Minimum-Time Reachability in Timed Games*

Thomas Brihaye[1], Thomas A. Henzinger[2], Vinayak S. Prabhu[3], and Jean-François Raskin[4]

[1] LSV-CNRS & ENS de Cachan; thomas.brihaye@lsv.ens-cachan.fr
[2]Department of Computer and Communication Sciences, EPFL; tah@epfl.ch
[3]Department of Electrical Engineering & Computer Sciences, UC Berkeley; vinayak@eecs.berkeley.edu
[4]Département d'Informatique, Université Libre de Bruxelles; jraskin@ulb.ac.be

Abstract. We consider the minimum-time reachability problem in concurrent two-player timed automaton game structures. We show how to compute the minimum time needed by a player to reach a location against all possible choices of the opponent We do not put any syntactic restriction on the game structure, nor do we require any player to guarantee time divergence. We only require players to use physically realizable strategies. The minimal time is computed in part using a fixpoint expression which we show can be used on equivalence classes of a non-trivial extension of the region equivalence relation.

1 Introduction

Timed automata [3], finite state automata enriched with clocks and clock constraints, are a well-established formalism for the modeling and analysis of timed systems. A large number of important and interesting theoretical results have been obtained on problems in the timed automata framework. In parallel with these theoretical results, efficient verification tools have been implemented and successfully applied to industrial relevant case studies.

Timed automata are models for closed systems, where every transition is controlled. If we want to distinguish between actions of several agents (for instance a *controller* and an *environment*) we have to consider games on timed automata, also known as *timed games*. In the sequel, we will focus on two-player games. In this context, the *reachability problem* asks whether player-1 has a *strategy* to force the timed game to reach a target set no matter how player-2 resolves his choices. These games were first introduced and studied in [22, 19]. In this framework, it is also natural to consider a *minimum-time reachability problem* — which asks what is the minimal time required by player-1 to reach a target set, no matter how player-2 resolves his choices. The work of [5] has shown this problem to be decidable for a partial class of timed automata games.

Any formalism involving timed systems has to face the problem of *zeno runs* — runs of the model where time converges. Zeno runs are of course not physically

* This research was supported in part by the NSF grants CCR-0208875, CCR-0225610, and CCR-0234690, and by the Swiss National Science Foundation.

meaningful. Avoidance of such runs has often been achieved by putting syntactic constraints on the cycles of timed automaton games [5, 15, 6, 20] or semantic conditions discretizing time. Other works on the existence of controllers [14, 7, 19, 9] have in general required that time divergence be ensured by the controller — an unfair view in settings where the player modelling the environment can also block time.

Recently, a more equitable treatment of zeno sequences has been proposed in [12]. This setting formulates a symmetric setup of the model, where all players are given equally powerful options for updating the state of the game, advancing time, or blocking time. Both players may block time, however, for a player to win for an objective, he must not be *responsible* for preventing time from diverging. It has been shown in [17] that this is equivalent to requiring that players use only *receptive* strategies [21, 4] in plays, so that zeno runs are avoided.

Our contribution: In this paper, we consider the *minimum time reachability problem (for timed games)*, in the framework of [12, 17]. We present an EXPTIME algorithm to compute minimum time needed by player-1 to force the game into a given set of locations, with both players restricted to using only receptive strategies (note that reachability in timed games is know to be EXPTIME-complete [16]). The proof technique builds on the techniques that have been introduced in [12, 17] and we show that the minimum time can be obtained by solving a μ-calculus fixed point equation. Proof of termination of the fixpoint algorithm requires an important new ingredient: an extension of the well-known region equivalence for timed automata. We show our extended region equivalence classes to be stable w.r.t. the monotone functions used in the fixed point equation. Using previous results of [17], we manage to restrict to finitely many regions and thus guarantee termination.

We note that standard regions do not suffice — the minimum-time reachability game has two components, the reachability part can be handled by discrete arguments based on the region graph, the minimum part requires minimization within regions in some sense (as in [11]). Unfortunately, both arguments are intertwined and cannot be argued about in isolation. Our extended regions manage to decouple these two threads in the proofs. We also note that region sequences which correspond to optimal runs may in general be *required* to contain region cycles in which time does *not* progress by an integer amount, thus an approach like that of [1] to reduce to a loop free region game also runs into problems.

Related work: Only special cases of the optimal reachability problem have been solved before, the work in [5] restricts its attention to the case where every cycle of the timed automaton ensures syntactically that a positive amount of time passes (a.k.a. strong non-zenoness assumption), [2] considers games restricted to a bounded number of moves, [17] presents an approximate computation of the minimum time (computation of the exact minimal time being left open). The general case for *weighted* timed automaton games is known to be undecidable [8]. The recent work of [18] presents a strategy improvement algorithm which computes the optimal time in all timed automaton games, but it does not require strategies to be receptive. Average reward games in the framework of [12] are considered in [1], but with the durations of time moves restricted to either 0 or 1. The non-game version of the minimum-time problem is presented in [11].

The rest of the paper is organised as follows. In Section 2, we recall the notations and definitions of timed games framework of [12]. The minimum-time reachability problem is formally defined in Section 3. Section 4 is the main section of the paper and is devoted to our algorithm which solves for minimum-time reachability and proof of its termination.

2 Timed Games

2.1 Timed Game Structures

We use the formalism of [12]. A *timed game structure* is a tuple $\mathcal{G} = \langle S, \Sigma, \sigma, A_1, A_2, \Gamma_1, \Gamma_2, \delta \rangle$ with the following components:

- S is a set of states.
- Σ is a finite set of propositions.
- $\sigma : S \mapsto 2^{\Sigma}$ is the observation map, which assigns to every state the set of propositions that are true in that state.
- A_1 and A_2 are two disjoint sets of actions for players 1 and 2, respectively. We assume that $\perp_i \notin A_i$, and write $A_i^{\perp}$ for $A_i \cup \{\perp_i\}$. We also assume $A_1^{\perp}$ and $A_2^{\perp}$ to be disjoint. The set of *moves* for player i is $M_i = \mathbb{R}_{\geq 0} \times A_i^{\perp}$. Intuitively, a move $\langle \Delta, a_i \rangle$ by player i indicates a waiting period of Δ time units followed by a discrete transition labeled with action a_i.
- $\Gamma_i : S \mapsto 2^{M_i} \setminus \emptyset$ are two move assignments. At every state s, the set $\Gamma_i(s)$ contains the moves that are available to player i. We require that $\langle 0, \perp_i \rangle \in \Gamma_i(s)$ for all states $s \in S$ and $i \in \{1, 2\}$. Intuitively, $\langle 0, \perp_i \rangle$ is a time-blocking stutter move.
- $\delta : S \times (M_1 \cup M_2) \mapsto S$ is the transition function. We require that for all time delays $\Delta, \Delta' \in \mathbb{R}_{\geq 0}$ with $\Delta' \leq \Delta$, and all actions $a_i \in A_i^{\perp}$, we have (1) $\langle \Delta, a_i \rangle \in \Gamma_i(s)$ iff both $\langle \Delta', \perp_i \rangle \in \Gamma_i(s)$ and $\langle \Delta - \Delta', a_i \rangle \in \Gamma_i(\delta(s, \langle \Delta', \perp_i \rangle))$; and (2) if $\delta(s, \langle \Delta', \perp_i \rangle) = s'$ and $\delta(s', \langle \Delta - \Delta', a_i \rangle) = s''$, then $\delta(s, \langle \Delta, a_i \rangle) = s''$.

The game proceeds as follows. If the current state of the game is s, then both players simultaneously propose moves $\langle \Delta_1, a_1 \rangle \in \Gamma_1(s)$ and $\langle \Delta_2, a_2 \rangle \in \Gamma_2(s)$. The move with the shorter duration "wins" in determining the next state of the game. If both moves have the same duration, then one of the two moves is chosen non-deterministically. Formally, we define the *joint destination function* $\delta_{\mathsf{jd}} : S \times M_1 \times M_2 \mapsto 2^S$ by

$$\delta_{\mathsf{jd}}(s, \langle \Delta_1, a_1 \rangle, \langle \Delta_2, a_2 \rangle) = \begin{cases} \{\delta(s, \langle \Delta_1, a_1 \rangle)\} & \text{if } \Delta_1 < \Delta_2; \\ \{\delta(s, \langle \Delta_2, a_2 \rangle)\} & \text{if } \Delta_2 < \Delta_1; \\ \{\delta(s, \langle \Delta_1, a_1 \rangle), \delta(s, \langle \Delta_2, a_2 \rangle)\} & \text{if } \Delta_1 = \Delta_2. \end{cases}$$

The time elapsed when the moves $m_1 = \langle \Delta_1, a_1 \rangle$ and $m_2 = \langle \Delta_2, a_2 \rangle$ are proposed is given by $\mathsf{delay}(m_1, m_2) = \min(\Delta_1, \Delta_2)$. The boolean predicate

