sigma^{\hat{\mathcal{T}}}$ there exists $\mu' \in \Sigma^{\mathcal{T}}$ such that for all $N \in \mathbb{N}$ we have

$$\mathcal{A}_N^{\mathcal{T}}(s, \mu') \leq \mathcal{A}_N^{\hat{\mathcal{T}}}([s], \mu) + \varepsilon. \quad (4)$$

Notice that from Proposition 12 and 7 we have that $\mathcal{A}_N^{\hat{\mathcal{T}}}([s]) = \mathcal{A}_N^{\mathcal{T}}(s)$ for all $s \in S$ and $N \in \mathbb{N}$. Combining (3) and (4) it is now easy to see that for every $s \in S$ and $\varepsilon > 0$ there exists a strategy $\mu_* \in \Sigma^{\mathcal{T}}$ and $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$ we have

$$\mathcal{A}_N^{\mathcal{T}}(s, \mu_*) \leq \mathcal{A}_N^{\mathcal{T}}(s) + \varepsilon.$$

The proof is now complete. \square

3.3 Proof of Theorem 9

LEMMA 14. *For every bounded priced timed automaton \mathcal{T} and every starting state s we have*

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s).$$

PROOF. From Proposition 12 we know that for every $s \in S$ we have that

$$\lim_{N \rightarrow \infty} \Phi(\lambda, \cdot)^N(\mathbf{0})(s) = \lim_{N \rightarrow \infty} \hat{\Phi}(\lambda, \cdot)^N(\mathbf{0})([s]).$$

Moreover since the reachable subgraph of $\hat{\mathcal{T}}$ from $[s]$ is a finite weighted state-transition system, from Theorem 3 it implies that the limit $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s)$ exists. Now for every $s \in S$ we show that it implies that $\lim_{N \rightarrow \infty} \mathcal{A}_N(s)$ exists and is equal to $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s)$ via the following equalities.

$$\begin{aligned} \lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) &= \lim_{\lambda \rightarrow 0^+} \lim_{N \rightarrow \infty} \Phi(\lambda, \cdot)^N(\mathbf{0})(s) \\ &= \lim_{\lambda \rightarrow 0^+} \lim_{N \rightarrow \infty} \hat{\Phi}(\lambda, \cdot)^N(\mathbf{0})([s]) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \hat{\Psi}^N(\mathbf{0})([s]) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \Psi^N(\mathbf{0})(s) \\ &= \lim_{N \rightarrow \infty} \mathcal{A}_N(s). \end{aligned}$$

The first and fifth equalities are from Proposition 7. The second and fourth equalities are from Proposition 12. Finally the third equality follows from Proposition 8 and the finiteness (Proposition 10) of the reachable state-transition system of \hat{T} from $[s]$.

Now we show that $\lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s)$, which follows from (5) and (6). First notice that for every strategy $\mu' \in \Sigma$ we have that $\inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \leq \mathcal{A}_N(s, \mu')$, for all $N \in \mathbb{N}$. It implies that $\lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \leq \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu')$. Since μ' was taken arbitrarily, it implies that

$$\lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \leq \inf_{\mu \in \Sigma} \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu). \quad (5)$$

The existence of ε -limit-average optimal strategy (Theorem 13) implies that for arbitrarily small $\varepsilon > 0$ there exists a strategy $\mu_* \in \Sigma$ and $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$ we have that

$$\inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \geq \mathcal{A}_N(s, \mu_*) - \varepsilon.$$

That implies that

$$\lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \geq \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu_*) - \varepsilon.$$

Now it trivially follows that for all $\varepsilon > 0$ we have

$$\lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma} \mathcal{A}_N(s, \mu) \geq \inf_{\mu \in \Sigma} \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu) - \varepsilon. \quad (6)$$

