

Computing Equilibria in Two-Player Timed Games via Turn-Based Finite Games*

Patricia Bouyer, Romain Brenguier, and Nicolas Markey

LSV, CNRS & ENS Cachan, France
{bouyer, brenguier, markey}@lsv.ens-cachan.fr

Abstract. We study two-player timed games where the objectives of the two players are not opposite. We focus on the standard notion of Nash equilibrium and propose a series of transformations that builds two finite turn-based games out of a timed game, with a precise correspondence between Nash equilibria in the original and in final games. This provides us with an algorithm to compute Nash equilibria in two-player timed games for large classes of properties.

1 Introduction

Timed games. Game theory (especially games played on graphs) has been used in computer science as a powerful framework for modelling interactions in embedded systems [16, 13]. Over the last fifteen years, games have been extended with the ability to depend on timing informations, taking advantage of the large development of timed automata [1]. Adding timing constraints allows for a more faithful representation of reactive systems, while preserving decidability of several important properties, such as the existence of a winning strategy for one of the agents to achieve her goal, whatever the other agents do [3]. Efficient algorithms exist and have been implemented, *e.g.* in the tool Uppaal-Tiga [4].

Zero sum vs. non-zero sum games. In this purely antagonist view, games can be seen as two-player games, where one agent plays against another one. Moreover, the objectives of those two agents are opposite: the aim of the second player is simply to prevent the first player from winning her own objective. More generally, a (positive or negative) payoff can be associated with each outcome of the game, which can be seen as the amount the second player will have to pay to the first player. Those games are said to be zero-sum.

In many cases, however, games can be non-zero-sum: the objectives of the two players are then no more complementary, and the aim of one player is no more to prevent the other player from winning. Such games appear *e.g.* in various problems in telecommunications, where the agents try to send data on a network [12]. Focusing only on surely-winning strategies in this setting may then be too narrow: surely-winning strategies must be winning against any behaviour of the other player, and do not consider the fact that the other player also tries to achieve her own objective.

* This work is partly supported by projects DOTS (ANR-06-SETI-003), QUASIMODO (FP7-ICT-STREP-214755) and GASICS (ESF-EUROCORES LogiCCC).

Nash equilibria. In the non-zero-sum game setting, it is then more interesting to look for equilibria. One of the most-famous and most-studied notion of equilibrium is that proposed by Nash in 1950 [14]: a *Nash equilibrium* is a behaviour of the players in which they act rationally, in the sense that no player can get a better payoff if she, alone, modifies her strategy [14]. Notice that a Nash equilibrium needs not exist in general, and may not be optimal, in the sense that several equilibria can coexist, and may have very different payoffs.

Our contribution. We extend the standard notion of Nash equilibria to timed games, where non-determinism naturally arises and has to be taken into account. We propose a whole chain of transformations that builds, given a two-player timed game, two turn-based finite games which, in some sense that we will make precise, preserve Nash equilibria. The first transformation consists in building a finite concurrent game with non-determinism based on the classical region abstraction; the second transformation decouples this concurrent game into two concurrent games, one per player: in each game, the preference relation of one of the players is simply dropped, but we have to consider “joint” equilibria. The last two transformations work on each of the two copies of the concurrent game: the first one solves the non-determinism by giving an advantage to the associated player, and the last one makes use of this advantage to build a turn-based game equivalent to the original concurrent game. This chain of transformations is valid for the whole class of two-player timed games, and Nash equilibria are preserved for a large class of objectives, for instance ω -regular objectives¹. These transformations allow to recover some known results about zero-sum games, but also to get new decidability results for Nash equilibria in two-player timed games.

Related work. Nash equilibria (and other related solution concepts such as *subgame-perfect equilibria*, *secure equilibria*, ...) have recently been studied in the setting of (untimed) games played on a graph [8–10, 15, 17–20]. None of them, however, focuses on timed games. In the setting of concurrent games, mixed strategies (*i.e.*, strategies involving probabilistic choices) are arguably more relevant than pure (*i.e.*, non-randomized) strategies. However, adding probabilities to timed strategies involves several important technical issues (even in zero-sum non-probabilistic timed games), and we defer the study of mixed-strategy Nash equilibria in two-player timed games to future works.

For lack of space, proofs are omitted and can be found in [5].

2 Preliminaries

2.1 Timed games

A *valuation* over a finite set of clocks Cl is a mapping $v: \text{Cl} \rightarrow \mathbb{R}_+$. If v is a valuation and $t \in \mathbb{R}_+$, then $v + t$ is the valuation that assigns to each $x \in \text{Cl}$ the value $v(x) + t$. If v is a valuation and $Y \subseteq \text{Cl}$, then $[Y \leftarrow 0]v$ is the valuation that assigns 0 to each $y \in Y$ and $v(x)$ to each $x \in \text{Cl} \setminus Y$. A *clock constraint* over Cl is a formula built on the grammar: $\mathcal{C}(\text{Cl}) \ni g ::= x \sim c \mid g \wedge g$, where x ranges over Cl , $\sim \in \{<, \leq, =, \geq, >\}$, and c is an integer. The semantics of clock constraints over valuations is natural, and we omit it.

¹ In the general case, the undecidability results on (zero-sum) priced timed games entail undecidability of the existence of Nash equilibria.

Definition 1. A timed automaton is a tuple $\langle \text{Loc}, \text{Cl}, \text{Inv}, \text{Trans} \rangle$ such that:

- Loc is a finite set of locations;
- Cl is a finite set of clocks;
- $\text{Inv}: \text{Loc} \rightarrow \mathcal{C}(\text{Cl})$ assigns an invariant to each location;
- $\text{Trans} \subseteq \text{Loc} \times \mathcal{C}(\text{Cl}) \times 2^{\text{Cl}} \times \text{Loc}$ is the set of transitions.