$\mathsf{blame}_i(s, m_1, m_2, s')$ indicates whether player i is "responsible" for the state change from s to s' when the moves m_1 and m_2 are proposed. Denoting the opponent of player $i \in \{1, 2\}$ by $\sim i = 3 - i$, we define

$$\mathsf{blame}_i(s, \langle \Delta_1, a_1\rangle, \langle \Delta_2, a_2\rangle, s') \;=\; \big(\Delta_i \leq \Delta_{\sim i} \;\wedge\; \delta(s, \langle \Delta_i, a_i\rangle) = s'\big).$$

A *run* of the timed game structure $\mathcal{G}$ is an infinite sequence $r = s_0, \langle m_1^0, m_2^0\rangle, s_1, \langle m_1^1, m_2^1\rangle, \ldots$ such that $s_k \in S$ and $m_i^k \in \Gamma_i(s_k)$ and $s_{k+1} \in \delta_{\mathsf{jd}}(s_k, m_1^k, m_2^k)$ for all $k \geq 0$ and $i \in 1, 2$. For $k \geq 0$, let $\mathsf{time}(r, k)$ denote the "time" at position k of the run, namely, $\mathsf{time}(r, k) = \sum_{j=0}^{k-1} \mathsf{delay}(m_1^j, m_2^j)$ (we let $\mathsf{time}(r, 0) = 0$). By $r[k]$ we denote the $(k+1)$-th state s_k of r. The run prefix $r[0..k]$ is the finite prefix of the run r that ends in the state s_k; we write $\mathsf{last}(r[0..k])$ for the ending state s_k of the run prefix. Let Runs be the set of all runs of $\mathcal{G}$, and let $\mathsf{FinRuns}$ be the set of run prefixes.

A *strategy* π_i for player $i \in \{1, 2\}$ is a function $\pi_i : \mathsf{FinRuns} \mapsto M_i$ that assigns to every run prefix $r[0..k]$ a move to be proposed by player i at the state $\mathsf{last}(r[0..k])$ if the history of the game is $r[0..k]$. We require that $\pi_i(r[0..k]) \in \Gamma_i(\mathsf{last}(r[0..k]))$ for every run prefix $r[0..k]$, so that strategies propose only available moves. The results of this paper are equally valid if strategies do not depend on past moves chosen by the players, but only on the past sequence of states and time delays [12]. For $i \in \{1, 2\}$, let Π_i be the set of player-i strategies. Given two strategies $\pi_1 \in \Pi_1$ and $\pi_2 \in \Pi_2$, the set of possible *outcomes* of the game starting from a state $s \in S$ is denoted $\mathsf{Outcomes}(s, \pi_1, \pi_2)$: it contains all runs $r = s_0, \langle m_1^0, m_2^0\rangle, s_1, \langle m_1^1, m_2^1\rangle, \ldots$ such that $s_0 = s$ and for all $k \geq 0$ and $i \in \{1, 2\}$, we have $\pi_i(r[0..k]) = m_i^k$.

We distinguish between *physical time* and *game time*. We allow moves with zero time delay, thus a physical time $t \in \mathbb{R}_{\geq 0}$ may correspond to several linearly ordered states, to which we assign the game times $\langle t, 0\rangle, \langle t, 1\rangle, \langle t, 2\rangle, \ldots$ For a run $r \in \mathsf{Runs}$, we define the set of game times as

$$\mathsf{GameTimes}(r) \;=\; \begin{array}{l} \{\langle t, k\rangle \in \mathbb{R}_{\geq 0} \times \mathbb{N} \mid 0 \leq k < |\{j \geq 0 \mid \mathsf{time}(r, j) = t\}|\} \,\cup \\ \{\langle t, 0\rangle \mid \mathsf{time}(r, j) \geq t \text{ for some } j \geq 0\}. \end{array}$$

The state of the run r at a game time $\langle t, k\rangle \in \mathsf{GameTimes}(r)$ is defined as

$$\mathsf{state}(r, \langle t, k\rangle) \;=\; \begin{cases} r[j + k] \;\; \text{if } \mathsf{time}(r, j) = t \text{ and for all } j' < j, \mathsf{time}(r, j') < t; \\ \delta(r[j], \langle t - \mathsf{time}(r, j), \perp_i\rangle) \;\; \text{if } \mathsf{time}(r, j) < t < \mathsf{time}(r, j + 1) \text{ and} \\ \qquad\qquad r[0..j + 1] = r[0..j], \langle m_1^j, m_2^j\rangle, r[j + 1] \text{ and} \\ \qquad\qquad \mathsf{blame}_i(r[j], m_1^j, m_2^j, r[j + 1]) \end{cases}$$

Note that if r is a run of the timed game structure $\mathcal{G}$, and $\mathsf{time}(r, j) < t < \mathsf{time}(r, j + 1)$, then $\delta(r[j], \langle t - \mathsf{time}(r, j), \perp_i\rangle)$ is a state in S, namely, the state that results from $r[j]$ by letting time $t - \mathsf{time}(r, j)$ pass. We say that the run r *visits* a set $X \subseteq S$ at time t if there is a $\tau = \langle t, k\rangle \in \mathsf{GameTimes}(r)$ such that $\mathsf{state}(r, \tau) \in X$. A run r visits a proposition $p \in \Sigma$ if it visits the set S_p defined as $\{s \mid p \in \sigma(s)\}$.

2.2 Timed Automaton Games

Timed automata [3] suggest a finite syntax for specifying infinite-state timed game structures. A *timed automaton game* is a tuple $\mathcal{T} = \langle L, \Sigma, \sigma, C, A_1, A_2, E, \gamma \rangle$ with the following components:

- L is a finite set of locations.
- Σ is a finite set of propositions.
- $\sigma : L \mapsto 2^{\Sigma}$ assigns to every location a set of propositions.
- C is a finite set of clocks. We assume that $z \in C$ for the unresettable clock z, which is used to measure the time elapsed since the start of the game.
- A_1 and A_2 are two disjoint sets of actions for players 1 and 2, respectively.
- $E \subseteq L \times (A_1 \cup A_2) \times \mathsf{Constr}(C) \times L \times 2^{C \setminus \{z\}}$ is the edge relation, where the set $\mathsf{Constr}(C)$ of *clock constraints* is generated by the grammar

  $$\theta ::= x \leq d \mid d \leq x \mid \neg\theta \mid \theta_1 \wedge \theta_2$$

  for clock variables $x \in C$ and nonnegative integer constants d. For an edge $e = \langle l, a_i, \theta, l', \lambda \rangle$, the clock constraint θ acts as a guard on the clock values which specifies when the edge e can be taken, and by taking the edge e, the clocks in the set $\lambda \subseteq C \setminus \{z\}$ are reset to 0. We require that for all edges $\langle l, a_i, \theta', l', \lambda' \rangle, \langle l, a_i, \theta'', l'', \lambda'' \rangle \in E$ with $l' \neq l''$, the conjunction $\theta' \wedge \theta''$ is unsatisfiable. This requirement ensures that a state and a move together uniquely determine a successor state.
- $\gamma : L \mapsto \mathsf{Constr}(C)$ is a function that assigns to every location an invariant for both players. All clocks increase uniformly at the same rate. When at location l, each player i must propose a move out of l before the invariant $\gamma(l)$ expires. Thus, the game can stay at a location only as long as the invariant is satisfied by the clock values.

A *clock valuation* is a function $\kappa : C \mapsto \mathbb{R}_{\geq 0}$ that maps every clock to a nonnegative real. The set of all clock valuations for C is denoted by $K(C)$. Given a clock valuation $\kappa \in K(C)$ and a time delay $\Delta \in \mathbb{R}_{\geq 0}$, we write $\kappa + \Delta$ for the clock valuation in $K(C)$ defined by $(\kappa + \Delta)(x) = \kappa(x) + \Delta$ for all clocks $x \in C$. For a subset $\lambda \subseteq C$ of the clocks, we write $\kappa[\lambda := 0]$ for the clock valuation in $K(C)$ defined by $(\kappa[\lambda := 0])(x) = 0$ if $x \in \lambda$, and $(\kappa[\lambda := 0])(x) = \kappa(x)$ if $x \notin \lambda$. A clock valuation $\kappa \in K(C)$ *satisfies* the clock constraint $\theta \in \mathsf{Constr}(C)$, written $\kappa \models \theta$, if the condition θ holds when all clocks in C take on the values specified by κ.

A *state* $s = \langle l, \kappa \rangle$ of the timed automaton game $\mathcal{T}$ is a location $l \in L$ together with a clock valuation $\kappa \in K(C)$ such that the invariant at the location is satisfied, that is, $\kappa \models \gamma(l)$. Let S be the set of all states of $\mathcal{T}$. In a state, each player i proposes a time delay allowed by the invariant map γ, together either with the action $\bot$, or with an action $a_i \in A_i$ such that an edge labeled a_i is enabled after the proposed time delay. We require that for $i \in \{1, 2\}$ and for all states $s = \langle l, \kappa \rangle$, if $\kappa \models \gamma(l)$, either $\kappa + \Delta \models \gamma(l)$ for all $\Delta \in \mathbb{R}_{\geq 0}$, or there exist a time delay $\Delta \in \mathbb{R}_{\geq 0}$ and an edge $\langle l, a_i, \theta, l', \lambda \rangle \in E$ such that (1) $a_i \in A_i$

and (2) $\kappa + \Delta \models \theta$ and for all $0 \leq \Delta' \leq \Delta$, we have $\kappa + \Delta' \models \gamma(l)$, and (3) $(\kappa + \Delta)[\lambda := 0] \models \gamma(l')$.

