

Mean-Payoff Games on Timed Automata*

Shibashis Guha¹, Marcin Jurdziński²,
Shankara Narayanan Krishna³, and Ashutosh Trivedi⁴

1 The Hebrew University of Jerusalem, Israel

shibashis@cs.huji.ac.il

2 The University of Warwick, UK

marcin@dcs.warwick.ac.uk

3 Indian Institute of Technology Bombay, India

krishnas@cse.iitb.ac.in

4 University of Colorado Boulder, USA

ashutosh.trivedi@colorado.edu

Abstract

Mean-payoff games on timed automata are played on the infinite weighted graph of configurations of priced timed automata between two players – Player Min and Player Max – by moving a token along the states of the graph to form an infinite run. The goal of Player Min is to minimize the limit average weight of the run, while the goal of the Player Max is the opposite. Brenguier, Cassez, and Raskin recently studied a variation of these games and showed that mean-payoff games are undecidable for timed automata with five or more clocks. We refine this result by proving the undecidability of mean-payoff games with three clocks. On a positive side, we show the decidability of mean-payoff games on one-clock timed automata with binary price-rates. A key contribution of this paper is the application of dynamic programming based proof techniques applied in the context of average reward optimization on an uncountable state and action space.

1998 ACM Subject Classification F.1.1 Models of Computation

Keywords and phrases Timed Automata, Mean-Payoff Games, Controller-Synthesis

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2016.44

1 Introduction

The classical mean-payoff games [24, 13, 16, 4] are two-player zero-sum games that are played on weighted finite graphs, where two players – Max and Min – take turn to move a token along the edges of the graph to jointly construct an infinite play. The objectives of the players Max and Min are to respectively maximize and minimize the limit average reward associated with the play. Mean-payoff games are well-studied in the context of optimal controller synthesis in the framework of Ramadge-Wonham [22], where the goal of the game is to find a control strategy that maximises the average reward earned during the evolution of the system. Mean-payoff games enjoy a special status in verification, since μ -calculus model checking and parity games can be reduced in polynomial-time to solving mean-payoff games. Mean-payoff objectives can also be considered as quantitative extensions [17] of classical Büchi objectives, where we are interested in the limit-average share of occurrences of

* The work of Shankara Narayanan Krishna is partly supported by CEFIPRA project AVeRTS. The work of Ashutosh Trivedi is based on research sponsored by DARPA under agreement number FA8750-15-2-0096. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.



accepting states rather than merely in whether or not infinitely many accepting states occur. For a broader discussion on quantitative verification, in general, and the transition from the classical qualitative to the modern quantitative interpretation of deterministic Büchi automata, we refer the reader to Henzinger’s excellent survey [17].

We study mean-payoff games played on an infinite configuration graph of timed automata. Asarin and Maler [3] were the first to study games on timed automata and they gave an algorithm to solve timed games with reachability time objective. Their work was later generalized and improved upon by Alur et al. [1] and Bouyer et al. [8]. Bouyer et al. [7, 5] also studied the more difficult average payoffs, but only in the context of scheduling, which in game-theoretic terminology corresponds to 1-player games. However, they left the problem of proving decidability of 2-player average reward games on priced timed automata open. Jurdziński and Trivedi [20] proved the decidability of the special case of average time games where all locations have unit costs. More recently, mean-payoff games on timed automata have been studied by Brenguier, Cassez and Raskin [10] where they consider average payoff per time-unit. Using the undecidability of *energy games* [9], they showed undecidability of mean-payoff games on weighted timed games with five or more clocks. They also gave a semi-algorithm to solve cycle-forming games on timed automata and characterized the conditions under which a solution of these games gives a solution for mean-payoff games.

On the positive side, we characterize general conditions under which dynamic programming based techniques can be used to solve the mean-payoff games on timed automata. As a proof-of-concept, we consider one-clock binary-priced timed games, and prove the decidability of mean-payoff games for this subclass. Our decidability result can be considered as the average-payoff analog of the decidability result by Brihaye et al. [11] for reachability-price games on timed automata. We strengthen the known undecidability results for mean-payoff games on timed automata in three ways: (i) we show that the mean-payoff games over priced timed games is undecidable for timed games with only three clocks; (ii) secondly, we show that undecidability can be achieved with binary price-rates; and finally, (iii) our undecidability results are applicable for problems where the average payoff is considered per move as well as for problems when it is defined per time-unit.

Howard [18, 21] introduced gain and bias optimality equations to characterize optimal average on one-player finite game arenas. Gain and bias optimality equations based characterization has been extended to two-player game arenas [14] as well as many subclasses of uncountable state and action spaces [12, 6]. The work of Bouyer et al. [6] is perhaps the closest to our approach – they extended optimality equations approach to solve games on hybrid automata with certain strong reset assumption that requires all continuous variables to be reset at each transition, which in the case of timed automata is akin to requiring all clocks to be reset at each transition. To the best of our knowledge, the exact decidability for timed games does not immediately follow from any previously known results.

Howard’s Optimality equations requires two variable per state: the gain of the state and the bias of the state. Informally speaking, the gain of a state corresponds to the optimal mean-payoff for games starting from that state, while the bias corresponds to the limit of transient sum of step-wise deviations from the optimal average. Hence, intuitively at a given point in a game, both players would prefer to first optimize the gain, and then choose to optimize bias among choices with equal gains. We give general conditions under which a solution of gain-bias equations for a finitary abstraction of timed games can provide a solution of gain-bias equations for the original timed game. For this purpose, we exploit a region-graph like abstraction of timed automata [19] called the boundary region abstraction (BRA). Our key contribution is the theorem that states that every solution of gain-bias

optimality equations for boundary region abstraction carries over to the original timed game, as long as for every region, the gain values are constant and the bias values are affine.

The paper is organized in the following manner. In Section 2 we describe mean-payoff games and introduce the notions of *gain* and *bias* optimality equations. This section also introduces mean-payoff games over timed automata and states the key results of the paper. Section 3 introduces the *boundary region abstraction* for timed automata and characterizes the conditions under which the solution of a game played over the boundary region abstraction can be lifted to a solution of mean payoff game over priced timed automata. In Section 4 we present the strategy improvement algorithm to solve optimality equations for mean-payoff games played over boundary region abstraction and connect them to solution of optimality equations over corresponding timed automata. Finally, Section 5 sketches the undecidability of mean-payoff games for binary-priced timed automata with three clocks.

2 Mean-Payoff Games on Timed Automata

We begin this section by introducing mean-payoff games on graphs with uncountably infinite vertices and edges, and show how, and under what conditions, gain-bias optimality equations characterize the value of mean-payoff games. We then set-up mean-payoff games for timed automata and state our key contributions.

2.1 Mean-Payoff Games

► **Definition 1** (Turn-Based Game Arena). A game arena Γ is a tuple $(S, S_{\text{Min}}, S_{\text{Max}}, A, T, \pi)$ where S is a (potentially uncountable) set of states partitioned between sets S_{Min} and S_{Max} of states controlled by Player Min and Player Max, respectively; A is a (potentially uncountable) set of *actions*; $T : S \times A \rightarrow S$ is a partial function called the *transition function*; and $\pi : S \times A \rightarrow \mathbb{R}$ is a partial function called the *price function*.

We say that a game arena is *finite* if both S and A are finite. For any state $s \in S$, we let $A(s)$ denote the set of actions available in s , i.e., the actions $a \in A$ for which $T(s, a)$ and $\pi(s, a)$ are defined. A transition of a game arena is a tuple $(s, a, s') \in S \times A \times S$ such that $s' = T(s, a)$ and we write $s \xrightarrow{a} s'$. A finite play starting at a state s_0 is a sequence of transitions $\langle s_0, a_1, s_1, a_2, \dots, s_n \rangle \in S \times (A \times S)^*$ such that for all $0 \leq i < n$ we have that $s_i \xrightarrow{a_{i+1}} s_{i+1}$ is a transition. For a finite play $\rho = \langle s_0, a_1, \dots, s_n \rangle$ we write $\text{Last}(\rho)$ for the final state of ρ , here $\text{Last}(\rho) = s_n$. The concept of an infinite play $\langle s_0, a_1, s_1, \dots \rangle$ is defined in an analogous way. We write $\text{Runs}(s)$ and $\text{Runs}_{\text{fin}}(s)$ for the set of plays and the set of finite plays starting at $s \in S$ respectively.