We assume the reader is familiar with timed automata [1], and in particular with states (pairs $(\ell, v) \in \text{Loc} \times \mathbb{R}_+^{\text{Cl}}$ such that $v \models \text{Inv}(\ell)$), runs (seen as infinite sequences of states for our purpose), etc. We now define the notion of two-player timed games. The two players will be called player 1 and player 2. Our definition follows that of [11].

Definition 2. A (two-player) timed game is a tuple $\mathcal{G} = \langle \text{Loc}, \text{Cl}, \text{Inv}, \text{Trans}, \text{Owner}, (\preceq_1, \preceq_2) \rangle$ where:

- $\langle \text{Loc}, \text{Cl}, \text{Inv}, \text{Trans} \rangle$ is a timed automaton;
- $\text{Owner}: \text{Trans} \rightarrow \{1, 2\}$ assigns a player to each transition;
- for each $i \in \{1, 2\}$, $\preceq_i \subseteq (\text{Loc} \times \mathbb{R}_+^{\text{Cl}})^\omega \times (\text{Loc} \times \mathbb{R}_+^{\text{Cl}})^\omega$ is a quasi-order on runs of the timed automaton, called the preference relation for player i .

A timed game is played as follows: from each state of the underlying timed automaton (starting from an initial state $s_0 = (\ell, \mathbf{0})$, where $\mathbf{0}$ maps each clock to zero), each player chooses a nonnegative real number d and a transition δ , with the intended meaning that she wants to delay for d time units and then fire transition δ . There are several (natural) restrictions on these choices:

- spending d time units in ℓ must be allowed² i.e., $v + d \models \text{Inv}(\ell)$;
- $\delta = (\ell, g, z, \ell')$ belongs to the current player (given by function **Owner**);
- the transition is fireable after d time units (i.e., $v + d \models g$), and the invariant is satisfied when entering ℓ' (i.e., $[z \leftarrow 0](v + d) \models \text{Inv}(\ell')$).

When there is no such possible choice for a player (for instance if there is no transition from ℓ belonging to that player), she chooses a special move, denoted by \perp .

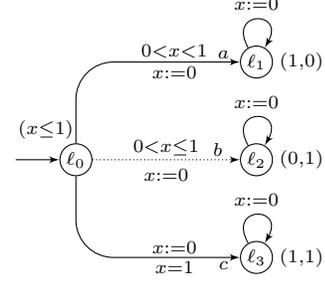
From a state (ℓ, v) and given a choice (m_1, m_2) for the two players, with $m_i \in (\mathbb{R}_+ \times \text{Trans}) \cup \{\perp\}$, an index i_0 such that $d_{i_0} = \min\{d_i \mid m_i = (d_i, \delta_i) \text{ and } i \in \{1, 2\}\}$ is selected (non-deterministically if both delays are identical), and the corresponding transition $\delta_{i_0} = (\ell, g, z, \ell')$ is applied, leading to a new state $(\ell', [z \leftarrow 0](v + d_{i_0}))$. To ensure well-definedness of the above semantics we assume in the sequel that timed games are *non-blocking*, that is, for any reachable state (ℓ, v) , at least one player has an allowed transition (this avoids that both players play the special action \perp).

The outcome of such a game when players have fixed their various choices is a run of the underlying timed automaton, that is an element of $(\text{Loc} \times \mathbb{R}_+^{\text{Cl}})^\omega$, and possible outcomes are compared by each player using their preference relations. In the examples, we will define the preference relation of a player by assigning a value (called a *payoff*) to each possible outcome of the game, and the higher the payoff, the better the run in the preference relation.

² Formally, this should be written $v + d' \models \text{Inv}(\ell)$ for all $0 \leq d' \leq d$, but this is equivalent to having only $v \models \text{Inv}(\ell)$ and $v + d \models \text{Inv}(\ell)$ since invariants are convex.

This semantics can naturally be formalized in terms of an infinite-state non-deterministic concurrent game and strategies, that we will detail in the next section.

Example 1. We give an example of a timed game, that we will use as a running example: consider the timed game \mathcal{G} on the right. When relevant the name of a transition is printed on the corresponding edge. Owners of the transitions are specified as follows: player 1 plays with plain edges, whereas player 2 plays with dotted edges. On the right of these locations we indicate payoffs for the two players (if a play ends up in ℓ_1 , player 1 gets payoff 1, whereas player 2 gets payoff 0). Hence player 1 will prefer runs ending in ℓ_1 or ℓ_3 than runs ending in ℓ_2 .



2.2 Concurrent games

In this section we define two-player concurrent games, which we then use to encode the formal semantics of timed games. A *transition system* is a 2-tuple $\mathcal{S} = \langle \text{States}, \text{Edg} \rangle$ where **States** is a (possibly uncountable) set of states, and $\text{Edg} \subseteq \text{States} \times \text{States}$ is the set of transitions. A *path* π in \mathcal{S} is a non-empty sequence $(s_i)_{0 \leq i < n}$ (where $n \in \mathbb{N} \cup \{+\infty\}$) of states of \mathcal{S} such that $(s_i, s_{i+1}) \in \text{Edg}$ for all $i < n - 1$. The *length* of π , denoted by $|\pi|$ is $n - 1$. The set of finite paths (also called *histories* in the sequel) of \mathcal{S} is denoted by³ $\text{Hist}_{\mathcal{S}}$, the set of infinite paths (also called *plays*) of \mathcal{S} is denoted by $\text{Play}_{\mathcal{S}}$, and $\text{Path}_{\mathcal{S}} = \text{Hist}_{\mathcal{S}} \cup \text{Play}_{\mathcal{S}}$ is the set of paths of \mathcal{S} . Given a path $\pi = (s_i)_{0 \leq i < n}$ and an integer $j \leq |\pi|$, the *j-th prefix* of π , denoted by $\pi_{\leq j}$, is the finite path $(s_i)_{0 \leq i < j+1}$. If $\pi = (s_i)_{0 \leq i < n}$ is a history, we write $\text{last}(\pi) = s_{|\pi|}$.