The proof is now complete. \square

3.4 Continuous Discounting

We say that a priced timed automaton is time-divergent if for every strategy σ the run $Run(s, \sigma) = \langle s_0, (t_1, a_1), s_1, \dots \rangle$ is such that $\sum_{i=1}^{\infty} t_i = \infty$. Fahrenberg and Larsen [9] studied bounded and time-divergent priced timed automata and defined continuous discounted cost criterion $\mathbb{D}_\lambda(r)$ of the run r as:

$$\mathbb{D}_\lambda(r) \stackrel{\text{def}}{=} \lambda \cdot \left(\sum_{i=0}^{\infty} e^{-\lambda \Delta_{i+1}} \text{pr}(a_{i+1}) + \int_0^{\infty} e^{-\lambda t} \text{pr}(r)(t) dt \right),$$

where $r = \langle (\ell_0, \nu_0), (t_1, a_1), (\ell_1, \nu_1), \dots \rangle$, $\Delta_i = \sum_{k=0}^i t_k$, $t_0=0$, and $\text{pr}(r) : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is defined as:

$$\text{pr}(r)(t) = \text{pr}(\ell_i) \text{ if } \Delta_i \leq t < \Delta_{i+1}.$$

They conjectured the relation between this continuous discounting criterion with the following average cost criterion introduced in [6]:

$$\mathbb{A}_N(r) \stackrel{\text{def}}{=} \frac{1}{\Delta_{N+1}} \cdot \left(\sum_{i=0}^N \text{pr}(a_{i+1}) + \int_0^{\Delta_{N+1}} \text{pr}(r)(t) dt \right).$$

We define $\mathbb{D}_\lambda(s, \sigma) \stackrel{\text{def}}{=} \mathbb{D}(Run(s, \sigma))$, $\mathbb{A}_N(s, \sigma) \stackrel{\text{def}}{=} \mathbb{A}_N(Run(s, \sigma))$, $\mathbb{A}(s, \sigma) \stackrel{\text{def}}{=} \liminf_{N \rightarrow \infty} \mathbb{A}_N(s, \sigma)$, and $\mathbb{A}(s) \stackrel{\text{def}}{=} \inf_{\sigma \in \Sigma} \mathbb{A}(s, \sigma)$.

Larsen and Fahrenberg [9] conjectured the following.

THEOREM 15. *For a bounded and time-divergent priced timed automaton \mathcal{T} we have that $\lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda(s) = \mathbb{A}(s)$.*

PROOF. The proof of this theorem can be shown using continuous versions of operators Φ and Ψ . However, for convenience we prove this theorem by exploiting known results [9, 6] of the correctness of the reductions of optimal continuous λ -discounted cost and optimal continuous average cost on a bounded and time-divergent priced timed automaton to its boundary region graph (corner-point abstraction).

We begin by observing that the limit $\lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda(s)$ exists since for every λ we have that $\mathbb{D}_\lambda(s)$ on \mathcal{T} is equal to $\mathbb{D}_\lambda([s])$ on \hat{T} , and finiteness (Proposition 10) of the reachable state-transition system

of \hat{T} from $[s]$. Now from the existence of positional Blackwell optimal strategies on finite MDPs we have:

$$\begin{aligned} \lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda(s) &= \lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda([s]) \\ &\stackrel{\text{def}}{=} \lim_{\lambda \rightarrow 0^+} \inf_{\sigma \in \Sigma^{\hat{T}}} \mathbb{D}_\lambda([s], \sigma) \\ &= \lim_{\lambda \rightarrow 0^+} \min_{\sigma \in \Pi^{\hat{T}}} \mathbb{D}_\lambda([s], \sigma) \\ &= \min_{\sigma \in \Pi^{\hat{T}}} \lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda([s], \sigma). \end{aligned}$$

The third equality follows from the existence of positional discount optimal strategies in a finite MDPs [26], while the proof of the last equality follows from the existence of Blackwell-optimal strategies (Theorem 3) in a finite MDP, and the proof is similar to that of the proof of (5) and (6) in the proof of Lemma 14.

For every positional strategy $\sigma \in \Pi^{\hat{T}}$ in \hat{T} we know that the limit $\lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda([s], \sigma)$ exists. Feller [10] showed that for a given Lebesgue-measurable and bounded real function g we have

$$\liminf_{N \rightarrow \infty} V_N \leq \liminf_{\lambda \rightarrow 0^+} V_\lambda \leq \liminf_{\lambda \rightarrow 0^+} V_\lambda \leq \liminf_{N \rightarrow \infty} V_N,$$

where $V_N = \frac{1}{N} \int_0^N g(t) dt$ and $V_\lambda = \lambda \int_0^\infty e^{-\lambda t} g(t) dt$. Moreover, if the middle inequality is an equality then all inequalities are equalities. Combining this fact with Theorem 1 we have that $\lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda([s], \sigma)$ is equal to $\liminf_{N \rightarrow \infty} \mathbb{A}_N([s], \sigma)$. Hence, it follows that

$$\lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda(s) = \min_{\sigma \in \Pi^{\hat{T}}} \liminf_{N \rightarrow \infty} \mathbb{A}_N([s], \sigma).$$