The timed automaton game $\mathcal{T}$ defines the following timed game structure $[\![\mathcal{T}]\!] = \langle S, \Sigma, \sigma^*, A_1, A_2, \Gamma_1, \Gamma_2, \delta \rangle$:

- S is defined above.
- $\sigma^*(\langle l, \kappa \rangle) = \sigma(l)$.
- For $i \in \{1, 2\}$, the set $\Gamma_i(\langle l, \kappa \rangle)$ contains the following elements:
  1. $\langle \Delta, \perp_i \rangle$ if for all $0 \leq \Delta' \leq \Delta$, we have $\kappa + \Delta' \models \gamma(l)$.
  2. $\langle \Delta, a_i \rangle$ if for all $0 \leq \Delta' \leq \Delta$, we have $\kappa + \Delta' \models \gamma(l)$, and $a_i \in A_i$, and there exists an edge $\langle l, a_i, \theta, l', \lambda \rangle \in E$ such that $\kappa + \Delta \models \theta$.
- $\delta(\langle l, \kappa \rangle, \langle \Delta, \perp_i \rangle) = \langle l, \kappa + \Delta \rangle$, and $\delta(\langle l, \kappa \rangle, \langle \Delta, a_i \rangle) = \langle l', (\kappa + \Delta)[\lambda := 0] \rangle$ for the unique edge $\langle l, a_i, \theta, l', \lambda \rangle \in E$ with $\kappa + \Delta \models \theta$.

2.3 Clock Regions

Timed automaton games can be solved using a region construction from the theory of timed automata [3]. For a real $t \geq 0$, let $\mathsf{frac}(t) = t - \lfloor t \rfloor$ denote the fractional part of t. Given a timed automaton game $\mathcal{T}$, for each clock $x \in C$, let c_x denote the largest integer constant that appears in any clock constraint involving x in $\mathcal{T}$ Two clock valuations $\kappa_1, \kappa_2 \in K(C)$ are *clock-region equivalent*, denoted $\kappa_1 \cong \kappa_2$, if the following three conditions hold:

1. For all $x \in C$, either $\lfloor \kappa_1(x) \rfloor = \lfloor \kappa_2(x) \rfloor$, or both $\lfloor \kappa_1(x) \rfloor > c_x$, $\lfloor \kappa_2(x) \rfloor > c_x$.
2. For all $x, y \in C$ with $\kappa_1(x) \leq c_x$ and $\kappa_1(y) \leq c_y$, we have $\mathsf{frac}(\kappa_1(x)) \leq \mathsf{frac}(\kappa_1(y))$ iff $\mathsf{frac}(\kappa_2(x)) \leq \mathsf{frac}(\kappa_2(y))$.
3. For all $x \in C$ with $\kappa_1(x) \leq c_x$, we have $\mathsf{frac}(\kappa_1(x)) = 0$ iff $\mathsf{frac}(\kappa_2(x)) = 0$.

Two states $\langle l_1, \kappa_1 \rangle, \langle l_2, \kappa_2 \rangle \in S$ are *clock-region equivalent*, denoted $\langle l_1, \kappa_1 \rangle \cong \langle l_2, \kappa_2 \rangle$, iff $l_1 = l_2$ and $\kappa_1 \cong \kappa_2$. It is not difficult to see that $\cong$ is an equivalence relation on S. A *clock region* is an equivalence class with respect to $\cong$. There are finitely many clock regions; more precisely, the number of clock regions is bounded by $|L| \cdot \prod_{x \in C}(c_x + 1) \cdot |C|! \cdot 2^{|C|}$. For a state $s \in S$, we write $[s] \subseteq S$ for the clock region containing s. These clock regions induce a time-abstract bisimulation.

3 The Minimum-Time Reachability Problem

Given a state s and a target proposition $p \in \Sigma$ in a timed game structure $\mathcal{G}$, the *reachability* problem is to determine whether starting from s, player-1 has a strategy for visiting the proposition p. We must make sure that player-2 does not prevent player-1 from reaching a target state by blocking time. We also require player-1 to not block time as it can lead to physically unmeaningful plays. These requirements can be achieved by requiring strategies to be *receptive* [21, 4]. Formally, we first define the following two sets of runs:

– $\mathsf{Timediv} \subseteq \mathsf{Runs}$ is the set of all time-divergent runs. A run r is *time-divergent* if $\lim_{k\to\infty} \mathsf{time}(r,k) = \infty$.
– $\mathsf{Blameless}_i \subseteq \mathsf{Runs}$ is the set of runs in which player i is responsible only for finitely many transitions. A run $s_0, \langle m_1^0, m_2^0\rangle, s_1, \langle m_1^1, m_2^1\rangle, \ldots$ belongs to the set $\mathsf{Blameless}_i$, for $i = \{1,2\}$, if there exists a $k \geq 0$ such that for all $j \geq k$, we have $\neg\, \mathsf{blame}_i(s_j, m_1^j, m_2^j, s_{j+1})$.

A strategy π_i for player $i \in \{1,2\}$ is *receptive* if for all opposing strategies $\pi_{\sim i}$, and all states $s \in S$, $\mathsf{Outcomes}(s, \pi_1, \pi_2) \subseteq \mathsf{Timediv} \cup \mathsf{Blameless}_i$. Thus, no what matter what the opponent does, a receptive player-i strategy should not be responsible for blocking time. Strategies that are not receptive are not physically meaningful (note that receptiveness is not sufficient for a strategy to be physically meaningful [10]). For $i \in \{1,2\}$, let Π_i^R be the set of player-i receptive strategies. A timed game structure is *well-formed* if both players have receptive strategies. We restrict our attention to well-formed timed game structures. Well-formedness of timed automaton games can be checked for (see [17]).

We say player-1 *wins* for the reachability objective p at state s, denoted $s \in \langle\langle 1\rangle\rangle \Diamond p$, if he has a receptive strategy π_1 such that for all player-2 receptive strategies π_2, we have that all runs $r \in \mathsf{Outcomes}(s, \pi_1, \pi_2)$ visit the proposition p.

Equivalently [17], we can define player-1 to be winning for the reachability objective p at state s if he has a strategy π_1 such that for all player-2 strategies π_2, for all runs $r \in \mathsf{Outcomes}(s, \pi_1, \pi_2)$:

– if $r \in \mathsf{Timediv}$, then r visits the proposition p;
– if $r \notin \mathsf{Timediv}$, then $r \in \mathsf{Blameless}_1$.

The *minimum-time reachability problem* is to determine the minimal time in which a player can force the game into a set of target states, using only receptive strategies. Formally, given a timed game structure $\mathcal{G}$, a target proposition $p \in \Sigma$, and a run r of $\mathcal{G}$, let

$$T_{\mathrm{visit}}(\mathcal{G}, r, p) = \begin{cases} \infty & \text{if } r \text{ does not visit } p; \\ \inf\{t \in \mathbb{R}_{\geq 0} \mid p \in \sigma(\mathsf{state}(r, \langle t, k\rangle)) \text{ for some } k\} & \text{otherwise.} \end{cases}$$

The *minimal time* for player-1 to force the game from a start state $s \in S$ to a visit to p is then

$$T_{\min}(\mathcal{G}, s, p) = \inf_{\pi_1 \in \Pi_1^R} \sup_{\pi_2 \in \Pi_2^R} \sup_{r \in \mathsf{Outcomes}(s, \pi_1, \pi_2)} T_{\mathrm{visit}}(\mathcal{G}, r, p)$$

We omit $\mathcal{G}$ when clear from the context.

4 Solving for Minimum-Time Reachability

We restrict our attention to well-formed timed automaton games. The definition of $T_{\min}$ quantifies strategies over the set of receptive strategies. Our algorithm will instead work over the set of *all* strategies. Theorem 1 presents this reduction. We will then present a game structure for the timed automaton game $\mathcal{T}$ in which $\mathsf{Timediv}$ and $\mathsf{Blameless}_1$ can be represented using Büchi and co-Büchi constraints. This builds on the framework of [12] in which a run satisfies the reachability objective p for player-1 iff it belongs in $(\mathsf{Timediv} \cap \mathsf{Reach}(p)) \cup (\neg\, \mathsf{Timediv} \cap \mathsf{Blameless}_1)$, where $\mathsf{Reach}(p)$ denotes the set of runs which visit p. In addition, our game structure will also have a backwards running clock, which will be used in the computation of the minimum time, using a μ-calculus algorithm on *extended regions*.