We extend the definition of concurrent games given e.g. in [2] with non-determinism:

Definition 3. A (two-player non-deterministic) concurrent game is a tuple $\mathcal{G} = \langle \text{States}, \text{Edg}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_1, \preceq_2) \rangle$ in which:

- $\langle \text{States}, \text{Edg} \rangle$ is a transition system;
- **Act** is a (possibly uncountable) set of actions;
- **Mov**: $\text{States} \times \{1, 2\} \rightarrow 2^{\text{Act}} \setminus \{\emptyset\}$ is a mapping indicating the actions available to each player in a given state;
- **Tab**: $\text{States} \times \text{Act}^2 \rightarrow 2^{\text{Edg}} \setminus \{\emptyset\}$ associates to each state and each pair of actions the set of resulting edges. It is required that if $(s', s'') \in \text{Tab}(s, (m_1, m_2))$, then $s' = s$.
- for each $i \in \{1, 2\}$, $\preceq_i \subseteq \text{States}^\omega \times \text{States}^\omega$ is a quasi-order called the preference relation for player i .

A *deterministic* concurrent game is a concurrent game where $\text{Tab}(s, (m_1, m_2))$ is a singleton for every $s \in \text{States}$ and $(m_1, m_2) \in \text{Mov}(s, 1) \times \text{Mov}(s, 2)$. A *turn-based* game is a concurrent game for which there exists a mapping **Owner**: $\text{States} \rightarrow \{1, 2\}$ such that, for every state $s \in \text{States}$, the set $\text{Mov}(s, i)$ is a singleton unless **Owner**(s) = i .

³ For this and the following definitions, we explicitly mention the underlying transition system as a subscript. In the sequel, we may omit this subscript when the transition system is clear from the context.

In a concurrent game, from some state s , each player i selects one action m_i among its set $\text{Mov}(s, i)$ of allowed actions (the resulting pair (m_1, m_2) is called a *move*). This results in a set of edges $\text{Tab}(s, (m_1, m_2))$, one of which is applied and gives the next state of the game. In the sequel, we abusively write $\text{Hist}_{\mathcal{G}}$, $\text{Play}_{\mathcal{G}}$ and $\text{Path}_{\mathcal{G}}$ for the corresponding set of paths in the underlying transition system of \mathcal{G} . We also write $\text{Hist}_{\mathcal{G}}(s)$, $\text{Play}_{\mathcal{G}}(s)$ and $\text{Path}_{\mathcal{G}}(s)$ for the respective subsets of paths starting in state s .

Definition 4. Let \mathcal{G} be a concurrent game, and $i \in \{1, 2\}$. A strategy for player i is a mapping $\sigma_i: \text{Hist}_{\mathcal{G}} \rightarrow \text{Act}$ such that $\sigma_i(\pi) \in \text{Mov}(\text{last}(\pi), i)$ for all $\pi \in \text{Hist}_{\mathcal{G}}$. A strategy profile is a pair (σ_1, σ_2) where σ_i is a player- i strategy. We write $\text{Strat}_{\mathcal{G}}^i$ for the set of strategies of player i in \mathcal{G} , and $\text{Prof}_{\mathcal{G}}$ for the set of strategy profiles in \mathcal{G} .

Notice that we only consider non-randomized (*pure*) strategies in this paper.

Let \mathcal{G} be a concurrent game, $i \in \{1, 2\}$, and σ_i be a player i -strategy. A path $\pi = (s_j)_{0 \leq j \leq |\pi|}$ is *compatible* with the strategy σ_i if, for all $k \leq |\pi| - 1$, there exists a pair of actions $(m_1, m_2) \in \text{Act}^2$ such that $m_j \in \text{Mov}(s_k, j)$ for all $j \in \{1, 2\}$, $m_i = \sigma_i(\pi_{\leq k})$, and $(s_k, s_{k+1}) \in \text{Tab}(s_k, (m_1, m_2))$. A path π is compatible with a strategy profile (σ_1, σ_2) whenever it is compatible with both strategies σ_1 and σ_2 . We write $\text{Out}_{\mathcal{G}, s}(\sigma_i)$ (resp. $\text{Out}_{\mathcal{G}, s}(\sigma_1, \sigma_2)$) for the set of paths from s (also called *outcomes*) in \mathcal{G} that are compatible with strategy σ_i (resp. strategy profile (σ_1, σ_2)). Notice that, in the case of deterministic concurrent games, a strategy profile has a single infinite outcome. This might not be the case for non-deterministic concurrent games.

Given a move (m_1, m_2) and a new action m' for player i , we write $(m_1, m_2)_{[i \rightarrow m']}$ for the move (n_1, n_2) with $n_i = m'$ and $n_{3-i} = m_{3-i}$. This notation is extended to strategies in a natural way.

In the context of non-zero-sum games, several notions of equilibria have been defined. We present a refinement of *Nash equilibria* towards non-deterministic concurrent games.

Definition 5. Let \mathcal{G} be a concurrent game, and s be a state of \mathcal{G} . A pseudo Nash equilibrium in \mathcal{G} from s is a tuple $((\sigma_1, \sigma_2), \pi)$ where $(\sigma_1, \sigma_2) \in \text{Prof}_{\mathcal{G}}$, and $\pi \in \text{Out}_{(\mathcal{G}, s)}(\sigma_1, \sigma_2)$ is such that for all $i \in \{1, 2\}$ and all $\sigma'_i \in \text{Strat}_{\mathcal{G}}^i$, it holds:

$$\forall \pi' \in \text{Out}_{(\mathcal{G}, s)}((\sigma_1, \sigma_2)_{[i \rightarrow \sigma'_i]}). \pi' \preceq_i \pi.$$

Such an outcome π is called a *best play* for the strategy profile (σ_1, σ_2) .