A *strategy* of Player Min is a function $\mu : \text{Runs}_{\text{fin}} \rightarrow A$ such that $\mu(\rho) \in A(\text{Last}(\rho))$ for all finite plays $\rho \in \text{Runs}_{\text{fin}}$, i.e. for any finite play, a strategy of Min returns an action available to Min in the last state of the play. A strategy χ of Max is defined analogously and we let Σ_{Min} and Σ_{Max} denote the sets of strategies of Min and Max, respectively. A strategy σ is *positional* if $\text{Last}(\rho) = \text{Last}(\rho')$ implies $\sigma(\rho) = \sigma(\rho')$ for all $\rho, \rho' \in \text{Runs}_{\text{fin}}$. This allows us to represent a positional strategy as a function in $[S \rightarrow A]$. Let Π_{Min} and Π_{Max} denote the set of positional strategies of Min and Max, respectively. For any state s and strategy pair $(\mu, \chi) \in \Sigma_{\text{Min}} \times \Sigma_{\text{Max}}$, let $\text{Run}(s, \mu, \chi)$ denote the unique infinite play $\langle s_0, a_1, s_1, \dots \rangle$ in which Min and Max play according to μ and χ , respectively, i.e. for all $i \geq 0$ we have that $s_i \in S_{\text{Min}}$ implies $a_{i+1} = \mu(\langle s_0, a_1, \dots, s_i \rangle)$ and $s_i \in S_{\text{Max}}$ implies $a_{i+1} = \chi(\langle s_0, a_1, \dots, s_i \rangle)$.

In a mean-payoff game on a game arena, players Min and Max move a token along the transitions indefinitely thus forming an infinite play $\rho = \langle s_0, a_1, s_1, \dots \rangle$ in the game graph.

The goal of player Min is to minimize $\mathcal{A}_{\text{Min}}(\rho) = \limsup_{n \rightarrow \infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} \pi(s_i, a_{i+1})$ and the goal of player Max is to maximize $\mathcal{A}_{\text{Max}}(\rho) = \liminf_{n \rightarrow \infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} \pi(s_i, a_{i+1})$. The *upper value* $\text{Val}^*(s)$ and the lower value $\text{Val}_*(s)$ of a state $s \in S$ are defined as:

$$\text{Val}^*(s) = \inf_{\mu \in \Sigma_{\text{Min}}} \sup_{\chi \in \Sigma_{\text{Max}}} \mathcal{A}_{\text{Min}}(\text{Run}(s, \mu, \chi)) \text{ and } \text{Val}_*(s) = \sup_{\chi \in \Sigma_{\text{Max}}} \inf_{\mu \in \Sigma_{\text{Min}}} \mathcal{A}_{\text{Max}}(\text{Run}(s, \mu, \chi))$$

respectively. It is always the case that $\text{Val}_*(s) \leq \text{Val}^*(s)$. A mean-payoff game is called *determined* if for every state $s \in S$ we have that $\text{Val}_*(s) = \text{Val}^*(s)$. Then, we write $\text{Val}(s)$ for this number and we call it the *value* of the mean-payoff game at state s . We say that a game is *positionally-determined* if for every $\varepsilon > 0$ we have strategies $\mu_\varepsilon \in \Pi_{\text{Min}}$ and $\chi_\varepsilon \in \Pi_{\text{Max}}$ such that for every initial state $s \in S$, we have that

$$\text{Val}_*(s) - \varepsilon \leq \inf_{\mu' \in \Sigma_{\text{Min}}} \mathcal{A}_{\text{Max}}(\text{Run}(s, \mu', \chi_\varepsilon)) \text{ and } \text{Val}^*(s) + \varepsilon \geq \sup_{\chi' \in \Sigma_{\text{Max}}} \mathcal{A}_{\text{Min}}(\text{Run}(s, \mu_\varepsilon, \chi')).$$

For a given ε we call each such strategy an ε -optimal strategy for the respective player.

Given two functions $G : S \rightarrow \mathbb{R}$ (gain) and $B : S \rightarrow \mathbb{R}$ (bias), we say that (G, B) is a solution to the optimality equations for mean-payoff game on $\Gamma = (S, S_{\text{Min}}, S_{\text{Max}}, A, T, \pi)$, denoted $(G, B) \models \text{Opt}(\Gamma)$ if

$$\begin{aligned} G(s) &= \begin{cases} \sup_{a \in A(s)} \{G(s') : s \xrightarrow{a} s'\} & \text{if } s \in S_{\text{Max}} \\ \inf_{a \in A(s)} \{G(s') : s \xrightarrow{a} s'\} & \text{if } s \in S_{\text{Min}}. \end{cases} \\ B(s) &= \begin{cases} \sup_{a \in A(s)} \{\pi(s, a) - G(s) + B(s') : s \xrightarrow{a} s' \text{ and } G(s) = G(s')\} & \text{if } s \in S_{\text{Max}} \\ \inf_{a \in A(s)} \{\pi(s, a) - G(s) + B(s') : s \xrightarrow{a} s' \text{ and } G(s) = G(s')\} & \text{if } s \in S_{\text{Min}}. \end{cases} \end{aligned}$$

We prove the following theorem connecting a solution of the optimality equations with mean-payoff games. We exploit this theorem to solve mean-payoff games on timed automata.

► **Theorem 2.** *If there exists a function $G : S \rightarrow \mathbb{R}$ with finite image and a function $B : S \rightarrow \mathbb{R}$ with bounded image such that $(G, B) \models \text{Opt}(\Gamma)$ then for every state $s \in S$, we have that $G(s) = \text{Val}(s)$ and for every $\varepsilon > 0$ both players have positional ε -optimal strategies.*

Proof. Assume that we are given the functions $G : S \rightarrow \mathbb{R}$ with finite image and $B : S \rightarrow \mathbb{R}$ with bounded image such that $(G, B) \models \text{Opt}(\Gamma)$. In order to prove the result we show, for every $\varepsilon > 0$, the existence of positional strategies μ_ε and χ_ε such that

$$G(s) - \varepsilon \leq \inf_{\mu' \in \Sigma_{\text{Min}}} \mathcal{A}_{\text{Max}}(\text{Run}(s, \mu', \chi_\varepsilon)) \text{ and } G(s) + \varepsilon \geq \sup_{\chi' \in \Sigma_{\text{Max}}} \mathcal{A}_{\text{Min}}(\text{Run}(s, \mu_\varepsilon, \chi')).$$

The proof is in two parts.

- Given $\varepsilon > 0$ we compute the positional strategy $\mu_\varepsilon \in \Pi_{\text{Min}}$ satisfying the following conditions: $\mu_\varepsilon(s) = a$ if

$$G(s) = G(s') \tag{1}$$

$$B(s) \geq \pi(s, a) - G(s) + B(s') - \varepsilon, \tag{2}$$

where $s \xrightarrow{a} s'$. Notice that it is always possible to find such strategy since (G, B) satisfies optimality equations and G is finite image.

Now consider an arbitrary strategy $\chi \in \Sigma_{\text{Max}}$ and consider the run $\text{Run}(s, \mu_\varepsilon, \chi) = \langle s_0, a_1, s_1, \dots, s_n, \dots \rangle$. Notice that for every $i \geq 0$ we have that $G(s_i) \geq G(s_{i+1})$ if $s_i \in S_{\text{Max}}$ and $G(s_i) = G(s_{i+1})$ if $s_i \in S_{\text{Min}}$. Hence $G(s_0), G(s_1), \dots$ is a non-increasing sequence. Since G is finite image, the sequence eventually becomes constant. Assume that for $i \geq N$ we have that $G(s_i) = g$. Now notice that for all $i \geq N$ we have that $B(s_i) \geq \pi(s_i, a_{i+1}) - g + B(s_{i+1})$ if $s_i \in S_{\text{Max}}$ and $B(s_i) \geq \pi(s_i, a_{i+1}) - g + B(s_{i+1}) - \varepsilon$

if $s_i \in S_{\text{Min}}$. Summing these equations sidewise from $i = N$ to $N + k$ we have that $B(s_N) \geq \sum_{i=N}^{N+k} \pi(s_i, a_{i+1}) - (k+1) \cdot g + B(s_{N+k+1}) - (k+1) \cdot \varepsilon$. Rearranging, we get

$$g \geq \frac{1}{k+1} \sum_{i=N}^{N+k} \pi(s_i, a_{i+1}) + \frac{1}{k+1} (B(s_{N+k+1}) - B(s_N)) - \varepsilon.$$

Hence

$$\begin{aligned} g &\geq \limsup_{k \rightarrow \infty} \frac{1}{k+1} \sum_{i=N}^{N+k} \pi(s_i, a_{i+1}) + \limsup_{k \rightarrow \infty} \frac{1}{k+1} (B(s_{N+k+1}) - B(s_N)) - \varepsilon \\ &= \limsup_{k \rightarrow \infty} \frac{1}{k} \sum_{i=0}^k \pi(s_i, a_{i+1}) - \varepsilon. \end{aligned}$$

Hence $G(s) + \varepsilon \geq \mathcal{A}_{\text{Min}}(\text{Run}(s, \mu_\varepsilon, \chi))$. Since χ is an arbitrary strategy in Σ_{Max} , we have $G(s) + \varepsilon \geq \sup_{\chi' \in \Sigma_{\text{Max}}} \mathcal{A}_{\text{Min}}(\text{Run}(s, \mu_\varepsilon, \chi'))$.