4.1 Allowing Players to Use All Strategies

To allow quantification over all strategies, we first modify the payoff function T_{visit}, so that players are maximally penalised on zeno runs:

$$T_{\mathrm{visit}}^{\mathrm{UR}}(r,p) = \begin{cases} \infty \;\text{ if } r \notin \mathsf{Timediv} \text{ and } r \notin \mathsf{Blameless}_i; \\ \infty \;\text{ if } r \in \mathsf{Timediv} \text{ and } r \text{ does not visit } p; \\ 0 \;\text{ if } r \notin \mathsf{Timediv} \text{ and } r \in \mathsf{Blameless}_i; \\ \inf \{t \in \mathbb{R}_{\geq 0} \mid p \in \sigma(\mathsf{state}(r, \langle t, k\rangle)) \text{ for some } k\} \;\text{ otherwise.} \end{cases}$$

It turns out that penalizing on zeno-runs is equivalent to penalising on non-receptive strategies:

Theorem 1. *For all well-formed timed game structures $\mathcal{G}$, for all states s and propositions p of $\mathcal{G}$, we have*

$$T_{\min}(s,p) = \inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} T_{\mathrm{visit}}^{\mathrm{UR}}(r,p)$$

4.2 Reduction to Reachability with Büchi and co-Büchi Constraints

We now decouple reachability from optimizing for minimal time, and show how reachability with time divergence can be solved for, using an appropriately chosen μ-calculus fixpoint.

Lemma 1 ([17]). *Given a state s, and a proposition p of a well-formed timed automaton game $\mathcal{T}$, 1)we can determine if $T_{\min}(s,p) < \infty$, and 2) If $T_{\min}(s,p) < \infty$, then $T_{\min}(s,p) < M = 8|L| \cdot \prod_{x\in C}(c_x+1) \cdot |C+1|! \cdot 2^{|C|}$. This upper bound is the same for all $s' \cong s$.*

Let M be the upper bound on $T_{\min}(s,p)$ as in Lemma 1 if $T_{\min}(s,p) < \infty$, and $M = 1$ otherwise. For a number N, let $\mathbb{R}_{[0,N]}$ and $\mathbb{R}_{[0,N)}$ denote $\mathbb{R} \cap [0,N]$ and $\mathbb{R} \cap [0,N)$ respectively. We first look at the enlarged game structure $\widehat{[\![\mathcal{T}]\!]}$ with the state space $\widehat{S} = S \times \mathbb{R}_{[0,1)} \times (\mathbb{R}_{[0,M]} \cup \{\perp\}) \times \{\mathrm{TRUE}, \mathrm{FALSE}\}^2$, and an augmented transition relation $\widehat{\delta} : \widehat{S} \times (M_1 \cup M_2) \mapsto \widehat{S}$. In an augmented state $\langle s, \mathfrak{z}, \beta, tick, bl_1\rangle \in \widehat{S}$, the component $s \in S$ is a state of the original game structure $[\![\mathcal{T}]\!]$, $\mathfrak{z}$ is value of a fictitious clock which gets reset every time it hits 1, β is the value of a fictitious clock which is running *backwards*, *tick* is true iff one time unit has passed since the last time it was true (so *tick* is true iff the last transition resulted in $\mathfrak{z} = 0$), and bl_1 is true if player-1 is to blame for the last transition.

Formally, $\langle s', \mathfrak{z}', \beta', tick', bl_1' \rangle = \widehat{\delta}(\langle s, \mathfrak{z}, \beta, tick, bl_1 \rangle, \langle \Delta, a_i \rangle)$ iff

1. $s' = \delta(s, \langle \Delta, a_i \rangle)$
2. $\mathfrak{z}' = (\mathfrak{z} + \Delta) \bmod 1$;
3. $\beta' = \beta \ominus \Delta$, where we define $\beta \ominus \Delta$ as $\beta - \Delta$ if $\beta \neq \bot$ and $\beta - \Delta \geq 0$, and $\bot$ otherwise ($\bot$ is an absorbing value for β).
4. $tick' = \text{TRUE}$ if $\mathfrak{z} + \Delta \geq 1$, and FALSE otherwise
5. $bl_1 = \text{TRUE}$ if $a_i \in A_1^{\perp}$ and FALSE otherwise.

Each run r of $[\![\mathfrak{T}]\!]$, and values $\mathfrak{z} \in \mathbb{R}_{\geq 0}, \beta \leq M$ can be mapped to a corresponding unique run $\widehat{r}_{\mathfrak{z},\beta}$ in $\widehat{[\![\mathfrak{T}]\!]}$, with $\widehat{r}_{\mathfrak{z},\beta}[0] = \langle r[0], \mathfrak{z}, \beta, \text{FALSE}, \text{FALSE} \rangle$. Similarly, each run $\widehat{r}$ of $\widehat{[\![\mathfrak{T}]\!]}$ can be projected to a unique run $\widehat{r} \downarrow \mathfrak{T}$ of $[\![\mathfrak{T}]\!]$. It can be seen that the run r is in $\mathsf{Timediv}$ iff *tick* is true infinitely often in $\widehat{r}_{\mathfrak{z},\beta}$, and that the set $\mathsf{Blameless}_1$ corresponds to runs along which bl_1 is true only finitely often.

Proposition 1. *Given a timed game structure $[\![\mathfrak{T}]\!]$, let $\widehat{X}_p = S_p \times \mathbb{R}_{[0,1)} \times \{0\} \times \{\text{TRUE}, \text{FALSE}\}^2$.*

1. *For a run r of the timed game structure $[\![\mathfrak{T}]\!]$, let $T_{\text{visit}}(r,p) < \infty$. Then, $T_{\text{visit}}(r,p) = \inf\{\beta \mid \beta \in \mathbb{R}_{[0,M]}$ and $\widehat{r}_{0,\beta}$ visits the set $\widehat{X}_p\}$.*
2. *Let $T_{\min}(s,p) < \infty$. Then,*
$T_{\min}(s,p) = \inf \left\{ \beta \mid \beta \in \mathbb{R}_{[0,M]} \text{ and } \langle s, 0, \beta, \text{FALSE}, \text{FALSE} \rangle \in \langle\langle 1 \rangle\rangle \Diamond \widehat{X}_p \right\}$
3. *If $T_{\min}(s,p) = \infty$, then for all β, we have $\langle s, 0, \beta, \text{FALSE}, \text{FALSE} \rangle \notin \langle\langle 1 \rangle\rangle \Diamond \widehat{X}_p$.*

The rechability objective can be reduced to a parity game: each state in $\widehat{S}$ is assigned an index $\Omega : \widehat{S} \mapsto \{0,1\}$, with $\Omega(\widehat{s}) = 1$ iff $\widehat{s} \notin \widehat{X_p}$; and $tick = \text{TRUE}$ or $bl_1 = \text{TRUE}$. We also modify the game structure so that the states in $\widehat{X_p}$ are absorbing.

Lemma 2. *For the timed game $\widehat{[\![\mathfrak{T}]\!]}$ with the reachability objective $\widehat{X}_p$, the state $\widehat{s} = \langle s, 0, \beta, \text{FALSE}, \text{FALSE} \rangle \in \langle\langle 1 \rangle\rangle \Diamond \widehat{X}_p$ iff player-1 has a strategy π_1 such that for all strategies π_2 of player-2, and all runs $\widehat{r}_{0,\beta} \in \mathsf{Outcomes}(\widehat{s}, \pi_1, \pi_2)$, the index 1 does not occur infinitely often in $\widehat{r}_{0,\beta}$.*

The fixpoint formula for solving the parity game in Lemma 2 is given by [13]

$$Y = \mu Y \nu Z \left[(\Omega^{-1}(1) \cap \mathsf{CPre}_1(Y)) \cup (\Omega^{-1}(0) \cap \mathsf{CPre}_1(Z)) \right]$$

The fixpoint expression uses the variables $Y, Z \subseteq \widehat{S}$ and the *controllable predecessor operator*, $\mathsf{CPre}_1 : 2^{\widehat{S}} \mapsto 2^{\widehat{S}}$ in its fixpoint computation, defined formally by $\mathsf{CPre}_1(X) \equiv \{\widehat{s} \mid \exists m_1 \in \Gamma_1(\widehat{s})\ \forall m_2 \in \Gamma_2(\widehat{s}) (\widehat{\delta}_{\mathsf{jd}}(\widehat{s}, m_1, m_2) \subseteq X)\}$. Intuitively, $\widehat{s} \in \mathsf{CPre}_1(X)$ iff player 1 can force the augmented game from $\widehat{s}$ into X in one move.