In the case of deterministic games, π is uniquely determined by (σ_1, σ_2) , and pseudo Nash equilibria coincide with *Nash equilibria* as defined in [14]: they are strategy profiles where no player has an incentive to unilaterally deviate from her strategy.

In the case of non-deterministic games, a strategy profile for an equilibrium may give rise to several outcomes. The choice of playing the best play π is then made cooperatively by both players: once both strategies are fixed, it is the interest of both players to cooperate and play “optimally”.

2.3 Back to timed games

Two comments are in order here: (i) non-determinism in timed games could be dropped by giving priority to one of the players, in case both of them play the same delay. Our

algorithm could of course be adapted in this case; (ii) even if the timed game were deterministic, our transformation to region games involves some extra non-determinism. As will be seen in the sequel, the above notion of *pseudo Nash equilibria* is the notion we need for our construction to preserve equilibria.

It is easy to see the semantics of a timed game as the semantics of an infinite-state concurrent game (see the research report [5]). Using that point-of-view, timed games inherit the notions of history, play, path, strategy, profile, outcome and pseudo Nash equilibrium. We illustrate some of these notions on the running example.

Example 1 (Cont'd). This game starts in configuration $(\ell_0, 0)$ (clock x is set to 0). A strategy profile is then determined by an initial choice for the first transition. If one of the players choose some delay smaller than 1, she will have payoff 1 but the other player will have payoff 0, hence the other player will be able to preempt this choice and choose a smaller delay that will improve her own payoff. Hence there will be no such pseudo Nash equilibrium. There is a single pseudo Nash equilibrium, where player 1 chooses $(1, c)$ (delay for 1 t.u. and take transition c) and player 2 chooses $(1, b)$. The best play for that strategy profile is the run taking transition c .

In this paper we will be interested in the computation of pseudo Nash equilibria in timed games. To do so we propose a sequence of transformations that will preserve equilibria (in some sense), yielding the construction of two turn-based finite games in which the initial problem will be reduced to the computation of *twin* Nash equilibria. All these transformations are presented in the next section. These transformations will also give a new point-of-view on timed games, which we will use in Section 4.2 to recover some decidability results. Many more results are expected.

3 From timed games to turn-based finite games

In this section we propose a chain of transformations of the timed game \mathcal{G} into two turn-based finite games, and reduce the computation of pseudo Nash equilibria in \mathcal{G} to the computation of ‘twin’ Nash equilibria in the two turn-based games. Notice that we will have to impose restrictions on the preference relations: indeed, price-optimal reachability is undecidable in two-player priced timed games, and these quantitative objectives can be encoded as a payoff function, see [7] for details.

3.1 From timed games to concurrent games...

We assume the reader is familiar with the region automaton abstraction for timed automata [1]. Let $\mathcal{G} = \langle \text{Loc}, \text{Cl}, \text{Inv}, \text{Trans}, \text{Owner}, (\preceq_1, \preceq_2) \rangle$ be a timed game. Let \mathfrak{R} be the set of regions for the timed automaton underlying \mathcal{G} , and $\pi_{\mathfrak{R}}$ be the projection over the regions \mathfrak{R} (for configurations, runs, etc.) We define the *region game* $\mathcal{R} = \langle \text{States}, \text{Edg}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_1^{\mathcal{R}}, \preceq_2^{\mathcal{R}}) \rangle$ as follows:

- States = $\{(\ell, r) \in \text{Loc} \times \mathfrak{R} \mid r \models \text{Inv}(\ell)\}$;
- Edg = $\{((\ell, r), (\ell', r')) \mid (\ell, r) \rightarrow (\ell', r') \text{ in the region automaton of } \mathcal{G}\}$;

- $\text{Act} = \{\perp\} \cup \{(r, \delta) \mid r \in \mathfrak{R} \text{ and } \delta \in \text{Trans}\}$;
- $\text{Mov}: \text{States} \times \{1, 2\} \rightarrow 2^{\text{Act}} \setminus \{\emptyset\}$ such that:

$$\text{Mov}((\ell, r), i) = \{(r', \delta) \mid r' \in \text{Succ}(r), r' \models \text{Inv}(\ell), \delta = (\ell, g, Y, \ell') \text{ is s.t.} \\ r' \models g \text{ and } [Y \leftarrow 0]r' \models \text{Inv}(\ell') \text{ and } \text{Owner}(\delta) = i\}$$

if this set is non-empty, and $\text{Mov}((\ell, r), i) = \{\perp\}$ otherwise.

- $\text{Tab}: \text{States} \times \text{Act}^2 \rightarrow 2^{\text{Edg}} \setminus \{\emptyset\}$ such that for every $(\ell, r) \in \text{States}$ and every $(m_1, m_2) \in \text{Mov}((\ell, r), 1) \times \text{Mov}((\ell, r), 2)$, if we write r' for $\min\{r_j \mid j \in \{1, 2\}\}$ and $m_j = (r_j, \delta_j)$,⁴ then we have:

$$\text{Tab}((\ell, r), (m_1, m_2)) = \{((\ell, r), (\ell_j, [Y_j \leftarrow 0]r_j)) \mid j \in \{1, 2\} \text{ and} \\ m_j = (r_j, \delta_j) \text{ with } r_j = r', (\ell, g_j, Y_j, \ell_i) = \delta_j \text{ and } r_j \models g_j\}$$

- The preference relation $\preceq_i^{\mathcal{R}}$ for player i is defined by saying that $\gamma \preceq_i^{\mathcal{R}} \gamma'$ iff there exists ρ and ρ' such that $\pi_{\mathfrak{R}}(\rho) = \gamma$, $\pi_{\mathfrak{R}}(\rho') = \gamma'$ and $\rho \preceq_i \rho'$.