■ This part is analogous to the first part of the proof and is omitted.

The proof is now complete. ◀

2.2 Timed Automata

Priced Timed Game Arenas (PTGAs) extend classical timed automata [2] with a partition of the actions between two players Min and Max. Before we present the syntax and semantics of PTGAs, we need to introduce the concept of clock variables and related notions.

Clocks. Let \mathcal{X} be a finite set of *clocks*. A *clock valuation* on \mathcal{X} is a function $\nu : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ and we write $V(\mathcal{X})$ (or just V when \mathcal{X} is clear from the context) for the set of clock valuations. Abusing notation, we also treat a valuation ν as a point in $(\mathbb{R}_{\geq 0})^{|\mathcal{X}|}$. Let $\mathbf{0}$ denote the clock valuation that assigns 0 to all clocks. 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 $c \in \mathcal{X}$. For $C \subseteq \mathcal{X}$, we write $\nu[C := 0]$ for the valuation where $\nu[C := 0](c)$ equals 0 if $c \in C$ and $\nu(c)$ otherwise. For $X \subseteq V(\mathcal{X})$, 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 a $K \in \mathbb{N}$ such that for every $c \in \mathcal{X}$ we have $\nu(c) \leq K$.

Clock Constraints. A *clock constraint* over \mathcal{X} 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 $c, c' \in \mathcal{X}$, $i \in \mathbb{N}$, $i \leq K$, and $\bowtie \in \{<, >, =, \leq, \geq\}$. For $\nu \in V(\mathcal{X})$ and $K \in \mathbb{N}$, let $\text{CC}(\nu, K)$ 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 $\nu(x)$.

Regions and Zones. Every clock region is an equivalence class of the indistinguishability-by-clock-constraints relation. For a given set of clocks \mathcal{X} and upper bound $K \in \mathbb{N}$ on clock constraints, a *clock region* is a maximal set $\zeta \subseteq V(\mathcal{X})$ such that $\text{CC}(\nu, K) = \text{CC}(\nu', K)$ for all $\nu, \nu' \in \zeta$. For the set of clocks \mathcal{X} and upper bound K we write $\mathcal{R}(\mathcal{X}, K)$ for the corresponding finite set of clock regions. We write $[\nu]$ for the clock region of ν . A *clock zone* is a convex set of clock valuations that satisfies constraints of the form $\gamma ::= c_1 \bowtie k \mid c_1 - c_2 \bowtie k \mid \gamma \wedge \gamma$, $k \in \mathbb{N}$, $c_1, c_2 \in \mathcal{X}$ and $\bowtie \in \{\leq, <, =, >, \geq\}$. We write $\mathcal{Z}(\mathcal{X}, K)$ for the set of clock zones over the set of clocks \mathcal{X} and upper bound K . When \mathcal{X} and K are clear from the context we write \mathcal{R} and \mathcal{Z} for the set of regions and zones. In this paper we fix a positive integer K , and work with K -bounded clocks and clock constraints.

2.3 Priced Timed Game Arena: Syntax and Semantics

► **Definition 3.** A priced timed game arena is a tuple $\mathsf{T} = (L_{\text{Min}}, L_{\text{Max}}, \text{Act}, \mathcal{X}, \text{Inv}, E, \rho, \delta, p)$ where L_{Min} and L_{Max} are sets of *locations* controlled by Player Min and Player Max and we write $L = L_{\text{Min}} \cup L_{\text{Max}}$; Act is a finite set of *actions*; \mathcal{X} is a finite set of *clocks*; $\text{Inv} : L \rightarrow \mathcal{Z}$ is an *invariant condition*; $E : L \times \text{Act} \rightarrow \mathcal{Z}$ is an *action enabledness function*; $\rho : \text{Act} \rightarrow 2^{\mathcal{C}}$ is a *clock reset function*; $\delta : L \times \text{Act} \rightarrow L$ is a *transition function*; and $p : L \cup L \times \text{Act} \rightarrow \mathbb{R}$ is a *price information function*. A PTGA is binary-priced when $p(\ell) \in \{0, 1\}$ for all $\ell \in L$.

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{X} , Inv , Act , E , ρ , δ and p . We draw the states of Min players as circles, while states of Max player as boxes.

Let $\mathsf{T} = (L_{\text{Min}}, L_{\text{Max}}, \text{Act}, \mathcal{X}, \text{Inv}, E, \rho, \delta, p)$ be a PTGA. A *configuration* of a PTGA is a pair (ℓ, ν) , where ℓ is a location and ν a clock valuation such that $\nu \in \text{Inv}(\ell)$. For any $t \in \mathbb{R}_{\geq 0}$, 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 $\text{Inv}(\ell)$ 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 transition relation δ and the reset function ρ , i.e. the clocks in $\rho(a)$ are reset and we move to the location $\delta(\ell, a)$.

A game on a PTGA starts in an *initial configuration* $(\ell, \nu) \in L \times V$ and players Min and Max construct an infinite play by taking turns to choose available timed actions (t, a) whenever the current location is controlled by them and the price $p(\ell) \cdot t + p(\ell, a)$ is paid to the Max by player Min. Formally, PTGA semantics is given as a game arena.

► **Definition 4 (PTGA Semantics).** Let $\mathsf{T} = (L_{\text{Min}}, L_{\text{Max}}, \text{Act}, \mathcal{X}, \text{Inv}, E, \rho, \delta, p)$ be a PTGA. The semantics of T is given by game arena $\llbracket \mathsf{T} \rrbracket = (S, S_{\text{Min}}, S_{\text{Max}}, A, T, \pi)$ where

- $S \subseteq L \times V$ is the set of states such that $(\ell, \nu) \in S$ if and only if $\nu \in \text{Inv}(\ell)$;
- $(\ell, \nu) \in S_{\text{Min}}$ (or $(\ell, \nu) \in S_{\text{Max}}$) if $(\ell, \nu) \in S$ and $\ell \in L_{\text{Min}}$ (or $\ell \in L_{\text{Max}}$, respectively).
- $A = \mathbb{R}_{\geq 0} \times \text{Act}$ is the set of *timed actions*;
- $T : S \times A \rightarrow S$ is the transition function such that for $(\ell, \nu) \in S$ and $(t, a) \in A$, we have $T((\ell, \nu), (t, a)) = (\ell', \nu')$ if and only if
 - $\nu + t' \in \text{Inv}(\ell)$ for all $t' \in [0, t]$; $\nu + t \in E(\ell, a)$; $(\ell', \nu') \in S$, $\delta(\ell, a) = \ell'$, $(\nu + t)[\rho(a) := 0] = \nu'$.
- $\pi : S \times A \rightarrow \mathbb{R}$ is the reward function where $\pi((\ell, \nu), (t, a)) = p(\ell) \cdot t + p(\ell, a)$.

We are interested in the mean-payoff decision problem for timed automata T that asks to decide whether the value of the mean-payoff game for a given state is below a given budget. For a PTGA T and budget $r \in \mathbb{R}$, we write $\text{MPG}(\mathsf{T}, r)$ for the r -mean payoff decision problem that asks whether the value of the game at the state $(\ell, \mathbf{0})$ is smaller than r . The following theorem summarizes the key contribution of this paper.

► **Theorem 5.** *The decision problem $\text{MPG}(\mathsf{T}, r)$ for binary-priced timed automata T is undecidable for automata with three clocks, and decidable for automata with one clock.*

3 Boundary Region Graph Abstraction

In this section we introduce an abstraction of priced timed games called the boundary region abstraction (that generalizes classical corner-point abstraction [7]), and characterize conditions under which a solution of optimality equations for the boundary region abstraction can be lifted to a solution of optimality equations for timed automata. Observe that in order

to keep our result as general as possible, we present the abstraction and corresponding results for timed automata with an arbitrary number of clocks. In the following section, we show that the required conditions hold for the case of one-clock binary-priced timed automata.