For time-divergent priced timed automata it was shown in [6] that there exist positional \mathbb{A} -optimal strategies, and hence

$$\min_{\sigma \in \Pi^{\hat{T}}} \liminf_{N \rightarrow \infty} \mathbb{A}_N([s], \sigma) = \inf_{\sigma \in \Sigma^{\hat{T}}} \liminf_{N \rightarrow \infty} \mathbb{A}_N([s], \sigma).$$

Now from the correctness [6] of the reduction of average price-per-reward problem for a bounded and reward diverging priced timed automaton to its boundary region graph, it follows that

$$\begin{aligned} \lim_{\lambda \rightarrow 0^+} \mathbb{D}_\lambda(s) &= \inf_{\sigma \in \Sigma^{\hat{T}}} \liminf_{N \rightarrow \infty} \mathbb{A}_N([s], \sigma) \\ &= \inf_{\sigma \in \Sigma^{\mathcal{T}}} \liminf_{N \rightarrow \infty} \mathbb{A}_N(s, \sigma) = \mathbb{A}(s). \end{aligned}$$

The proof is now complete. \square

4. PROBABILISTIC TIMED AUTOMATA

Probabilistic timed automata naturally extend both timed automata and Markov decision processes, and can model uncertainty in real-time systems. Building upon our results on timed automata, we relate expected discounted cost optimization problem to expected average cost problem on probabilistic timed automata.

DEFINITION 6. *A priced probabilistic timed automaton is a tuple $\mathcal{T} = (L, \mathcal{C}, Inv, Act, Enb, \delta, pr)$ where*

- L is a finite set of locations;
- \mathcal{C} is a finite set of clocks;
- $Inv : L \rightarrow \mathcal{Z}$ is an invariant condition;
- Act is a finite set of actions;
- $Enb : L \times Act \rightarrow \mathcal{Z}$ is an action guard function;
- $\delta : L \times Act \rightarrow \mathcal{P}(2^{\mathcal{C}} \times L)$ is a probabilistic transition function; and $pr : L \cup Act \rightarrow \mathbb{N}$ is a price function.

The semantics of a priced probabilistic timed automaton \mathcal{T} is a Markov decision process $\llbracket \mathcal{T} \rrbracket = (S, A, T, \pi)$ where the set S , set A , set $A(s)$ for $s \in S$, and price function π are defined in the same way as in Definition 4. The only difference is the probabilistic transition function $T : S \times A \rightarrow \mathcal{P}(S)$ and is defined as:

$$T((\ell, \nu), (t, a))(\ell', \nu') = \sum \{ \delta(\ell, a)(C, \ell') : C \subseteq \mathcal{C} \text{ and } \nu + t[C:=0] = \nu' \},$$

For $(\ell, \nu), (\ell', \nu') \in S$ and $(t, a) \in A((\ell, \nu))$.

We have the following result analogous to Theorem 9 for probabilistic timed automata.

THEOREM 16. *For every bounded priced probabilistic timed automaton \mathcal{T} and every starting state s we have*

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s).$$

Moreover, there exist ε -Blackwell optimal and ε -limit-average optimal strategies for every $\varepsilon > 0$.

The proof of the existence of ε -Blackwell and ε -limit-average optimal strategies is very similar to the proof of Theorem 13. It uses following results from [16]: for every $s \in S$, strategy μ of a boundary region graph \hat{T} (now a Markov decision process), and $\varepsilon > 0$ there exists a strategy μ' in \mathcal{T} whose N -average cost is at most ε worse than that of μ for all $N \in \mathbb{N}$, and an analogous result for λ -discounted cost. The proof for the equivalence of $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s)$, $\lim_{N \rightarrow \infty} \mathcal{A}_N(s)$ and $\mathcal{A}(s)$ follows the proof of Lemma 14, as also in the case of priced probabilistic timed automata, the operators $\hat{\Phi}(\lambda, f)$ and $\hat{\Psi}(f)$ stay regionally concave for every regionally concave function $f : \hat{S} \rightarrow \mathbb{R}_{\geq 0}$.