Note that the game \mathcal{R} is non-deterministic, even if the original timed game is not. Indeed, non-determinism appears when players want to play delays leading to the same region. The (relative) order of the choices for the delays chosen by the two players cannot be distinguished by the region abstraction.

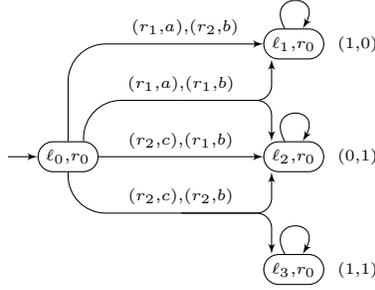
Definition 6. A preference relation \preceq_i is said to be region-uniform when for all plays ρ and ρ' , if the sequence of regions seen in both paths are the same, then they are equivalent, i.e. $\rho \preceq_i \rho'$ and $\rho' \preceq_i \rho$.

Proposition 7. Let \mathcal{G} be a timed game, and assume that the two preference relations of \mathcal{G} are region-uniform. Let \mathcal{R} be its associated region game. Then there is a pseudo Nash equilibrium in \mathcal{G} from $(\ell_0, \mathbf{0})$ with best play ρ iff there is a pseudo Nash equilibrium in \mathcal{R} from $(\ell_0, [\mathbf{0}]_{\mathfrak{R}})$ with best play $\pi_{\mathfrak{R}}(\rho)$. Furthermore, this equivalence is constructive.

Example 1 (Cont'd). We illustrate the construction and the previous notions on the running example. We write r_0 (resp. r_1, r_2) for the region $x = 0$ (resp. $0 < x < 1$, $x = 1$). The region game \mathcal{R} is as depicted on Fig. 1. In this region game, there are two non-deterministic transitions. First when the two players choose to wait until region r_2 , in which case the game can turn to either ℓ_2 or to ℓ_3 . Then when both players choose to move within the region r_1 (there is an uncertainty on whether player 1 or player 2 was faster), and depending on who was faster, the game will move to either ℓ_1 or ℓ_2 . The first non-determinism is inherent to the game (and could be removed by construction assuming one player is more powerful, see Subsection 2.3 for explanations), whereas the second non-determinism is (somehow) artificial and comes from the region abstraction.

In \mathcal{G} , there is a single pseudo Nash equilibrium, where both players wait until $x = 1$ (region r_2), and propose to move respectively to ℓ_3 (resp. ℓ_2). The best play is then $(\ell_0, 0)(\ell_3, 0)^*$. This corresponds to the unique pseudo Nash equilibrium that we find in the region game.

⁴ This is well-defined because both r_j 's are time-successors of r .



The transition table from (ℓ_0, r_0)
(i.e., $\text{Tab}((\ell_0, r_0), (m_1, m_2))$):

	$m_2 = (r_1, b)$	$m_2 = (r_2, b)$
$m_1 = (r_1, a)$	$(\ell_1, r_0), (\ell_2, r_0)$	(ℓ_1, r_0)
$m_1 = (r_2, c)$	(ℓ_2, r_0)	$(\ell_2, r_0), (\ell_3, r_0)$

Fig. 1. The region game from our original automaton.

3.2 ... next to two twin concurrent games...

Given a concurrent non-deterministic finite game $\mathcal{R} = \langle \text{States}, \text{Edg}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_1^{\mathcal{R}}, \preceq_2^{\mathcal{R}}) \rangle$, we construct two concurrent games \mathcal{R}_1 and \mathcal{R}_2 where we simply forget the preferences of one player. Formally for $i \in \{1, 2\}$, we define the game $\mathcal{R}_i = \langle \text{States}, \text{Edg}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_1^i, \preceq_2^i) \rangle$, where \preceq_i^i is the quasi-order \preceq_i , and \preceq_{3-i}^i is the trivial quasi-order where all runs are equivalent.

Definition 8. A twin pseudo Nash equilibrium for the two games \mathcal{R}_1 and \mathcal{R}_2 is a tuple $((\sigma_1^{\mathcal{R}_1}, \sigma_2^{\mathcal{R}_1}), (\sigma_1^{\mathcal{R}_2}, \sigma_2^{\mathcal{R}_2}), \rho)$ such that $((\sigma_1^{\mathcal{R}_1}, \sigma_2^{\mathcal{R}_1}), \rho)$ is a pseudo Nash equilibrium in the game \mathcal{R}_1 and $((\sigma_1^{\mathcal{R}_2}, \sigma_2^{\mathcal{R}_2}), \rho)$ is a pseudo Nash equilibrium in the game \mathcal{R}_2 . We furthermore say that ρ is a best play for the twin pseudo equilibrium.

We relate pseudo Nash equilibria in \mathcal{R} with twin pseudo Nash equilibria in \mathcal{R}_1 and \mathcal{R}_2 . Note that we require best plays be the same, but not strategies.

Proposition 9. Let \mathcal{R} be the region game associated with some timed game \mathcal{G} . Then there is a pseudo Nash equilibrium in \mathcal{R} from s with best play γ if and only if there is a twin pseudo equilibrium for the corresponding games \mathcal{R}_1 and \mathcal{R}_2 from s with best play γ . Furthermore this equivalence is constructive.

3.3 ... next to concurrent deterministic games...