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

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

Intuition for the Boundary Region Graph (BRG). Recall that K is an upper bound on clock values and let $\llbracket K \rrbracket_{\mathbb{N}} = \{0, 1, \dots, K\}$. For any $\nu \in V, b \in \llbracket K \rrbracket_{\mathbb{N}}$ and $c \in \mathcal{X}$, we define $\text{time}(\nu, (b, c)) \stackrel{\text{def}}{=} b - \nu(c)$ if $\nu(c) \leq b$, and $\text{time}(\nu, (b, c)) \stackrel{\text{def}}{=} 0$ if $\nu(c) > b$. Intuitively, $\text{time}(\nu, (b, c))$ returns the amount of time that must elapse in ν before the clock c reaches the integer value b . Observe that, for any $\zeta' \in \mathcal{R}_{\text{Thin}}$, there exists $b \in \llbracket K \rrbracket_{\mathbb{N}}$ and $c \in \mathcal{X}$, such that $\nu \in \zeta$ implies $(\nu + (b - \nu(c))) \in \zeta'$ for all $\zeta \in \mathcal{R}$ in the past of ζ' and write $\zeta \rightarrow_{b,c} \zeta'$. The boundary region abstraction is motivated by the following. Consider $a \in \text{Act}, (\ell, \nu)$ and $\zeta \xrightarrow{*} \zeta'$ such that $\nu \in \zeta, [\zeta, \zeta'] \subseteq \text{Inv}(\ell)$ and $\nu' \in E(\ell, a)$. (For illustration, see Figure 2 in the appendix in [15]).

- If $\zeta' \in \mathcal{R}_{\text{Thick}}$, then there are infinitely many $t \in \mathbb{R}_{\geq 0}$ such that $\nu + t \in \zeta'$. However, amongst all such t 's, for one of the boundaries of ζ' , the closer $\nu + t$ is to this boundary, the ‘better’ the timed action (t, a) becomes for a player’s objective. However, since ζ' is a thick region, the set $\{t \in \mathbb{R}_{\geq 0} \mid \nu + t \in \zeta'\}$ is an open interval, and hence does not contain its boundary values. Let the closest boundary of ζ' from ν be defined by the hyperplane $c = b_{\text{inf}}$ and the farthest boundary of ζ' from ν be defined by the hyperplane $c = b_{\text{sup}}$. $b_{\text{inf}}, b_{\text{sup}} \in \mathbb{N}$ are such that $b_{\text{inf}} - \nu(c)$ ($b_{\text{sup}} - \nu(c)$) is the infimum (supremum) of the time spent to reach the lower (upper) boundary of region ζ' . Let the zones that correspond to these boundaries be denoted by ζ'_{inf} and ζ'_{sup} respectively. Then $\zeta \rightarrow_{b_{\text{inf}}, c} \zeta'_{\text{inf}} \rightarrow \zeta'$ and $\zeta \rightarrow_{b_{\text{sup}}, c} \zeta'_{\text{sup}} \leftarrow \zeta'$. In the boundary region abstraction we include these ‘best’ timed actions through $(b_{\text{inf}}, c, a, \zeta')$ and $(b_{\text{sup}}, c, a, \zeta')$.
- If $\zeta' \in \mathcal{R}_{\text{Thin}}$, then there exists a unique $t \in \mathbb{R}_{\geq 0}$ such that $\nu + t \in \zeta'$. Moreover since ζ' is a thin region, there exists a clock $c \in \mathcal{C}$ and a number $b \in \mathbb{N}$ such that $\zeta \rightarrow_{b,c} \zeta'$ and $t = b - \nu(c)$. In the boundary region abstraction we summarise this ‘best’ timed action from region ζ via region ζ' through the action (b, c, a, ζ') .

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

► **Definition 6.** For a priced timed game arena $\mathsf{T} = (L_{\text{Min}}, L_{\text{Max}}, \text{Act}, \mathcal{X}, \text{Inv}, E, \rho, \delta, p)$ the boundary region abstraction of T is given by the game arena $\widehat{\mathsf{T}} = (\widehat{S}, \widehat{S}_{\text{Min}}, \widehat{S}_{\text{Max}}, \widehat{A}, \widehat{T}, \widehat{\pi})$

- $\widehat{S} \subseteq L \times V \times \mathcal{R}$ is the set of states such that $(\ell, \nu, \zeta) \in \widehat{S}$ if and only if $\zeta \subseteq \text{Inv}(\ell)$ and $\nu \in \bar{\zeta}$ (recall that $\bar{\zeta}$ denotes the closure of ζ);
- $(\ell, \nu, \zeta) \in \widehat{S}_{\text{Min}}$ (or $(\ell, \nu, \zeta) \in \widehat{S}_{\text{Max}}$) if $(\ell, \nu, \zeta) \in \widehat{S}$ and $\ell \in L_{\text{Min}}$ (or $\ell \in L_{\text{Max}}$, resp.);
- $\widehat{A} = (\llbracket K \rrbracket_{\mathbb{N}} \times \mathcal{X} \times \text{Act} \times \mathcal{R})$ is the set of actions;
- For $\hat{s} = (\ell, \nu, \zeta) \in \widehat{S}$ and $\alpha = (b_{\alpha}, c_{\alpha}, a_{\alpha}, \zeta_{\alpha}) \in \widehat{A}$, function $\widehat{T}(\hat{s}, \alpha)$ is defined if $[\zeta, \zeta_{\alpha}] \subseteq \text{Inv}(\ell)$ and $\zeta_{\alpha} \subseteq E(\ell, a_{\alpha})$ and it equals $(\ell', \nu', \zeta') \in \widehat{S}$ where $\delta(\ell, a_{\alpha}) = \ell', \nu_{\alpha}[C := 0] = \nu'$ and

- $\zeta_\alpha[C:=0] = \zeta'$ with $\nu_\alpha = \nu + \text{time}(\nu, (b_\alpha, c_\alpha))$ and one of the following conditions holds:
 $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_\alpha$; $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_{\text{inf}} \rightarrow \zeta_\alpha$ for some $\zeta_{\text{inf}} \in \mathcal{R}$; $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_{\text{sup}} \leftarrow \zeta_\alpha$ for some $\zeta_{\text{sup}} \in \mathcal{R}$;
 ■ for $(\ell, \nu, \zeta) \in \widehat{S}$ and $(b_\alpha, c_\alpha, a_\alpha, \zeta_\alpha) \in \widehat{A}$ the reward function $\widehat{\pi}$ is given by:
 $\widehat{\pi}((\ell, \nu, \zeta), (b_\alpha, c_\alpha, a_\alpha, \zeta_\alpha)) = p(\ell, a_\alpha) + p(\ell) \cdot (b_\alpha - \nu(c_\alpha))$

Although the boundary region abstraction is not a finite game arena, every state has only finitely many time successors (the boundaries of the regions) and for a fixed initial state we can restrict attention to a finite game arena due to the following observation.

► **Lemma 7** ([23]). *Let T be a priced timed game arena and $\widehat{\mathsf{T}}$ the corresponding BRA. For any state of $\widehat{\mathsf{T}}$, its reachable sub-graph is finite and can be constructed in time exponential in the size of T when T has more than one clock. For one clock T , the reachable sub-graph of $\widehat{\mathsf{T}}$ can be constructed in time polynomial in the size of T . Moreover, the reachable sub-graph from the initial location and clock valuation is precisely the corner-point abstraction.*

3.1 Reduction to Boundary Region Abstraction

In what follows, unless specified otherwise, we fix a PTGA $\mathsf{T} = (L_{\text{Min}}, L_{\text{Max}}, \text{Act}, \mathcal{X}, \text{Inv}, E, \rho, \delta, p)$ with semantics $\llbracket \mathsf{T} \rrbracket = (S, S_{\text{Min}}, S_{\text{Max}}, A, T, \pi)$ and BRA $\widehat{\mathsf{T}} = (\widehat{S}, \widehat{S}_{\text{Min}}, \widehat{S}_{\text{Max}}, \widehat{A}, \widehat{T}, \widehat{\pi})$. Let $G : \widehat{S} \rightarrow \mathbb{R}$ and $B : \widehat{S} \rightarrow \mathbb{R}$ be such that $(G, B) \models \text{Opt}(\widehat{\mathsf{T}})$, i.e. for every $\hat{s} \in \widehat{S}$ we have that

$$G(\hat{s}) = \begin{cases} \max_{\alpha \in \widehat{A}(\hat{s})} \{G(\hat{s}') : \hat{s} \xrightarrow{\alpha} \hat{s}'\} & \text{if } \hat{s} \in \widehat{S}_{\text{Max}} \\ \min_{\alpha \in \widehat{A}(\hat{s})} \{G(\hat{s}') : \hat{s} \xrightarrow{\alpha} \hat{s}'\} & \text{if } \hat{s} \in \widehat{S}_{\text{Min}}. \end{cases}$$

$$B(\hat{s}) = \begin{cases} \max_{\alpha \in \widehat{A}(\hat{s})} \{\pi(\hat{s}, \alpha) - G(\hat{s}) + B(\hat{s}') : \hat{s} \xrightarrow{\alpha} \hat{s}' \text{ and } G(\hat{s}) = G(\hat{s}')\} & \text{if } \hat{s} \in \widehat{S}_{\text{Max}} \\ \min_{\alpha \in \widehat{A}(\hat{s})} \{\pi(\hat{s}, \alpha) - G(\hat{s}) + B(\hat{s}') : \hat{s} \xrightarrow{\alpha} \hat{s}' \text{ and } G(\hat{s}) = G(\hat{s}')\} & \text{if } \hat{s} \in \widehat{S}_{\text{Min}}. \end{cases}$$