5. TIMED GAMES

Two player zero-sum games on timed automata and probabilistic timed automata can model controller synthesis problems for real-time systems. For simplicity we consider turn-based games on probabilistic timed automata. The results presented here can partially be extended to concurrent timed games model of [12] with some effort. A turn based timed game is played between two players Min (the controller) and Max (the environment) who construct an infinite run of the (probabilistic) timed automaton by choosing a timed action when the play (finite execution so far) reaches a location controlled by them. Players choose their moves in order to optimize their respective payoffs. The discussion below is for games on probabilistic timed automata, but results imply similar results for games on timed automata.

DEFINITION 7. *A probabilistic timed game arena is a tuple $\Gamma = (\mathcal{T} = (L, \mathcal{C}, \text{Inv}, \text{Act}, \text{Enb}, \delta, \text{pr}), L_{\text{Min}}, L_{\text{Max}})$, where*

- \mathcal{T} is a priced probabilistic timed automaton with a restricted price function $\text{pr} : L \cup \text{Act} \rightarrow \mathbb{N}$ such that $\text{pr}(\ell)=1$ and $\text{pr}(a)=0$ for all $\ell \in L$ and $a \in \text{Act}$.
- L_{Min} and L_{Max} are the set of locations controlled by player Min and player Max, respectively, such that L_{Min} and L_{Max} form a partition of the set L . i.e. $L_{\text{Min}} \cup L_{\text{Max}} = L$ and $L_{\text{Min}} \cap L_{\text{Max}} = \emptyset$.

The semantics of a probabilistic timed game arena Γ is defined as a stochastic game arena $\llbracket \Gamma \rrbracket = (\llbracket \mathcal{T} \rrbracket = (S, A, T, \pi), S_{\text{Min}}, S_{\text{Max}})$ where $\llbracket \mathcal{T} \rrbracket$ is the semantics of \mathcal{T} , the set $S_{\text{Min}} = \{(\ell, \nu) : \ell \in L_{\text{Min}}\}$ and the set $S_{\text{Max}} = S \setminus S_{\text{Min}}$.

The following theorem is the main result of this section.

THEOREM 17. *For every bounded probabilistic timed game arena Γ and every starting states $s \in S$ we have*

$$\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s) = \lim_{N \rightarrow \infty} \mathcal{A}_N(s) = \mathcal{A}(s),$$

Moreover, there exist ε -Blackwell optimal and ε -limit-average optimal strategies of both players for every $\varepsilon > 0$.

For such game arena it was shown in [12] that \mathcal{D}_λ games are positionally determined for all $\lambda \in (0, 1]$. Using similar techniques, we show that \mathcal{A}_N games are determined for all $N \in \mathbb{N}$. The proof for the equivalence of $\lim_{\lambda \rightarrow 0^+} \mathcal{D}_\lambda(s)$ and $\lim_{N \rightarrow \infty} \mathcal{A}_N(s)$ follows the structure of the proof of Lemma 14. However, the proof of Proposition 12 can not be lifted directly as the functions $\hat{\Phi}(\lambda, f)$ and $\hat{\Psi}(f)$ are no longer regionally concave for a regionally concave function f , due to both minimum and maximum appearing on the right hand sides of the operators $\hat{\Phi}$ and $\hat{\Psi}$. We use quasi-simple functions, introduced by [12], in the place of concave functions to show analog of Proposition 12 for probabilistic timed games.

Due to space constraints, we only sketch the existence of ε -limit-average optimal strategies. It follows closely the structure of the proof of Theorem 13, and uses Theorem 5 and the following fact.

PROPOSITION 18. *For every bounded probabilistic timed game arena Γ , state $s \in S$, strategy $\mu \in \Sigma^{\hat{T}}$ of player Min in \hat{T} , and $\varepsilon > 0$ there exists a strategy $\mu' \in \Sigma^{\hat{T}}$ and $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$ we have that*

$$\mathcal{A}_N^{\hat{T}}(s, \mu') \leq \mathcal{A}_N^{\hat{T}}([s], \mu) + \varepsilon.$$

An analogous result holds for strategies of player Max.