We transform each game \mathcal{R}_i into a concurrent deterministic game \mathcal{C}_i . Game \mathcal{C}_i will give priority to player i , in that it will be the role of player i to solve non-determinism. The game $\mathcal{C}_i = \langle \text{States}, \text{Edg}, \text{Act}', \text{Mov}_i, \text{Tab}_i, (\preceq_1^i, \preceq_2^i) \rangle$ is defined as follows:

- $\text{Act}' = \text{Act} \cup ((\text{Act} \setminus \{\perp\}) \times \{\bullet, \circ\})$;
- $\text{Mov}_i: \text{States} \times \{1, 2\} \rightarrow 2^{\text{Act}'} \setminus \{\emptyset\}$ such that:

$$\text{Mov}_i(s, i) = \begin{cases} \{\perp\} & \text{if } \text{Mov}(s, i) = \{\perp\} \\ \text{Mov}(s, i) \times \{\bullet, \circ\} & \text{otherwise} \end{cases}$$

$$\text{Mov}_i(s, 3-i) = \text{Mov}(s, 3-i)$$

- Given $(m_1, m_2) \in \text{Mov}(s, 1) \times \text{Mov}(s, 2)$ we have that $\text{Tab}(s, (m_1, m_2))$ has at least one element, and at most two elements.⁵
 - In case it has only one element, then setting $m'_{3-i} = m_{3-i}$ and picking $m'_i \in \{(m_i, \bullet), (m_i, \circ)\}$, we define: $\text{Tab}_i(s, (m'_1, m'_2)) = \text{Tab}(s, (m_1, m_2))$;
 - In case it has two elements, say (s, s_\bullet) and (s, s_\circ) , one of them comes from a transition of player i in \mathcal{G} and the other comes from a transition of player $3-i$ in \mathcal{G} . Hence w.l.o.g. we can assume that (s, s_\bullet) belongs to player i . We now define $m'_{3-i} = m_{3-i}$ and for any $m'_i \in \{(m_i, \bullet), (m_i, \circ)\}$, we define:

$$\text{Tab}_i(s, (m'_1, m'_2)) = \begin{cases} \{(s, s_\bullet)\} & \text{if } m'_i = (m_i, \bullet) \\ \{(s, s_\circ)\} & \text{if } m'_i = (m_i, \circ) \end{cases}$$

By construction, the two games \mathcal{C}_1 and \mathcal{C}_2 are deterministic, and they share the same structure. Only decisions on how to solve non-determinism are made by different players. Our aim will be to compute equilibria in these two similar games.

Proposition 10. *Assume \mathcal{C}_i (with $i \in \{1, 2\}$) is the deterministic concurrent game defined from the concurrent game \mathcal{R}_i . Then there is a pseudo Nash equilibrium in \mathcal{R}_i from s with best play γ iff there is a Nash equilibrium in \mathcal{C}_i from s with best play γ .⁶ Furthermore this equivalence is constructive.*

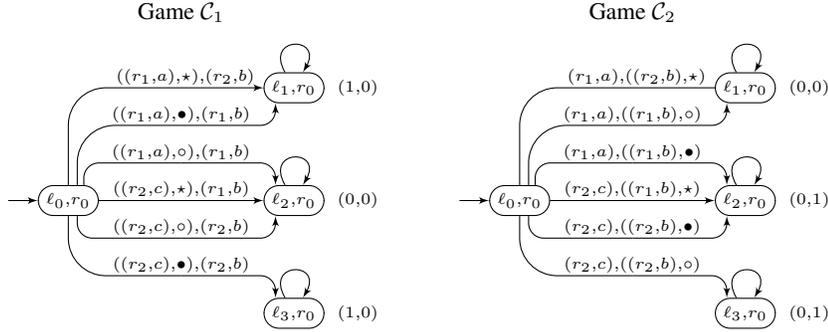


Fig. 2. Two concurrent games \mathcal{C}_1 and \mathcal{C}_2 from our original automaton.

Example 1 (Cont'd). We build on the previous example, and give the two games \mathcal{C}_1 and \mathcal{C}_2 in Fig. 2. An action (m, \star) denotes either (m, \bullet) or (m, \circ) . There are several Nash equilibria in game \mathcal{C}_1 : one where the first player chooses $((r_2, c), \bullet)$ and the second player chooses (r_2, b) , which leads to (ℓ_3, r_0) with payoff $(1, 0)$; and one where both players play a pair of actions leading to (ℓ_1, r_0) , in which case the payoff is also $(1, 0)$.

⁵ This is because the game \mathcal{G} is non-blocking, and in this game, each player proposes her choice for a transition, and one of these two transitions will be chosen.

⁶ Remember that \mathcal{C}_i 's and \mathcal{R}_i 's share the same structure and have the same runs.

Similarly there are several Nash equilibria in game \mathcal{C}_2 : one where the second player chooses $((r_2, b), \circ)$ and the first player chooses (r_2, c) , which leads to (ℓ_3, r_0) with payoff $(0, 1)$; the second one where both players play a pair of actions leading to (ℓ_2, r_0) , in which case the payoff is also $(0, 1)$.

There is a single twin equilibrium in \mathcal{C}_1 and \mathcal{C}_2 , namely the one leading to state (ℓ_3, r_0) , which coincides with those equilibria already found in \mathcal{G} and \mathcal{R} .

3.4 ... and finally to two turn-based games

In the (deterministic) concurrent game \mathcal{C}_i , the advantage is given to player i , who has the ability to solve non-determinism. We can give a slightly different interpretation to that mechanism, which takes into account an interpretation of the new actions. Indeed, actions have a timed interpretation in the original timed game, and can be ordered w.r.t. their delay. Taking advantage of this order on actions, we build a turn-based game \mathcal{T}_i .