For a function $F : \widehat{S} \rightarrow \mathbb{R}$ we define a function $F^\boxplus : S \rightarrow \mathbb{R}$ as $(\ell, \nu) \mapsto F(\ell, \nu, [\nu])$. In this section we show under what conditions we can lift a solution (G, B) of optimality equations of BRA to (G^\boxplus, B^\boxplus) for priced timed game arena. Given a set of valuations $X \subseteq V$, a function $f : X \rightarrow \mathbb{R}_{\geq 0}$ is *affine* if for any valuations $\nu_x, \nu_y \in X$ we have that for all $\lambda \in [0, 1]$, $f(\lambda\nu_x + (1-\lambda)\nu_y) = \lambda f(\nu_x) + (1-\lambda)f(\nu_y)$. We say that a function $f : \widehat{S} \rightarrow \mathbb{R}_{\geq 0}$ is *regionally affine* if $f(\ell, \cdot, \zeta)$ is affine over a region for all $\ell \in L$ and $\zeta \in \mathcal{R}$, and f is *regionally constant* if $f(\ell, \cdot, \zeta)$ is constant over a region for all $\ell \in L$ and $\zeta \in \mathcal{R}$. Some properties of affine functions that are useful in the proof of the key lemma are given in Lemma 8.

► **Lemma 8.** *Let $X \subseteq V$ and $Y \subseteq \mathbb{R}_{\geq 0}$ be convex sets. Let $f : X \rightarrow \mathbb{R}$ and $w : X \times Y \rightarrow \mathbb{R}$ be affine functions. Then for $C \subseteq \mathcal{X}$ we have that $\phi_C(\nu, t) = w(\nu, t) + f((\nu + t)[C:=0])$ is also an affine function, and $\inf_{t_1 < t < t_2} \phi_C(\nu, t) = \min\{\overline{\phi}_C(\nu, t_1), \overline{\phi}_C(\nu, t_2)\}$ and $\sup_{t_1 < t < t_2} \phi_C(\nu, t) = \max\{\overline{\phi}_C(\nu, t_1), \overline{\phi}_C(\nu, t_2)\}$, $\overline{\phi}$ is the unique continuous closure of ϕ .*

► **Theorem 9.** *Let $G : \widehat{S} \rightarrow \mathbb{R}$ and $B : \widehat{S} \rightarrow \mathbb{R}$ are such that $(G, B) \models \text{Opt}(\widehat{\mathsf{T}})$ and G is regionally constant and B is regionally affine, then $(G^\boxplus, B^\boxplus) \models \text{Opt}(\mathsf{T})$.*

Proof. We need to show that $(G^\boxplus, B^\boxplus) \models \text{Opt}(\mathsf{T})$, i.e. for every

$$G^\boxplus(s) = \begin{cases} \sup_{(t,a) \in A(s)} \{G^\boxplus(s') : s \xrightarrow{(t,a)} s'\} & \text{if } s \in S_{\text{Max}} \\ \inf_{(t,a) \in A(s)} \{G^\boxplus(s') : s \xrightarrow{(t,a)} s'\} & \text{if } s \in S_{\text{Min}}. \end{cases}$$

$$B^\boxplus(s) = \begin{cases} \sup_{(t,a) \in A(s)} \{\pi(s, (t,a)) - G^\boxplus(s) + B^\boxplus(s') : s \xrightarrow{(t,a)} s' \text{ and } G^\boxplus(s) = G^\boxplus(s')\} & \text{if } s \in S_{\text{Max}} \\ \inf_{(t,a) \in A(s)} \{\pi(s, (t,a)) - G^\boxplus(s) + B^\boxplus(s') : s \xrightarrow{(t,a)} s' \text{ and } G^\boxplus(s) = G^\boxplus(s')\} & \text{if } s \in S_{\text{Min}}. \end{cases}$$

Consider the case when $s = (\ell, \nu) \in S_{\text{Min}}$ and consider the right side of the gain equations.

$$\begin{aligned}
& \inf_{(t,a) \in A(s)} \{G^{\boxplus}(s') : s \xrightarrow{(t,a)} s'\} \\
&= \min_{\substack{\zeta'' : [\nu] \rightarrow^* \zeta'' \\ [\zeta, \zeta''] \in \text{Inv}(\ell)}} \min_{a \in \text{Act}} \inf_{\substack{t: \\ \nu+t \in \zeta''}} \{G(\delta(\ell, a), (\nu+t)[\rho(a):=0], [(\nu+t)][\rho(a):=0])\} \\
&= \min_{\alpha \in \hat{A}(\ell, \nu, [\nu])} \{G(\ell', \nu', \zeta') : (\ell, \nu, \zeta) \xrightarrow{\alpha} (\ell', \nu', \zeta')\} = G(\ell, \nu, [\nu]) = G^{\boxplus}(\ell, \nu).
\end{aligned}$$

The first equality holds since $(G, B) \models \text{Opt}(\hat{\mathbb{T}})$. The second equality follows since G is regionally constant and hence it suffices to consider the delay $\text{time}(\nu, (b, c))$ that corresponds to either left or right boundary of the region ζ'' , i.e. for fixed ν, ζ'' and $a \in \text{Act}$ we have that $\inf_{\substack{t: \\ \nu+t \in \zeta''}} \{G(\ell', (\nu+t)[\rho(a):=0], \zeta')\} = G(\ell', \nu_\alpha[C:=0], \zeta')$ where $\nu_\alpha = \nu + \text{time}(\nu, (b_\alpha, c_\alpha))$, $\zeta''[C:=0] = \zeta'$ with $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta''$ if ζ'' is thin, and $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_{\text{inf}} \rightarrow \zeta''$ for some $\zeta_{\text{inf}} \in \mathcal{R}$ if ζ'' is thick. Similarly, for the bias equations, we need to show:

$$\begin{aligned}
& \inf_{\substack{t: \\ \nu+t \in \zeta''}} \{\pi((\ell, \nu), (t, a)) - G(\ell, \nu) + B(\ell', (\nu+t)[\rho(a):=0], \zeta')\} \\
&= \pi((\ell, \nu, [\nu]), (\text{time}(\nu, (b_\alpha, c_\alpha)))) - G(\ell, \nu, [\nu]) + B(\ell', \nu_\alpha[C:=0], \zeta')
\end{aligned}$$

where $\nu_\alpha = \nu + \text{time}(\nu, (b_\alpha, c_\alpha))$, $\zeta''[C:=0] = \zeta'$ with $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta''$ if ζ'' is thin; and $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_{\text{inf}} \rightarrow \zeta''$ for some $\zeta_{\text{inf}} \in \mathcal{R}$ or $\zeta \rightarrow_{b_\alpha, c_\alpha} \zeta_{\text{sup}} \rightarrow \zeta''$ for some $\zeta_{\text{sup}} \in \mathcal{R}$ if ζ'' is thick. Given B is regionally affine (and hence linear in t) and the price function is linear in t , the whole expression $\pi((\ell, \nu), (t, a)) - G(\ell, \nu) + B(\ell', (\nu+t)[\rho(a):=0], \zeta')$ is linear in t and from Lemma 8 it attains its infimum or supremum on either boundary of the region. \blacktriangleleft

4 Decidability for One Clock Binary-priced PTGA

Given the undecidability with 3 or more clocks, we focus on one clock PTGA. We provide a strategy improvement algorithm to compute a solution $G : \hat{S} \rightarrow \mathbb{R}$ and $B : \hat{S} \rightarrow \mathbb{R}$ of the optimality equations, i.e. $(G, B) \models \text{Opt}(\hat{\mathbb{T}})$ for the BRA $\hat{\mathbb{T}} = (\hat{S}, \hat{S}_{\text{Min}}, \hat{S}_{\text{Max}}, \hat{A}, \hat{T}, \hat{\pi})$ of one-clock binary-priced PTGAs with certain “integral payoff” restriction. Further, we show that for one clock binary-priced integral-payoff PTGA, the solution of optimality equations of corresponding BRG is such that the gains are regionally constant and biases are regionally affine. Hence by Theorem 9, the algorithm can be applied to solve mean-payoff games for one-clock binary-priced integral-payoff PTGAs. We also show how to lift the integral-payoff restriction to recover decidability for one-clock binary-priced PTGA.