4.3 Termination of the μ-Calculus Fixpoint Iteration

We prove termination of the μ-calculus algorithm by demonstrating that we can work on a finite partition of the state space. Let an equivalence relation $\cong_e$ on the states in $\widehat{S}$ be defined as: $\langle\langle l^1, \kappa^1\rangle, \mathfrak{z}^1, \beta^1, tick^1, bl_1^1\rangle \cong_e \langle\langle l^2, \kappa^2\rangle, \mathfrak{z}^2, \beta^2, tick^2, bl_1^2\rangle$ iff

1. $l^1 = l^2, tick^1 = tick^2$, and $bl^1 = bl^2$.
2. $\widehat{\kappa^1} \cong \widehat{\kappa^2}$ where $\widehat{\kappa^i} : C \cup \{z\} \mapsto \mathbb{R}_{\geq 0}$ is a clock valuation such that $\widehat{\kappa^i}(c) = \kappa^i(c)$ for $c \in C$, $\widehat{\kappa^i}(z) = \mathfrak{z}^i$,and $c_z = 1$ (c_z is the maximum value of the clock z in the definition of $\cong$) for $i \in \{1, 2\}$.
3. $\beta^1 = \perp$ iff $\beta^2 = \perp$.
4. If $\beta^1 \neq \perp, \beta^2 \neq \perp$ then
   - $\lfloor\beta^1\rfloor = \lfloor\beta^2\rfloor$
   - $\mathsf{frac}(\beta^1) = 0$ iff $\mathsf{frac}(\beta^2) = 0$.
   - For each clock $x \in C \cup \{z\}$ with $\kappa^1(x) \leq c_x$ and $\kappa^2(x) \leq c_x$, we have $\mathsf{frac}(\kappa^1(x)) + \mathsf{frac}(\beta^1) \sim 1$ iff $\mathsf{frac}(\kappa^2(x)) + \mathsf{frac}(\beta^2) \sim 1$ with $\sim \in \{<, =, >\}$.

The number of equivalence classes induced by $\cong_e$ is again finite $\left(O\left((|L| \cdot \prod_{x \in C}(c_x + 1) \cdot |C + 1|! \cdot 2^{|C|})^2 \cdot |C|\right)\right)$. We call each equivalence class an *extended region*. An extended region Y of $\widehat{[\![\mathcal{T}]\!]}$ can be specified by the tuple $\langle l, tick, bl_1, h, \mathcal{P}, \beta_i, \beta_f, C_<, C_=, C_>\rangle$ where for a state $\widehat{s} = \langle\langle l, \kappa\rangle, \mathfrak{z}, \beta, tick, bl_1\rangle$ in Y,

- $l, tick, bl_1$ correspond to $l, tick, bl_1$ in $\widehat{s}$.
- h is a function which specifies the integer values of clocks: $h(x) = \lfloor\kappa(x)\rfloor$ if $\kappa(x) < C_x + 1$, and $h(x) = C_x + 1$ otherwise.
- $\mathcal{P} \subseteq 2^{C \cup \{z\}}$ is a partition of the clocks $\{C_0, \ldots, C_n \mid \uplus C_i = C \cup \{z\}, C_i \neq \emptyset$ for $i > 0\}$, such that 1)for any pair of clocks x, y, we have $\mathsf{frac}(\kappa(x)) < \mathsf{frac}(\kappa(y))$ iff $x \in C_j, y \in C_k$ for $j < k$; and 2)$x \in C_0$ iff $\mathsf{frac}(\kappa(x)) = 0$.
- $\beta_i \in \mathbb{N} \cap \{0, \ldots, M\} \cup \{\perp\}$ indicates the integral value of β.
- $\beta_f \in \{\text{TRUE}, \text{FALSE}\}$ indicates whether the fractional value of β is greater than 0, $\beta_f = \text{TRUE}$ iff $\beta \neq \perp$ and $\mathsf{frac}(\beta) > 0$.
- For a clock $x \in C \cup \{z\}$ and $\beta \neq \perp$, we have $\mathsf{frac}(\kappa(x)) + \mathsf{frac}(\beta) \sim 1$ iff $x \in C_\sim$ for $\sim \in \{<, =, >\}$.

Pictorially, the relationship between $\widehat{\kappa}$ and β can be visualised as in Fig. 1. The figure depicts an extended region for $C_0 = \emptyset, \beta_i \in \mathbb{N} \cap \{0, \ldots, M\}, \beta_f = \text{TRUE}, C_< = C \cup \{z\}, C_= = \emptyset, C_> = \emptyset$. The vertical axis is used for the fractional value of β. The horizontal axis is used for the fractional values of the clocks in C_i. Thus, given a disjoint partition $\{C_0, \ldots, C_n\}$ of the clocks, we pick $n+1$ points on a line parallel to the horizontal axis, $\{\langle C_0^f, \mathsf{frac}(\beta)\rangle, \ldots, \langle C_n^f, \mathsf{frac}(\beta)\rangle\}$, with C_i^f being the fractional value of the clocks in the set C_i at $\widehat{\kappa}$. The following lemma states that the equivalence relation $\cong_e$ induces a time-abstract bisimulation.

Lemma 3. *Let Y, Y' be extended regions in a timed game structure $\widehat{[\![\mathcal{T}]\!]}$. Suppose player-$i$ has a move from $s_1 \in Y$ to $s_1' \in Y'$, for $i \in \{1, 2\}$. Then, for any $s_2 \in Y$, player-i has a move from s_2 to some $s_2' \in Y'$.*

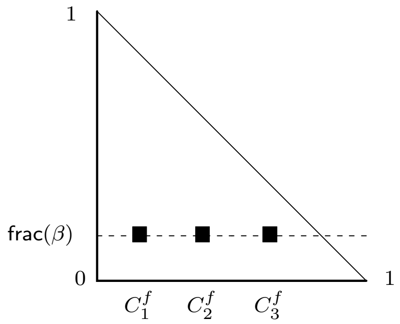

Fig. 1. An extended region with $C_< = C \cup \{z\}, C_= = \emptyset, C_> = \emptyset$

Lemma 4. *Let Y, Y_1', Y_2' be extended regions in a timed game structure $\widehat{[\![\mathcal{T}]\!]}$. Suppose player-i has a move from $s_1 \in Y$ to $s_1' \in Y'$, for $i \in \{1,2\}$. Then, one of the following cases must hold:*

1. *From all states $\widehat{s} \in Y$, player-1 has some move $m_1^{\widehat{s}}$ with $\widehat{\delta}(\widehat{s}, m_1^{\widehat{s}}) \in Y_1'$ such that for all moves $m_2^{\widehat{s}}$ of player-2 with $\widehat{\delta}(\widehat{s}, m_2^{\widehat{s}}) \in Y_2'$, we have $\mathsf{blame}_1(\widehat{s}, m_1^{\widehat{s}}, m_2^{\widehat{s}}, \widehat{\delta}(\widehat{s}, m_1^{\widehat{s}})) = \text{TRUE}$ and $\mathsf{blame}_2(\widehat{s}, m_1^{\widehat{s}}, m_2^{\widehat{s}}, \widehat{\delta}(\widehat{s}, m_2^{\widehat{s}})) = \text{FALSE}$.*
2. *From all states $\widehat{s} \in Y$, for all moves $m_1^{\widehat{s}}$ of player-1 with $\widehat{\delta}(\widehat{s}, m_1^{\widehat{s}}) \in Y_1'$, player-2 has some move $m_2^{\widehat{s}}$ with $\widehat{\delta}(\widehat{s}, m_2^{\widehat{s}}) \in Y_2'$ such that $\mathsf{blame}_2(\widehat{s}, m_1^{\widehat{s}}, m_2^{\widehat{s}}, \widehat{\delta}(\widehat{s}, m_2^{\widehat{s}})) = \text{TRUE}$.*

Intuitively, Lemma 4 says that either player-1 can force the game in one step from every state in Y so that the next state lies in Y', or player-2 can always foil player-1 from going to the extended region Y'. Thus moves to some extended regions always "beat" moves to other extended regions.

Corollary 1. *Let $X \subseteq \widehat{S}$ consist of a union of extended regions in a timed game structure $\widehat{[\![\mathcal{T}]\!]}$. Then $\mathsf{CPre}_1(X)$ is again a union of extended regions.*

Corollary 1 demonstrates that the sets in the fixpoint computation of the μ-calculus algorithm which computes winning states for player-1 for the reachability objective $\widehat{X}_p$ consist of unions of extended regions. Since the number of extended regions is finite, the algorithm terminates.