Let $\mathcal{C}_i = \langle \text{States}, s_0, \text{Edg}, \text{Act}', \text{Mov}_i, \text{Tab}_i, (\preceq_1^i, \preceq_2^i) \rangle$ be the games obtained from the previous construction. Let $s \in \text{States}$. We naturally order the set $\text{Mov}_i(s, 1) \cup \text{Mov}_i(s, 2)$ with a relation $<_s$ so that:

- (i) if $\perp \in \text{Mov}_i(s, 1) \cup \text{Mov}_i(s, 2)$ then \perp is maximal w.r.t. $<_s$;
- (ii) for every $m \in \text{Mov}_i(s, j)$, there exists $s' \in \text{States}$ such that for every $m' \in \text{Mov}_i(s, 3-j)$, $m <_s m'$ implies $\text{Tab}_i(s, (m, m')) = \{(s, s')\}$.

This is possible due to the definition of game \mathcal{C}_i : when (r, δ_{3-i}) is allowed to player $3-i$ from s , and $((r, \delta_i), \bullet)$ and $((r, \delta_i), \circ)$ are allowed to player i from s , then the three actions are totally ordered by $<_s$ as follows:⁷ $((r, \delta_i), \bullet) <_s (r, \delta_{3-i}) <_s ((r, \delta_i), \circ)$. Intuitively an action with marker \bullet means that player i can play her own transition faster than player $3-i$ can play her own transition, but also that she can decide to play more slowly (role of action with marker \circ).

We can also define an equivalence relation $=_s$ compatible with this order, by saying $m =_s m' \Leftrightarrow m, m' \in \text{Mov}_i(s, 1) \cup \text{Mov}_i(s, 2)$, $m \not<_s m'$ and $m' \not<_s m$. It is worth noticing that $m =_s m'$ implies that they belong to the same player. This can be the case if two transitions are available to a player from the same region, and also if a player can only play action \perp . We will write $[m]_s$ for the equivalence class associated to m . We next say that $[m]_s$ belongs to player j whenever all actions in $[m]_s$ belong to player j .

Example 1 (Cont'd). Consider games \mathcal{C}_1 and \mathcal{C}_2 depicted in Fig. 2. In game \mathcal{C}_1 , the order on actions (written simply as $<$) from (ℓ_0, r_0) is given by:

$$\begin{array}{cccccc} ((r_1, a), \bullet) & < & (r_1, b) & < & ((r_1, a), \circ) & < & ((r_2, c), \bullet) & < & (r_2, b) & < & ((r_2, c), \circ) \\ \zeta & & \zeta \\ (\ell_1, r_0) & & (\ell_2, r_0) & & (\ell_1, r_0) & & (\ell_3, r_0) & & (\ell_2, r_0) \end{array}$$

Below each action we write the target state when this action is played, provided an action smaller (for the order $<$) is not played by the other player. There is no target with action $((r_2, c), \circ)$ because it is always preempted by some ‘faster’ action (no \perp action is available in our example).

In game \mathcal{C}_2 , the order on actions (also written $<$) from (ℓ_0, r_0) is given by:

⁷ This is due to the fact that we have assumed edge (s, s_\bullet) belong to player i , see the construction of game \mathcal{C}_i .

$$\begin{array}{ccccccccc}
((r_1, b), \bullet) & < & (r_1, a) & < & ((r_1, b), \circ) & < & ((r_2, b), \bullet) & < & (r_2, c) & < & ((r_2, b), \circ) \\
\downarrow & & \downarrow \\
(\ell_2, r_0) & & (\ell_1, r_0) & & (\ell_2, r_0) & & (\ell_2, r_0) & & (\ell_3, r_0)
\end{array}$$

We will take advantage of this order on actions to build turn-based games that will in some sense be equivalent with the previous concurrent (deterministic) games. The idea will be to take the smallest action(s) in the order, and ask the corresponding player whether or not she wants to play that action; if yes, we proceed with this action in the game, otherwise we do the same with the second action in the order until one of the players plays her action; The meaning in the context of timed games is actually also the following: we see that if the two players want to play in the same region, then in game \mathcal{C}_i the advantage of player i is that we first ask her whether she wishes to play her action (role of action labelled with \bullet), then if not, the other player will be asked to decide whether she wants to play her own action, and finally, if not, we ask a last time player i whether she wants to play her action (now she has the additional knowledge that the other player didn't choose her own action).

Formally we define the turn-based game \mathcal{T}_i as follows: $\mathcal{T}_i = \langle \text{States}_i, \text{Edg}_i, \text{Act}' \cup \{\text{del}\}, \text{Mov}'_i, \text{Tab}'_i, (\preceq^i_1, \preceq^i_2) \rangle$ where:

- $\text{States}_i = \{(s, [m]_s) \mid s \in \text{States and } m \in (\text{Mov}_i(s, 1) \cup \text{Mov}_i(s, 2)) \setminus \{\perp\}\}$;
- The set Edg_i is defined as follows:

$$\begin{aligned}
\text{Edg}_i = & \{((s, [m]_s), (s, [m']_s)) \mid m' \neq \perp \text{ is next after } m \text{ w.r.t. } <_s\} \\
& \cup \{((s, [m]_s), (s', [m']_{s'})) \mid \{(s, s')\} = \text{Tab}_i(s, (m, m'')) \\
& \text{for every } m <_s m'' \text{ and } m' \text{ is minimal w.r.t. } <_{s'} \text{ from } s'\};
\end{aligned}$$