The proof of this proposition uses the fact that the $\lim_{N \rightarrow \infty} \mathcal{A}_N$ is regionally constant and is along the same lines as the proof of Proposition 7 in [19].

Now we show the equality of $\lim_{N \rightarrow \infty} \mathcal{A}_N(s)$ and $\mathcal{A}(s)$. Since for all $N \in \mathbb{N}$ we know that \mathcal{A}_N -games are determined, we have

$$\inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{A}_N(s, \mu, \chi) = \sup_{\chi \in \Sigma_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \mathcal{A}_N(s, \mu, \chi).$$

Let $\chi_\varepsilon \in \Sigma_{\text{Max}}$ be ε -limit-average optimal strategy of player Max and hence there exists $N_0 \in \mathbb{N}$ s.t. χ_ε is ε -optimal for player Max in all \mathcal{A}_N games for $N \geq N_0$. It follows that

$$\begin{aligned} \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{A}_N(s, \mu, \chi) &\leq \inf_{\mu \in \Sigma_{\text{Min}}} \mathcal{A}_N(s, \mu, \chi_\varepsilon) + \varepsilon \\ &\leq \mathcal{A}_N(s, \mu', \chi_\varepsilon) + \varepsilon, \end{aligned}$$

for any arbitrary strategy $\mu' \in \Sigma_{\text{Min}}$. It is now immediate that

$$\begin{aligned} \lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{A}_N(s, \mu, \chi) &\leq \sup_{\chi \in \Sigma_{\text{Max}}} \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu', \chi) + \varepsilon \\ &\leq \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \limsup_{N \rightarrow \infty} \mathcal{A}_N(s, \mu, \chi) + \varepsilon, \end{aligned}$$

since μ' was an arbitrary strategy. Similarly, we show that

$$\begin{aligned} \lim_{N \rightarrow \infty} \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{A}_N(s, \mu, \chi) &\geq \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \liminf_{N \rightarrow \infty} \mathcal{A}_N(s, \mu, \chi) - \varepsilon, \end{aligned}$$

In a similar manner we show the equality of the lower values. Now the determinacy of the \mathcal{A} -games on bounded probabilistic timed game arenas follows from the determinacy of \mathcal{A}_N -games on these arenas. The \mathcal{A} -games on probabilistic timed game arenas are also

known as expected average-time games [12, 17]. Theorem 17, along with the correctness of the reduction [12] of \mathcal{D}_λ -games on a probabilistic timed game arenas to \mathcal{D}_λ -games on corresponding boundary region graphs, shows a reduction from expected average-time games to \mathcal{A} -games on the finite subgraph of reachable states of the corresponding boundary region graph. The following result is now immediate.

THEOREM 19. *Expected average-time games [12] on bounded probabilistic timed game arenas are decidable.*

6. CONCLUSION

Competitive optimization problems on timed automata are central to model based performance evaluation and optimal controller design for embedded systems. Discounted cost and average cost criteria are two standard performance metrics for reactive systems. We studied the limits of λ -discounted cost and N -average cost performance criteria for timed automata and its generalizations. We showed the equivalence of these limits to the optimal average cost criterion for priced timed automata, priced probabilistic timed automata, and probabilistic timed games. Our results show a new proof of determinacy and decidability of expected average-time games on probabilistic timed automata, while simplifying an already known result for average-time games on timed automata.

Acknowledgment

This research was partially supported by NSF awards CNS 0931239, CNS 1035715, and CCF 0915777.