Theorem 2. *Let $s \in S$ be a state and $p \in \Sigma$ a proposition in a timed automaton game $\mathcal{T}$.*

1. *The minimum time for player-1 to visit p starting from s (denoted $T_{\min}(s,p)$) is computable in time $O\left((|L| \cdot \prod_{x \in C}(c_x + 1) \cdot |C+1|! \cdot 2^{|C|})^2 \cdot |C|\right)$.*
2. *For every region R of $[\![\mathcal{T}]\!]$, either there is constant $d_R \in \mathbb{N} \cup \{\infty\}$ such that for every state $s \in R$ we have $T_{\min}(s,p) = d_R$, or there is an integer constant d_R and a clock $x \in C$ such that for every state $s \in R$ we have $T_{\min}(s,p) = d_R - \mathsf{frac}(\kappa(x))$, where $\kappa(x)$ is the value of clock x in s.*

References

1. B. Adler, L. de Alfaro, and M. Faella. Average reward timed games. In *FORMATS 05*, LNCS 3829, pages 65–80. Springer, 2005.
2. R. Alur, M. Bernadsky, and P. Madhusudan. Optimal reachability for weighted timed games. In *ICALP 04*, LNCS 3142, pages 122–133. Springer, 2004.
3. R. Alur and D.L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994.
4. R. Alur and T.A. Henzinger. Modularity for timed and hybrid systems. In *CONCUR 97*, LNCS 1243, pages 74–88. Springer, 1997.
5. E. Asarin and O. Maler. As soon as possible: Time optimal control for timed automata. In *HSCC 99*, LNCS 1569, pages 19–30. Springer, 1999.
6. P. Bouyer, F. Cassez, E. Fleury, and K.G. Larsen. Optimal strategies in priced timed game automata. In *FSTTCS 04*, LNCS 3328, pages 148–160. Springer, 2004.
7. P. Bouyer, D. D'Souza, P. Madhusudan, and A. Petit. Timed control with partial observability. In *CAV 03*, LNCS 2725, pages 180–192. Springer, 2003.
8. T. Brihaye, V. Bruyère, and J.F. Raskin. On optimal timed strategies. In *FORMATS 05*, LNCS 3829, pages 49–64. Springer, 2005.
9. F. Cassez, A. David, E. Fleury, K.G. Larsen, and D. Lime. Efficient on-the-fly algorithms for the analysis of timed games. In *CONCUR 05*, LNCS 3653, pages 66–80. Springer, 2005.
10. F. Cassez, T.A. Henzinger, and J.-F. Raskin. A comparison of control problems for timed and hybrid systems. In *HSCC 02*, LNCS 2289, pages 134–148. Springer, 2002.
11. C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1(4):385–415, 1992.
12. L. de Alfaro, M. Faella, T.A. Henzinger, R. Majumdar, and M. Stoelinga. The element of surprise in timed games. In *CONCUR 03*, LNCS 2761, pages 144–158. Springer, 2003.
13. L. de Alfaro, T.A. Henzinger, and R. Majumdar. From verification to control: Dynamic programs for omega-regular objectives. In *LICS 01*, pages 279–290. IEEE Computer Society Press, 2001.
14. D. D'Souza and P. Madhusudan. Timed control synthesis for external specifications. In *STACS 02*, LNCS 2285, pages 571–582. Springer, 2002.
15. M. Faella, S. La Torre, and A. Murano. Dense real-time games. In *LICS 02*, pages 167–176. IEEE Computer Society, 2002.
16. T.A. Henzinger and P.W. Kopke. Discrete-time control for rectangular hybrid automata. *Theoretical Computer Science*, 221:369–392, 1999.
17. T.A. Henzinger and V.S. Prabhu. Timed alternating-time temporal logic. In *FORMATS 06*, LNCS 4202, pages 1–17. Springer, 2006.
18. M. Jurdziński and A. Trivedi. Reachability-time games on timed automata. In *ICALP 07*, LNCS. Springer, 2007.
19. O. Maler, A. Pnueli, and J. Sifakis. On the synthesis of discrete controllers for timed systems (an extended abstract). In *STACS 95*, pages 229–242, 1995.
20. A. Pnueli, E. Asarin, O. Maler, and J. Sifakis. Controller synthesis for timed automata. In *Proc. System Structure and Control*. Elsevier, 1998.
21. R. Segala, R. Gawlick, J.F. Søgaard-Andersen, and N.A. Lynch. Liveness in timed and untimed systems. *Inf. Comput.*, 141(2):119–171, 1998.
22. H. Wong-Toi and G. Hoffmann. The control of dense real-time discrete event systems. In *Proc. of 30th Conf. Decision and Control*, pages 1527–1528, 1991.

A Appendix

Proof of Theorem 1

We restrict our attention to strategies for plays starting from state s. The proof of the theorem relies on Lemmas 5,6 and 7.

Lemma 5. *Consider a timed game structure $\mathcal{G}$ and a state $s \in S$. Let $\pi_1 \in \Pi_1^R$ and $\pi_2^R \in \Pi_2^R$ be player-1 and player-2 receptive strategies, and let $\pi_2 \in \Pi_2$ be any player-2 strategy. Consider a player-2 strategy π_2^* be defined as, $\pi_2^*(r[0..k]) = \pi_2(r[0..k])$ for all run prefixes $r[0..k]$ of $\mathsf{Outcomes}(s, \pi_1, \pi_2)$, and $\pi_2^*(r[0..k]) = \pi_2^R(r[k'..k])$ otherwise, where k' is the first position such that $r[0..k']$ is not a run prefix of $\mathsf{Outcomes}(s, \pi_1, \pi_2)$. Then, π_2^* is a receptive strategy.*

Proof. Intuitively, the strategy π_2^* acts like π_2 on runs of $\mathsf{Outcomes}(s, \pi_1, \pi_2)$, and like π_2^R otherwise. Consider any player-1 strategy $\pi_1' \in \Pi_1$. If $\pi_1' = \pi_1$, then $\mathsf{Outcomes}(s, \pi_1', \pi_2^*) = \mathsf{Outcomes}(s, \pi_1, \pi_2) \subseteq \mathsf{Timediv}$. Suppose $\pi_1' \neq \pi_1$. Let $k' \geq 0$ be the first step in the game (with player-2 strategy π_2) which witnesses the fact that $\pi_1' \neq \pi_1$, that is, 1)for all runs $r \in \mathsf{Outcomes}(s, \pi_1', \pi_2)$, we have $r[0..k'-1]$ is a run prefix of some run in $\mathsf{Outcomes}(s, \pi_1, \pi_2)$ and 2)for some run $r' \in \mathsf{Outcomes}(s, \pi_1', \pi_2)$, we have $r'[0..k']$ to be not a run prefix of any run in $\mathsf{Outcomes}(s, \pi_1, \pi_2)$.

We thus have $r[0..k'-1]$ to be a run prefix of some run in $\mathsf{Outcomes}(s, \pi_1, \pi_2)$ for every $r \in \mathsf{Outcomes}(s, \pi_1', \pi_2^*)$. Consider state $s_{k'} = r^*[k']$ for $r^* \in \mathsf{Outcomes}(s, \pi_1', \pi_2^*)$. We have $r^*[0..k'-1] = r[0..k'-1]$ for some run $r \in \mathsf{Outcomes}(s, \pi_1', \pi_2)$. After this point (ie., from $r^*[0..k']$ onwards), the strategy π_2^* behaves like π_2^R when "started" from $s_{k'}$. Since π_2^R is a receptive player-2 strategy, we have $\mathsf{Outcomes}(s_{k'}, \pi_1', \pi_2^*) \subseteq \mathsf{Timediv} \cup \mathsf{Blameless}_2$. Thus, $\mathsf{Outcomes}(s, \pi_1', \pi_2^*) \subseteq \mathsf{Timediv} \cup \mathsf{Blameless}_2$ (finite prefixes of runs do not change membership in these sets). Hence π_2^* is a receptive player-2 strategy.

Lemma 6. *Consider a timed game structure $\mathcal{G}$ and a state $s \in S$. We have,*

$$\inf_{\pi_1 \in \Pi_1} \sup_{\pi_2 \in \Pi_2} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \{T_{\text{visit}}^{\text{UR}}(r,p)\} = \inf_{\pi_1 \in \Pi_1^R} \sup_{\pi_2 \in \Pi_2} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \{T_{\text{visit}}^{\text{UR}}(r,p)\}$$

Proof. Consider any $\pi_1 \in \Pi_1 \setminus \Pi_1^R$. There exists $\pi_2 \in \Pi_2$ such that $\mathsf{Outcomes}(s, \pi_1, \pi_2) \not\subseteq \mathsf{Timediv} \cup \mathsf{Blameless}_1$. Thus, $\inf_{\pi_1 \in \Pi_1 \setminus \Pi_1^R} \sup_{\pi_2 \in \Pi_2} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \{T_{\text{visit}}^{\text{UR}}(r,p)\} = \infty$.

Lemma 7. *Consider a timed game structure $\mathcal{G}$ and a state $s \in S$. For every player-1 receptive strategy $\pi_1 \in \Pi_1^R$, we have $\sup_{\pi_2 \in \Pi_2} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \{T_{\text{visit}}^{\text{UR}}(r,p)\} = \sup_{\pi_2 \in \Pi_2^R} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \{T_{\text{visit}}^{\text{UR}}(r,p)\}$.*

Proof. Let $\pi_2 \in \Pi_2$.

Consider $r \in \mathsf{Outcomes}(s, \pi_1, \pi_2)$. Since π_1 is receptive, we cannot have $r \notin \mathsf{Timediv}$ and $r \notin \mathsf{Blameless}_1$.