Regionally constant positional strategies. Standard strategy improvement algorithms iterate over a finite set of strategies such that the value of the subgame at each iteration gets strictly improved. However, since there are infinitely many positional strategies in a boundary region abstraction, we focus on “regionally constant” positional strategies (RCPSs). We say that a positional strategy $\mu : \hat{S} \rightarrow \hat{A}$ of player Min is regionally-constant if for all $(\ell, \nu, \zeta), (\ell, \nu', \zeta) \in \hat{S}_{\text{Min}}$ we have that $[\nu]=[\nu']$ implies that $\mu(\ell, \nu, \zeta) = \mu(\ell, \nu', \zeta)$. We similarly define RCPSs for player Max. In other words, in an RCPS a player chooses the same boundary action for every valuation of a region – as a side-result we show that optimal strategies for both players have this form. Observe that there are finitely many RCPSs for both players. We write $\hat{\Pi}_{\text{Min}}$ and $\hat{\Pi}_{\text{Max}}$ for the set of RCPSs for player Min and player Max, respectively. For a BRA $\hat{\mathbb{T}}$, $\chi \in \hat{\Pi}_{\text{Max}}$, and $\mu \in \hat{\Pi}_{\text{Min}}$ we write $\hat{\mathbb{T}}(\chi)$ and $\hat{\mathbb{T}}(\mu)$ for the “one-player” game on the sub-graph of BRAs where the strategies of player Max and Min

Algorithm 1: COMPUTEVALUEZEROPLAYER(\mathbb{T}, μ, χ)

```

1 Consider  $\widehat{\mathbb{T}}(\mu, \chi)$  as a (single successor) weighted graph  $\mathcal{G} = (V, E, w)$  where
  ■  $V = L \times \mathcal{R} \times \mathcal{R}$  (with an order  $\preceq$ ) and  $E \subseteq V \times \widehat{A} \times V$ 
  ■  $(v_1, \alpha, v_2) \in E$  if  $v_1 = (\ell_1, \zeta_1, \zeta'_1)$ ,  $v_2 = (\ell_2, \zeta_2, \zeta'_2)$ , and  $\mu(\ell_1, \nu_1, \zeta'_1) = \alpha$  (or
     $\chi(\ell_1, \nu_1, \zeta'_1) = \alpha$ ) for all  $\nu_1 \in \zeta_1$  and  $(\ell_1, \nu_1, \zeta'_1) \xrightarrow{\alpha} (\ell_2, \nu_2, \zeta'_2)$  for some  $\nu_2 \in \zeta_2$ .
  ■  $w(v_1, \alpha, v_2)$  is the expression  $\nu \mapsto b_\alpha - \nu(c_\alpha)$ ;
  for every cycle  $C$  of  $\mathcal{G}$  do
    Let  $\text{Reach}(C)$  be set of vertices that reach  $C$ ;
    Let  $\gamma$  be the average weight of the cycle ( $w$  is constant on cycles);
    For every vertex  $V$  in  $\text{Reach}(C)$  set  $G(V) = \gamma$  and  $B(V) = \perp$ ;
    For the smallest  $\preceq$ -vertex  $V_*$  in  $C$ . Set  $B(V_*) = 0$ ;
    while there is  $V' \in \text{Reach}(C)$  with  $B(V') = \perp$  do
      Let  $(V', \alpha, V'') \in E$  with  $B(V'') \neq \perp$ ;
       $B(V') := \nu \mapsto (w(V', \alpha, V''))(\nu) - G + B(V'')$ ;
  return  $(G, B)$ ;

```

have been fixed to RCPSs χ and μ , respectively. Similarly we define the zero-player game $\widehat{\mathbb{T}}(\mu, \chi)$ where strategies of both players are fixed to RCPSs μ and χ .

Let $\widehat{\mathbb{T}}(\chi, \mu)$ be a zero-player game on the subgraph where strategies of player Max (and Min) are fixed to RCPSs χ (and μ). Observe that for $\widehat{\mathbb{T}}(\mu, \chi)$ the unique runs originating from states $\hat{s}_0 = (\ell, \nu, \zeta)$ and $\hat{s}'_0 = (\ell, \nu', \zeta)$ with $[\nu] = [\nu']$ follow the same “lasso” after one step, i.e. the unique runs $\hat{s}_0 \xrightarrow{\alpha_1} \hat{s}_1 \cdots \hat{s}_k \xrightarrow{\alpha_{k+1}} \cdots \hat{s}_{k+N-1} \xrightarrow{\alpha_{k+N}} \hat{s}_k^*$ and $\hat{s}'_0 \xrightarrow{\alpha_1} \hat{s}'_1 \cdots \hat{s}'_k \xrightarrow{\alpha_{k+1}} \cdots \hat{s}'_{k+N-1} \xrightarrow{\alpha_{k+N}} \hat{s}'_k^*$ are such that for $\hat{s}_i = (\ell_i, \nu_i, \zeta_i)$ and $\hat{s}'_i = (\ell'_i, \nu'_i, \zeta'_i)$ we have that $\ell_i = \ell'_i$, $\zeta_i = \zeta'_i$ and $\nu_i = \nu'_i$ for all $i \in [1, k+N-1]$. This is so because for one-clock timed automata the successors of the states $\hat{s}_0 = (\ell, \nu, \zeta)$ and $\hat{s}'_0 = (\ell, \nu', \zeta)$ for action $\alpha_1 = (b, c, a, \zeta')$ is the same (ℓ'', ν'', ζ'') where $\nu''(c) = \nu(c) + (b - \nu(c)) = b = \nu'(c) + (b - \nu'(c))$ if $c \notin \rho(a)$ and $\nu''(c) = 0$ otherwise. Consider the optimality equations (See Appendix C.3 in [15]) for the lasso. Observe that the gain for the states $\hat{s}_0, \dots, \hat{s}_{k+N-1}$ is the same, and let's call it g . If we add the bias equations side-wise for the cycle, we get $g = \frac{1}{N} \sum_{i=0}^{N-1} \pi(\hat{s}_{k+i}, \alpha_{k+i+1})$. It follows from the previous observation that the gains are regionally constant.

Integral Payoff PTGA. The gain in a zero-player game, $\widehat{\mathbb{T}}(\chi, \mu)$, although regionally-constant, may not be a whole number. We say that a PTGA is integral-payoff if for every pair $(\mu, \chi) \in \widehat{\Pi}_{\text{Min}} \times \widehat{\Pi}_{\text{Max}}$ of RCPSs the gain as defined above is a whole number. Observe that the denominator in the gains correspond to the number of edges in a simple cycle of the BRA $\widehat{\mathbb{T}}$. If there are N simple cycles in the region graph of length n_1, n_2, \dots, n_N , then let \mathcal{L} be the least-common multiple of n_1, n_2, \dots, n_N . We multiply the constants appearing in the guards and invariants of the original PTGA \mathbb{T} by \mathcal{L} to obtain a PTGA $\Upsilon_{\mathbb{T}}$. It is easy to observe that mean-payoff of any state in \mathbb{T} is the mean-payoff in $\Upsilon_{\mathbb{T}}$ divided by \mathcal{L} . For notational convenience, we assume that the given PTGA is an integral-payoff PTGA and hence for RCPS strategy profile (μ, χ) the gain is regionally constant and integral.