7. REFERENCES

- [1] R. Alur and D.L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126:183–235, 1994.
- [2] R. Alur, S. La Torre, and G. Pappas. Optimal paths in weighted timed automata. In *International Workshop on Hybrid Systems: Computation and Control (HSCC)*, volume 2034 of *LNCS*, pages 49–62. Springer, 2001.
- [3] E. Asarin and O. Maler. As soon as possible: Time optimal control for timed automata. In *International Workshop on Hybrid Systems: Computation and Control (HSCC)*, volume 1569 of *LNCS*, pages 19–30. Springer, 1999.
- [4] G. Behrmann, K. G. Larsen, and J. I. Rasmussen. Priced timed automata: Algorithms and applications. In *International Symposium Formal Methods for Components and Objects (FMCO)*, pages 162–182, 2004.
- [5] D. Blackwell. Discrete dynamic programming. *Annals of Mathematical Statistics*, 33:719–726, 1962.
- [6] P. Bouyer, E. Brinksma, and K. G. Larsen. Staying alive as cheaply as possible. In *International Workshop on Hybrid Systems: Computation and Control (HSCC)*, volume 2993 of *LNCS*, pages 203–218. Springer, 2004.
- [7] T. Brihaye, T. A. Henzinger, V. S. Prabhu, and J. Raskin. Minimum-time reachability in timed games. In *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 4596 of *LNCS*, pages 825–837, 2007.
- [8] F. Cassez, J. Jessen, K. Larsen, J. Raskin, and P. Reynier. Automatic synthesis of robust and optimal controllers: an industrial case study. In *International Workshop on Hybrid Systems: Computation and Control (HSCC)*, volume 5469 of *LNCS*, pages 90–104. Springer, 2009.
- [9] U. Fahrenberg and K. G. Larsen. Discount-optimal infinite runs in priced timed automata. *ENTCS*, 239:179 – 191, 2009. Joint Proceedings of INFINITY 2006, 2007, 2008.
- [10] W. Feller. *An Introduction to Probability Theory and its Applications*, volume II. John Wiley & Sons, New York, second edition, 1971.
- [11] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. 1997.
- [12] V. Forejt, M. Kwiatkowska, G. Norman, and A. Trivedi. Expected reachability-time games. In *International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 6246 of *LNCS*, pages 122–136. Springer, September 2010.
- [13] G. H. Hardy and J. E. Littlewood. Tauberian theorems concerning power series and Dirichlet’s series whose coefficients are positive. *Proc. of London Math. Soc.*, 13:174–191, 1914.
- [14] M. Hendriks and M. Verhoef. Timed automata based analysis of embedded system architectures. In *International Parallel and Distributed Processing Symposium (IPDPS)*, pages 179–187. IEEE Computer Society, 2006.
- [15] P. Herber, J. Fellmuth, and S. Glesner. Model checking system C designs using timed automata. In *Proceedings of the Hardware/Software Codesign and System Synthesis*, pages 131–136. ACM, 2008.
- [16] M. Jurdziński, M. Kwiatkowska, G. Norman, and A. Trivedi. Concavely-priced probabilistic timed automata. In *International Conference on Concurrency Theory (CONCUR)*, volume 5710 of *LNCS*, pages 415–430. Springer, 2009.
- [17] M. Jurdziński and A. Trivedi. Average-time games. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 2 of *Leibniz International Proceedings in Informatics*. Schloss Dagstuhl, 2008.
- [18] M. Jurdziński and A. Trivedi. Concavely-priced timed automata. In *International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 5215 of *LNCS*, pages 48–62. Springer, 2008.
- [19] Marcin Jurdziński and Ashutosh Trivedi. Average-time games on timed automata. *CoRR*, abs/0910.2891, 2009.
- [20] Kronos. <http://www-verimag.imag.fr/TEMPORISE/kronos/>.
- [21] M. Kwiatkowska, G. Norman, R. Segala, and J. Sproston. Automatic verification of real-time systems with discrete probability distributions. *Theoretical Computer Science*, 282:101–150, 2002.
- [22] E. Lehrer and S. Sorin. A uniform Tauberian theorem in dynamic programming. *Mathematics of Operation Research*, 17:303–307, 1992.
- [23] O. Maler, K. G. Larsen, and B.H. Krogh. On zone-based analysis of duration probabilistic automata. In *Proc. of International Workshop on Verification of Infinite-State Systems*, volume 39 of *EPTCS*, pages 33–46, 2010.
- [24] J.-F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10:53–66, 1981.
- [25] Prism. <http://www.prismmodelchecker.org/>.
- [26] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.
- [27] RED. <http://cc.ee.ntu.edu.tw/~farn/red/>.
- [28] D. Rosenberg and S. Sorin. An operator approach to zero-sum stochastic games. *Israel Journal of Mathematics*, 121(1):221–246, 2001.
- [29] UPPAAL. <http://www.uppaal.com/>.