Suppose $r \notin \mathsf{Timediv}$. Then $r \in \mathsf{Blameless}_1$. In this case, $0 = T^{\mathrm{UR}}_{\mathrm{visit}}(r,p) \leq T^{\mathrm{UR}}_{\mathrm{visit}}(r',p)$ for any $r' \in \mathsf{Outcomes}(s,\pi_1,\pi_2^R)$ and π_2^R any player-2 receptive strategy (as we have a well-formed time game structure, there exists some receptive strategy π_2^R).

Suppose $r \in \mathsf{Timediv}$ and r does not visit p. Consider the strategy π_2^* which acts like π_2 on $\mathsf{Outcomes}(s,\pi_1,\pi_2$, and like π_2^R otherwise, as formally defined in Lemma 5. We have π_2^* to be receptive. Clearly $r \in \mathsf{Outcomes}(s,\pi_1,\pi_2^*)$ does not visit p, and hence $\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\} = \sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2^*)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\} = \infty$.

Finally, let r visit p and be in $\mathsf{Timediv}$ for all $r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)$. Let π_2^* be a player-2 receptive strategy as in Lemma 5. We again have $\mathsf{Outcomes}(s,\pi_1,\pi_2) = \mathsf{Outcomes}(s,\pi_1,\pi_2^*)$, and hence $\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\} = \sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2^*)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\}$.

Thus, $\sup_{\pi_2\in\Pi_2}\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\} = \sup_{\pi_2\in\Pi_2^R}\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\}$.

Lemmas 6 and 7 together imply

$$\inf_{\pi_1\in\Pi_1}\sup_{\pi_2\in\Pi_2}\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\} = \inf_{\pi_1\in\Pi_1^R}\sup_{\pi_2\in\Pi_2^R}\sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T^{\mathrm{UR}}_{\mathrm{visit}}(r,p)\}$$

Theorem 1 follows from the fact that for $\pi_1 \in \Pi_1^R, \pi_2 \in \Pi_2^R$ and $r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)$, we have $T^{\mathrm{UR}}_{\mathrm{visit}}(r,p) = T_{\mathrm{visit}}(r,p)$.

Proof of Proposition 1

The first claim is a corollary of the following proposition:

Proposition 2. *Consider the set S_p for a proposition p in a timed game structure $[\![\mathfrak{T}]\!]$.*

1. *If a run r of $[\![\mathfrak{T}]\!]$ visits S_p at time $t \leq M$, then, the run $\widehat{r}_{0,\beta}$ visits $S_p \times \mathbb{R}_{[0,1)} \times \{0\} \times \{\text{TRUE}, \text{FALSE}\}^2$, for $\beta = t$.*
2. *If for some $\beta \in \mathbb{R}$, a run $\widehat{r}$ of $\widehat{[\![\mathfrak{T}]\!]}$ with $\widehat{r}[0] = \langle s, 0, \beta, \text{FALSE}, \text{FALSE}\rangle$ visits $S_p \times \mathbb{R}_{[0,1)} \times \{0\} \times \{\text{TRUE}, \text{FALSE}\}^2$, then the corresponding run $r = \widehat{r} \downarrow \mathfrak{T}$ of $[\![\mathfrak{T}]\!]$ visits S_p at time $t = \beta$.*

Proposition 2 is a straightforward result of the fact that β is kept decrementing at rate -1 till it hits 0.

The second claim of Proposition 1 essentially follows from the fact that the additional components in the states do not help the players in creating more powerful strategies.

$T_{\min}(s,p)$

$= \inf_{\pi_1\in\Pi_1^R}\sup_{\pi_2\in\Pi_2^R}\ \sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\{T_{\mathrm{visit}}([\![\mathfrak{T}]\!],r,p)\}$

$= \inf_{\pi_1\in\Pi_1^R}\sup_{\pi_2\in\Pi_2^R}\ \sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\left\{\begin{cases}\infty \text{ if } r \text{ does not visit } p;\\ \inf\{\beta \mid \beta\in\mathbb{R}_{[0,M]} \text{ and } \widehat{r}_{0,\beta} \text{ visits the set } \widehat{X}_p\} \text{ o.w.}\end{cases}\right\}$

$= \inf_{\pi_1\in\Pi_1^R}\sup_{\pi_2\in\Pi_2^R}\ \sup_{r\in\mathsf{Outcomes}(s,\pi_1,\pi_2)}\ \inf_{\beta\in\mathbb{R}_{[0,M]}}$

$$\left\{ g(r,\beta) \,\middle|\, g(r,\beta) = \infty \text{ if } \widehat{r}_{0,\beta} \text{ does not visit } \widehat{X}_p;\ \beta \text{ otherwise} \right\}$$

$$= \inf_{\beta \in \mathbb{R}_{[0,M]}} \inf_{\pi_1 \in \Pi_1^R} \sup_{\pi_2 \in \Pi_2^R} \sup_{r \in \mathsf{Outcomes}(s,\pi_1,\pi_2)} \left\{ g(r,\beta) \,\middle|\, g(r,\beta) = \infty \text{ if } \widehat{r}_{0,\beta} \text{ does not visit } \widehat{X}_p;\ \beta \text{ otherwise} \right\}$$

Now, considering plays in $\widehat{[\![\mathcal{T}]\!]}$ which start from state $\widehat{s} = \langle s, z, \beta, \mathit{tick}, \mathit{bl}_1 \rangle$, every strategy $\widehat{\pi_i} \in \widehat{\Pi_i}$ is equivalent to a strategy $\pi_i \in \Pi_i$ in which player-i "guesses" the values of $z, \beta, \mathit{tick}, \mathit{bl}_1$. Once these initial values have been guessed, each player can keep on deterministically updating the values at each step. Hence observation of the additional components in states of $\widehat{[\![\mathcal{T}]\!]}$ do not help the players in their strategies. Therefore,

$$T_{\min}(s,p) = \inf_{\beta \in \mathbb{R}_{[0,M]}} \inf_{\widehat{\pi_1} \in \widehat{\Pi_1}^R} \sup_{\widehat{\pi_2} \in \widehat{\Pi_2}^R} \sup_{\widehat{r}_{0,\beta} \in \mathsf{Outcomes}(s,\widehat{\pi_1},\widehat{\pi_2})} \left\{ g(r,\beta) \,\middle|\, g(r,\beta) = \infty \text{ if } \widehat{r}_{0,\beta} \text{ does not visit } \widehat{X}_p;\ \beta \text{ otherwise} \right\}$$

Proof of Lemma 2

We first note that the states in $\widehat{X}_p$ can be absorbing as $\widehat{[\![\mathcal{G}]\!]}$ is a well-formed time game structure, and hence player-1 has a receptive strategy which does not block time when the game starts at state $\widehat{s}$ for every state $\widehat{s} \in \widehat{X}_p$. Consider a run $\widehat{r}$ such that $\widehat{r}$ visits $\widehat{X}_p$. We can assume without loss of generality that either time diverges in $\widehat{r}$, or time converges but player-1 is not to blame (player-1 can play a receptive strategy upon reaching $\widehat{X}_p$). Thus this run satisfies the winning condition for player-1. And since $\widehat{X}_p$ is absorbing in our parity game, we see 1 only finitely often.

Consider a run $\widehat{r}$ such that $\widehat{r}$ does not visit $\widehat{X}_p$. Let time diverge in this run. This run violates the winning condition for player-1, and correspondingly we also see the index 1 infinitely often (due to *tick* being true infinitely often). Now let time converge in this run (so *tick* is true only finitely often). If player-1 is to blame for blocking time, then the index 1 will again be true infinitely often. If player-1 is not to blame, then bl_1 will only be true finitely often in this run, and hence we will see the index 1 only finitely often.

Proof of Lemma 3

The result can be proved using the following lemma:

Lemma 8. *Let Y, Y' be extended regions in a timed game structure $\widehat{[\![\mathcal{T}]\!]}$. Consider a state $\widehat{s} \in Y$ and $t \in \mathbb{R}_{>0}$ Suppose $(0,t] = T^Y \cup T^{Y'}$, such that for all $\tau \in T^Y$ we have $\widehat{s} + \tau \in Y$, and for all $\tau \in T^{Y'}$ we have $\widehat{s} + \tau \in Y'$ ($Y \rightarrow Y'$ is the first extended region change due to the passage of time). Then, for all states $\widehat{s}_2 \in Y$, there exists $t_2 \in \mathbb{R}_{>0}$ such that for some $T_2^Y, {T^{Y'}}_2$ with $(0,t_2] = T_2^Y \cup {T^{Y'}}_2$, for all $\tau_2 \in T_2^Y$ we have $\widehat{s}_2 + \tau_2 \in Y$, and for all $\tau_2 \in {T^{Y'}}_2$ we have $\widehat{s}_2 + \tau_2 \in Y'$.*

Proof. We outline a sketch of the proof. For simplicity, consider the values of each clock x to be less than $C_x + 1$. We look at the time successors of states $\widehat{s}$ in Y. The following cases for $Y = \langle l, \mathit{tick}, \mathit{bl}_1, h, \mathcal{P}, \beta_i, \beta_f, C_<, C_=, C_> \rangle$ can arise:

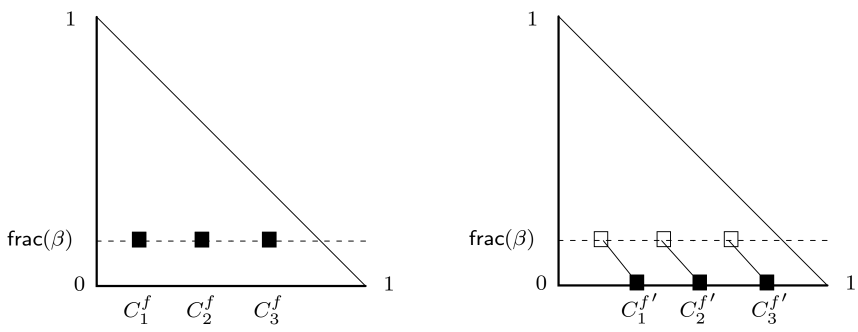

Fig. 2. An extended region with $C_< = C \cup \{z\}, C_= = \emptyset$ and its time successor.