4.1 Strategy Improvement Algorithm for Binary-Priced PTGA

Let \mathbb{T} be a one-clock integral-payoff binary-priced PTGA \mathbb{T} and $\widehat{\mathbb{T}}$ be its boundary region graph. For a given RCPS profile $(\mu, \chi) \in \widehat{\Pi}_{\text{Min}} \times \widehat{\Pi}_{\text{Max}}$, Algorithm 1 computes the solution for the optimality equations $\text{Opt}(\mathbb{T}(\mu, \chi))$. This algorithm considers $\widehat{\mathbb{T}}(\mu, \chi)$ as a graph whose

Algorithm 2: COMPUTEVALUE TWOPLAYER(\mathbb{T})

```

1 Choose an arbitrary regionally constant positional strategy  $\chi' \in \Pi_{\text{Max}}$ ;
2 repeat
3    $\chi := \chi'$ ;
4   Choose an arbitrary regionally constant positional strategy  $\mu' \in \Pi_{\text{Min}}$ ;
5   repeat
6      $\mu := \mu'$ ;
7      $(G, B) := \text{COMPUTEVALUEZEROPLAYER}(\mathbb{T}, \mu, \chi)$ ;
8      $\mu' := \text{IMPROVEMINSTRATEGY}(\mathbb{T}, \mu, G, B)$ ;
9   until  $\mu = \mu'$ ;
10   $\chi' := \text{IMPROVEMAXSTRATEGY}(\mathbb{T}, \chi, G, B)$ ;
11 until  $\chi = \chi'$ ;
12 return  $(G, B)$ ;

```

vertices are “regions” $(\ell, [\nu], \zeta)$ corresponding to state $(\ell, \nu, \zeta) \in \widehat{S}$ of the boundary region graph, edges are boundary actions between them determined by the regionally constant strategy profile, and weight of an edge is the time function associated with the boundary action. Observe that every cycle in this graph will have constant weight on the edges since taking boundary actions in a loop will require going from an integral valuation to another integral valuation, and the average cost of such a cycle can be easily computed.

Also observe that, not unlike standard convention [21], our algorithm chooses a vertex in a cycle arbitrarily and fixes the bias of all of the states in that vertex to 0. This is possible since optimality equations over a cycle are underdetermined, and we exploit this flexibility to achieve solution to biases in a particularly “simple” structure. We say that a function $f : \widehat{S} \rightarrow \mathbb{R}_{\geq 0}$ is regionally simple [3] if for all $\ell \in L$, $\zeta, \zeta' \in \mathcal{R}$ either i) there exists a $d \in \mathbb{N}$ such that $f(\ell, \nu, \zeta') = d$ for all $\nu \in \zeta$; or ii) there exists $d \in \mathbb{N}$ and $c \in \mathcal{X}$ such that $f(\ell, \nu, \zeta') = d - \nu(c)$ for all $\nu \in \zeta$. Key properties of regionally simple functions (Lemma 20 in Appendix C.2 in [15]) include that they are also regionally affine, closed under minimum and maximum, and if $B : \widehat{S} \rightarrow \mathbb{R}$ is a regionally simple function and $G : \widehat{S} \rightarrow \mathbb{N}$ is a regionally constant function, then $\hat{s} \mapsto \pi(\hat{s}, \alpha) - G(\hat{s}) + B(\hat{s}')$, with $\hat{s} \xrightarrow{\alpha} \hat{s}'$, is a regionally simple function. Using these properties and induction on the distance to \preceq -minimal element in the reachable cycle, we prove the correctness and following property of Algorithm 1.

► **Lemma 10.** *Algorithm 1 computes solution of optimality equations $(G, B) \models \text{Opt}(\widehat{\mathbb{T}}(\mu, \chi))$ for $\mu \in \widehat{\Pi}_{\text{Min}}$ and $\chi \in \widehat{\Pi}_{\text{Max}}$. Moreover, G is regionally constant and B is regionally simple.*

The strategy improvement algorithm to solve optimality equations is given as Algorithm 2. It begins by choosing an arbitrary regionally constant positional strategy χ' and at every iteration of the loop (2–11) the algorithm computes (5–9) the value (G, B) of the current RCPS χ and based on the value, the function `IMPROVEMAXSTRATEGY` returns an improved strategy by picking boundary action that lexicographically maximizes gain and bias respecting the policy that switches a decision only for a strict improvement. We formally define the function `IMPROVEMAXSTRATEGY` as follows: for $\chi \in \widehat{\Sigma}_{\text{Max}}$, $G : \widehat{S} \rightarrow \mathbb{R}$, and $B : \widehat{S} \rightarrow \mathbb{R}$ we let strategy `IMPROVEMAXSTRATEGY`(\mathbb{T}, χ, G, B) be such that for all $\hat{s} \in \widehat{S}_{\text{Max}}$ we have

$$\text{IMPROVEMAXSTRATEGY}(\mathbb{T}, \chi, G, B)(\hat{s}) = \begin{cases} \chi(\hat{s}) & \text{if } \chi(\hat{s}) \in M^*(\hat{s}, G, B) \\ \text{Choose}(M^*(\hat{s}, G, B)) & \text{Otherwise.} \end{cases}$$

where $M^*(\hat{s}, G, B) = \text{argmax}_{\alpha \in \widehat{A}}^{\text{lex}} \{(G(\hat{s}'), \pi(\hat{s}, \alpha) - G(\hat{s}) + B(\hat{s}')) : \hat{s} \xrightarrow{\alpha} \hat{s}'\}$ and `Choose` picks an arbitrary element from a set. `IMPROVEMAXSTRATEGY` satisfies the following.

► **Lemma 11.** *If $\chi \in \widehat{\Pi}_{Max}$, G is regionally constant, and B is regionally simple, then function $\text{IMPROVEMAXSTRATEGY}(\mathbb{T}, \chi, G, B)$ returns a regionally constant positional strategy.*

The lines (5–9) compute the value of the strategy χ of Player Max via a strategy improvement algorithm. This sub-algorithm works by starting with an arbitrary strategy of Player Min and computing the value (G, B) of the zero-player PTGA $\widehat{\mathbb{T}}(\mu, \chi)$. Based on the value, the function $\text{IMPROVEMINSTRATEGY}$ returns an improved strategy of Min. The function $\text{IMPROVEMINSTRATEGY}$ is defined as a dual of the function $\text{IMPROVEMAXSTRATEGY}$ where χ is replaced by μ and argmax by argmin . $\text{IMPROVEMINSTRATEGY}$ satisfies the following.

► **Lemma 12.** *If $\mu \in \widehat{\Pi}_{Min}$, G is regionally constant, and B is regionally simple, then function $\text{IMPROVEMINSTRATEGY}(\mathbb{T}, \mu, G, B)$ returns a regionally constant positional strategy.*

It follows from Lemma 11 and Lemma 12 that at every iteration of the strategy improvement the strategies μ and χ are RCPSs. Together with finiteness of the set of RCPSs and strict improvement at every step (see Lemmas 21 and 22 in [15] for the formal statements), we get following result.

► **Theorem 13.** *Algorithm 2 computes solution of optimality equations $(G, B) \models \text{Opt}(\widehat{\mathbb{T}})$ for integral payoff PTGA \mathbb{T} . Moreover, G is regionally constant and B is regionally affine.*

This theorem – together with Theorem 9 and Theorem 2 – gives a proof of decidability for mean-payoff games for integral-payoff binary-priced one-clock timed automata.

5 Undecidability Results

► **Theorem 14.** *The mean-payoff problem $\text{MPG}(\mathbb{T}, r)$ is undecidable for PTGA \mathbb{T} with 3 clocks having location-wise price-rates $\pi(\ell) \in \{0, 1, -1\}$ for all $\ell \in L$ and $r = 0$. Moreover, it is undecidable for binary-priced \mathbb{T} with 3 clocks and $r > 0$.*

Proof. We first show the undecidability result of the mean-payoff problem $\text{MPG}(\mathbb{T}, 0)$ with location prices $\{1, 0, -1\}$ and no edge prices. We prove the result by reducing the non-halting problem of 2 counter machines. Our reduction uses a PTGA with 3 clocks x_1, x_2, x_3 , location prices $\{1, 0, -1\}$, and no edge prices. Each counter machine instruction (increment, decrement, zero check) is specified using a PTGA module. The main invariant in our reduction is that on entry into any module, we have $x_1 = \frac{1}{5c_1 7c_2}$, $x_2 = 0$ and $x_3 = 0$, where c_1, c_2 are the values of counters C_1, C_2 . We outline the construction for the decrement instruction of counter C_1 in Figure 2. For conciseness, we present here modules using arbitrary location prices. However, we can redraw these with extra locations and edges using only the location prices from $\{1, 0, -1\}$ as shown for WD_1^1 in Figure 1.

The role of the Min player is to faithfully simulate the two counter machine, by choosing appropriate delays to adjust the clocks to reflect changes in counter values. Player Max will have the opportunity to verify that player Min did not cheat while simulating the machine.