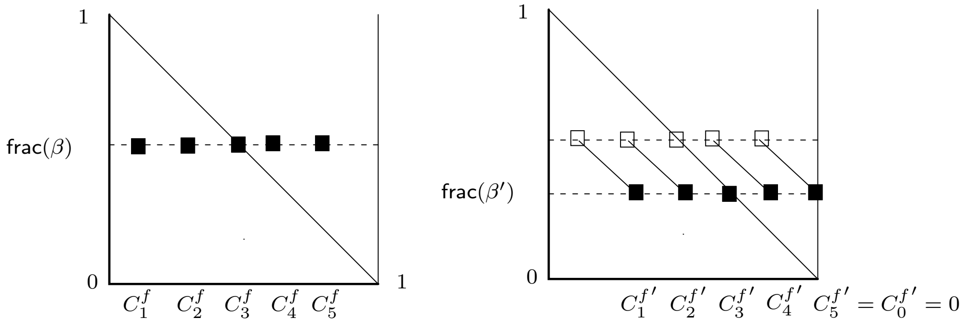

Fig. 3. An extended region with $C_< \neq \emptyset, C_= \neq \emptyset, C_> \neq \emptyset$ and its time successor.

Case 1 $C_0 = \emptyset, \beta_i \in \mathbb{N} \cap \{0, \ldots, M\}, \beta_f = \text{TRUE}, C_< = C \cup \{z\}, C_= = \emptyset, C_> = \emptyset$.

For any state in Y, the next extended region Y' can *only* be $\langle l, tick, bl_1, h, \mathcal{P}, \beta_i, \beta'_f = \text{FALSE}, C_<, C_=, C_> \rangle$, which is hit after a time of $\mathsf{frac}(\beta_f)$ (note that $C_n^f + \mathsf{frac}(\beta) < 1$ implies $\mathcal{P}$ is going to be unchanged in the time successor extended region).

Case 2 $C_0 = \emptyset, \beta_i \in \mathbb{N} \cap \{0, \ldots, M\}, \beta_f = \text{TRUE}, C_< \neq \emptyset, C_= \neq \emptyset, C_> \neq \emptyset$. Pictorially, this can be depicted as in Fig. 3.

Consider any state in Y. The extended region changes after a time of $1 - C_n^f$. The new state then lies in an extended region such that $C'_i = C_i$ for $0 < i < n$, and $C'_0 = C_n$. Also, $C_i^{f'} = C_i^f + (1 - C_n^f)$ for $0 < i < n$, and $\mathsf{frac}\,\beta' = \mathsf{frac}(\beta) - (1 - C_n^f)$. We also have that if $C_i^f + \mathsf{frac}(\beta) \sim 1$, then $C_i^{f'} + \mathsf{frac}\,\beta' = C_i^f + \mathsf{frac}(\beta) \sim 1$ for $\sim \in \{<, =, >\}, 0 < i < n$. Thus the new state lies in the region $\langle l, tick', bl_1, h', \mathcal{P}' = \{C'_0, \ldots C'_{n-1} \mid C'_i = C_i \text{ for } 0 < i < n, C'_0 = C_n\}, \beta_i, \beta_f, C'_< = C_< \cup C_n, C_= \rangle$, with $tick' = \text{TRUE}$ iff $z \in C_n$, and h' is h with the integer values for clocks in $C_n \setminus \{z\}$ incremented by

1. This analysis holds for *all* the states in Y. Thus the extended region Y' following Y is unique.

Case 3 $C_0 \neq \emptyset, \beta_i \in \mathbb{N} \cap \{0, \ldots, M\}, \beta_f = \text{TRUE}$

All the states in Y then move to $\langle l, tick, bl_1, h, \mathcal{P}' = \{C'_0, \ldots, C'_{n+1} \mid C'_0 = \emptyset \text{ and } C'_{i+1} = C_i, 0 \leq i \leq n\}, \beta_i, \beta_f, C_<, C_=, C_> \rangle$.

Case 4 $C_0 \neq \emptyset, \beta_i \in \mathbb{N} \cap \{1, \ldots, M\}, \beta_f = \text{FALSE}$

The time successor in this case is $\langle l, tick, bl_1, \mathcal{P}' = \{C'_0, \ldots, C'_{n+1} \mid C'_0 = \emptyset \text{ and } C'_{i+1} = C_i, 0 \leq i \leq n\}, \beta'_i = \beta_i - 1, \beta'_f = \text{TRUE}, C'_<, C'_=, C'_> \rangle$. We show $C'_<, C'_=, C'_>$ to be unique as follows: the new state $\widehat{s} + t$ has the constraints 1)$\mathsf{frac}(\beta') = 1 - t$ and 2)${C^f_{i+1}}' = C^f_i + t$ for $i \leq n$. Thus, $\mathsf{frac}(\beta') + {C^f_{i+1}}' = (1 - t) + C^f_i + t = 1 + C^f_i$. Hence, $C'_< = \emptyset$ and $C'_= = C'_1 = C_0$.

Case 5 $\beta_i = 0, \beta_f = \text{FALSE}$

We get $\beta' = \perp$ in the next state (and hence $C_< = C_= = \emptyset, \beta_i = \perp, \beta_f = \text{FALSE}$). The rest of the components of the extended region have a unique value as in the time successors of standard regions.

Case 6 $\beta_i = \perp$

The value of $\mathcal{P}'$ gets updated as in the time successors of standard regions.

The analysis of the remaining cases proceeds in a similar vein to the above cases.

Proof of Theorem 2

We prove the second part of the claim.

Proof. Let M be the upper bound on $T_{\min}(s, p)$ as in Lemma 1 if $T_{\min}(s, p) < \infty$, and $M = 1$ otherwise. From the comments after Corollary 1, the states in $\widehat{S}$ from which player-1 has a winning strategy for reaching $\widehat{X}_p$ are computable, and consist of a union of extended regions $\cup_{k=1}^{n} Y_k$. Suppose this union is non-empty. Using Proposition 1, the minimum time for player-1 to reach p from s is then $\min_k \left\{ \inf \left\{ \beta \mid \beta \in \mathbb{R}_{[0,M]} \text{ and } \langle s, 0, \beta, \text{FALSE}, \text{FALSE} \rangle \in Y_k \right\} \right\}$. Note that $s = \langle l, \kappa \rangle$ is fixed here, only β can be varied. We also have that $\inf\{\beta \mid \beta \in \mathbb{R}_{[0,M]} \text{ and } \langle \langle l, \kappa \rangle, 0, \beta, \text{FALSE}, \text{FALSE} \rangle \in Y_k\}$ is equal to (letting $Y_k = \langle l, \text{FALSE}, \text{FALSE}, h, \mathcal{P}, \beta_i, \beta_f, C_<, C_=, C_> \rangle$):

1. An integer when $C_> = C_= = \emptyset$ or when $\beta_f = \text{FALSE}$. The infimum value for β is reached when $\beta_f = \text{FALSE}$ (for then the set of β's is a singleton). Thus, player-1 has an optimal strategy when $\beta_f = \text{FALSE}$.
2. $d_k - \mathsf{frac}(\kappa(x))$ when $C_= = C_j \neq \emptyset$, and where $x \in C_j$. The infimum value is actually attained by player-1 with some strategy π_1 in this case.
3. $d_k - \mathsf{frac}(\kappa(x))$ when $C_= = \emptyset, C_> \neq \emptyset$, where $x \in C_j$ for $C_> = \{C_j, \ldots, C_n\}$. The infimum value is not attained by player-1 in this case – he can only get arbitrarily close to the optimum.

Note that $z \in C_0$ in every Y_k (for, $\widehat{\kappa}(z) = 0$). Finally, $\min_k\{e_k \mid e_k = d_k \text{ or } d_k - x_k\}$ is again an expression of the form d_r or $d_r - x$ over a region.