We enter location ℓ_k with $x_1 = \frac{1}{5c_1 7c_2}$, $x_2 = 0$ and $x_3 = 0$. Let's denote by x_{old} the value $\frac{1}{5c_1 7c_2}$. To correctly decrement C_1 , player Min should choose a delay of $4x_{old}$ at location ℓ_k . At location Check, there is no time elapse and player Max has three possibilities : (i) to go to ℓ_{k+1} and continue the simulation, or (ii) to enter the widget WD_1^1 , or (iii) to enter the widget WD_2^1 . If player Min makes an error, and delays $4x_{old} + \varepsilon$ or $4x_{old} - \varepsilon$ at ℓ_k ($\varepsilon > 0$), then player Max can enter one of the widgets and punish player Min. Player Max enters widget WD_1^1 if the error made by player Min is of the form $4x_{old} + \varepsilon$ at ℓ_k and enters widget WD_2^1 if the error made by player Min is of the form $4x_{old} - \varepsilon$ at ℓ_k .

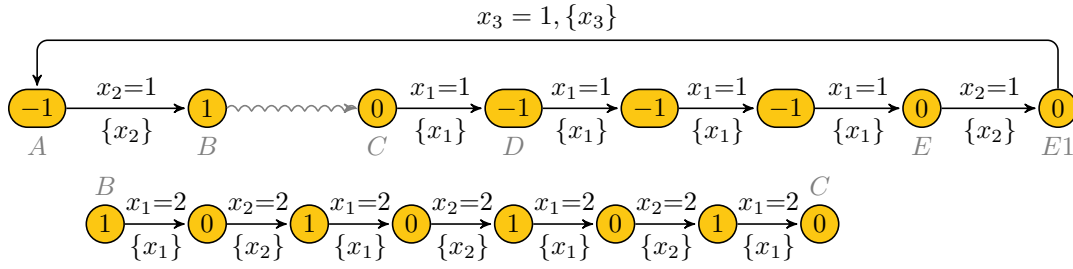


Figure 1 WD_1^1 redrawn with location prices from $\{1, 0, -1\}$. Every location has a self loop with the guard $x_2, x_3 = 1$, reset x_2, x_3 , which is not shown here for conciseness. The curly edge from B to C is shown below. The mean-payoff incurred in one transit from A to A via E is $\frac{\epsilon}{14}$. If Min makes no error, this is 0.

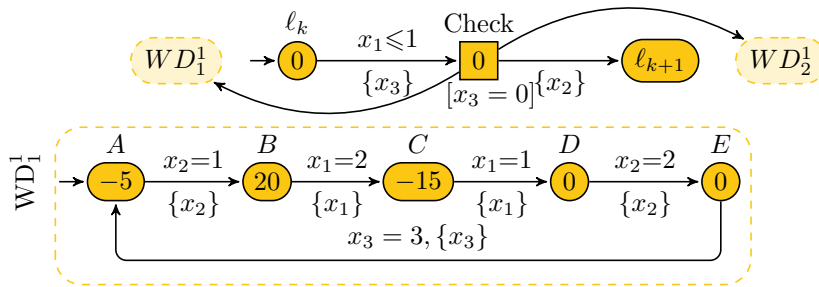


Figure 2 Simulation to decrement counter C_1 , mean cost is ϵ for error ϵ . The widget WD_2^1 has exactly the same structure and guards on all transitions as WD_1^1 , but the price signs are reversed.

Let us examine the widget WD_1^1 . When we enter WD_1^1 for the first time, we have $x_1 = x_{old} + 4x_{old} + \epsilon$, $x_2 = 4x_{old} + \epsilon$ and $x_3 = 0$. In WD_1^1 , the cost of going once from location A to E is 5ϵ . Also, when we get back to A after going through the loop once, the clock values with which we entered WD_1^1 are restored; thus, each time, we come back to A , we restore the starting values with which we enter WD_1^1 . The third clock is really useful for this purpose only. It can be seen that the mean cost of transiting from A to A through E is ϵ . In a similar way, it can be checked that the mean cost of transiting from A to A through E in widget WD_2^1 is ϵ when player Min chooses a delay $4x_{old} - \epsilon$ at ℓ_k . Thus, if player Min makes a simulation error, player Max can always choose to goto one of the widgets, and ensure that the mean pay-off is not ≤ 0 . Note that when $\epsilon = 0$, then player Min will achieve his objective: the mean pay-off will be 0. Details of other gadgets are in Appendix D in [15].

For the $MPG(T, r)$ problem ($r > 0$) we reduce the non-halting problem by constructing a PTGA with 3 clocks and location prices in $\{0, 1\}$ such that the meanpayoff is $\leq \frac{1}{3}$ iff Min does a faithful simulation. Again, details can be found in Appendix D in [15]. ◀

In Appendix D.2 in [15], we show how this undecidability results extends (with the same parameters) if one defines mean payoff per time unit instead of per step. This way of averaging across time spent was considered in [10], where the authors show the undecidability of $MPG(T, 0)$ with 5 clocks. We improve this result to show undecidability already in 3 clocks.

References

- 1 R. Alur, M. Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In *Proc. of ICALP*, pages 122–133. Springer, 2004.
- 2 R. Alur and D. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
- 3 E. Asarin and O. Maler. As soon as possible: Time optimal control for timed automata. In F. W. Vaandrager and J. H. van Schuppen, editors, *Proc. of HSCC*, pages 19–30, 1999.
- 4 H. Björklund, S. Sandberg, and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean payoff games. In *Proc. of MFCS*, pages 673–685, 2004.
- 5 P. Bouyer. Weighted timed automata: Model-checking and games. In *Proc. of MFPS*, volume 158, pages 3–17, 2006.
- 6 P. Bouyer, T. Brihaye, M. Jurdzinski, R. Lazic, and M. Rutkowski. Average-price and reachability-price games on hybrid automata with strong resets. In *FORMATS*, volume 5215 of *LNCS*, pages 63–77, 2008.
- 7 P. Bouyer, E. Brinksma, and K. G. Larsen. Staying alive as cheaply as possible. In *Proc. of HSCC*, volume 2993 of *LNCS*, pages 203–218. Springer, 2004.
- 8 P. Bouyer, F. Cassez, E. Fleury, and K. G. Larsen. Optimal strategies in priced timed game automata. In *Proc. of FSTTCS*, volume 3328 of *LNCS*, pages 148–160. Springer, 2004.
- 9 P. Bouyer, K. G. Larsen, and N. Markey. Lower-bound constrained runs in weighted timed automata. In *Proc. of QEST*, pages 128–137, 2012. doi:10.1109/QEST.2012.28.
- 10 R. Brenguier, F. Cassez, and J. F. Raskin. Energy and mean-payoff timed games. In *Proc. of HSCC*, pages 283–292, 2014. doi:10.1145/2562059.2562116.
- 11 T. Brihaye, G. Geeraerts, S. N. Krishna, L. Manasa, B. Monmege, and A. Trivedi. Adding negative prices to priced timed games. In *Proc. of CONCUR*, pages 560–575, 2014.
- 12 E. Dynkin and A. Yushkevich. *Controlled Markov Processes*. Springer, 1979.
- 13 A. Ehrenfeucht and A. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8:109–113, 1979.
- 14 J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
- 15 S. Guha, M. Jurdziński, S. N. Krishna, and A. Trivedi. Mean-payoff games on timed automata. *CoRR*, abs/1607.08480, 2016. URL: <http://arxiv.org/abs/1607.08480>.
- 16 V. A. Gurvich, A. V. Karzanov, and L. G. Khachiyan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28:85–91, 1988.
- 17 Thomas A. Henzinger. Quantitative reactive modeling and verification. *Computer Science – Research and Development*, 28(4):331–344, 2013.
- 18 R. A. Howard. *Dynamic Programming and Markov Processes*. MIT Press, 1960.
- 19 M. Jurdziński and A. Trivedi. Reachability-time games on timed automata. In *Proc. of ICALP*, pages 838–849. Springer, 2007.
- 20 M. Jurdziński and A. Trivedi. Average-time games. In R. Hariharan, M. Mukund, and V. Vinay, editors, *Proc. of FSTTCS*, Dagstuhl Seminar Proceedings, 2008.
- 21 M. L. Puterman. *Markov Decision Processes: Disc. Stoc. Dynamic Prog.* Wiley, 1994.
- 22 P. J. Ramadge and W. M. Wonham. The control of discrete event systems. In *IEEE*, volume 77, pages 81–98, 1989.
- 23 A. Trivedi. *Competitive Optimisation on Timed Automata*. PhD thesis, Department of Computer Science, The University of Warwick, 2009.
- 24 U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.