- The set of available actions is defined as follows:
 - if $[m]_s$ belongs to player j , then we use the new action del (for *delay*):
$$\text{Mov}'_i((s, [m]_s), j) = \begin{cases} \text{Mov}_i(s, j) \cap [m]_s & \text{if } m \text{ is maximal w.r.t.} \\ <_s \text{ in } \text{Mov}_i(s, j) \\ (\text{Mov}_i(s, j) \cap [m]_s) \cup \{\text{del}\} & \text{otherwise} \end{cases}$$
 - if $[m]_s$ belongs to player $3 - j$, then $\text{Mov}'_i((s, [m]_s), j) = \{\perp\}$.
- The transition table is defined as follows:
 - if $[m]_s$ belongs to player 1:

$$\left\{ \begin{array}{l} \text{Tab}'_i((s, [m]_s), (m, \perp)) = \{((s, [m]_s), (s', [m']_s)) \mid \{(s, s')\} = \text{Tab}_i(s, (m, m'')) \\ \text{for every } m <_s m'' \text{ and } m' \text{ is minimal w.r.t. } <_{s'} \text{ from } s'\} \\ \text{Tab}'_i((s, [m]_s), (\text{del}, \perp)) = \{((s, [m]_s), (s, [m']_s)) \mid m' \text{ is next after } m \text{ w.r.t. } <_s\} \end{array} \right.$$

- the second case ($[m]_s$ belongs to player 2) is similar, just swap m or del with \perp .
- In order to define the preference relations we first define a projection from plays in the turn-based game \mathcal{T}_i onto plays in the concurrent game \mathcal{C}_i . Pick a run ν in \mathcal{T}_i , and define its projection $\psi_i(\nu)$ in \mathcal{C}_i as follows: if

$$\nu = (s_1, \mathbf{m}_1^1)(s_1, \mathbf{m}_1^2) \dots (s_1, \mathbf{m}_1^{k_1})(s_2, \mathbf{m}_2^1) \dots (s_2, \mathbf{m}_2^{k_2}) \dots (s_p, \mathbf{m}_p^1) \dots (s_p, \mathbf{m}_p^{k_p}) \dots$$

with \mathbf{m}_i^1 minimal w.r.t. $<_{s_i}$ for every $1 \leq i$, then $\psi(\nu) = s_1 s_2 \dots s_p \dots$. The preference relations are then defined according to this projection:

$$\nu \preceq_j^i \nu' \Leftrightarrow \psi_i(\nu) \preceq_j^i \psi_i(\nu')$$

Note that the game \mathcal{T}_i is turn-based⁸ and that a state $(s, [m]_s)$ belongs to player j such that $m \in \text{Mov}_i(s, j)$ (as already mentioned this is independent of the choice of m in $[m]_s$). The structure of the turn-based games \mathcal{T}_1 and \mathcal{T}_2 are now slightly different from that of the previous concurrent deterministic games \mathcal{C}_1 and \mathcal{C}_2 .

Proposition 11. *Let \mathcal{C}_i (with $i \in \{1, 2\}$) be the previous deterministic concurrent game, and let \mathcal{T}_i be the associated turn-based game. There is a Nash equilibrium in \mathcal{C}_i from s with best play $\psi_i(\nu)$ iff there is a Nash equilibrium in \mathcal{T}_i from $(s, [m]_s)$ with best play ν , where m is a minimal action w.r.t. $<_s$. Furthermore this equivalence is constructive.*

Example 1 (Cont'd). We build on our running example, and compute the corresponding games \mathcal{T}_1 and \mathcal{T}_2 . They are displayed on Fig. 3. Plain states and plain edges belong to

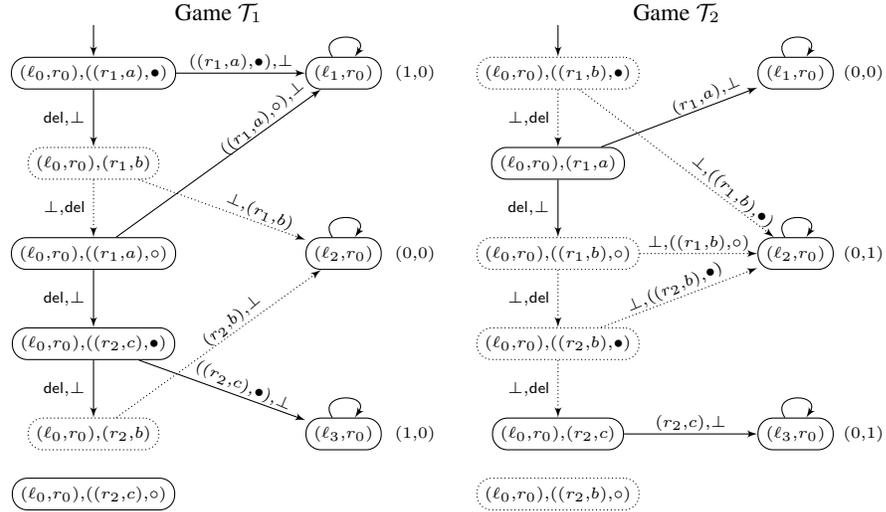


Fig. 3. Final turn-based games from our original timed game

player 1 whereas dotted states and dotted edges belong to player 2. We do recognize here the various Nash equilibria that we described in the concurrent deterministic games, and only one is “common” to both games, namely the one leading to (ℓ_3, r_0) .

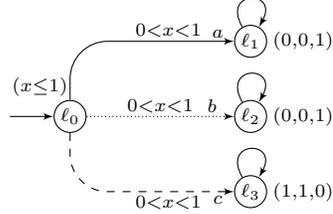
3.5 Summary of the construction

The following theorem summarizes our construction:

⁸ By construction, in any state $(s, [m]_s)$, one of $\text{Mov}'_i((s, [m]_s), j)$ with $j \in \{1, 2\}$ equals $\{\perp\}$.

Theorem 12. *Let \mathcal{G} be a timed game with region-uniform preference relations. Assume \mathcal{T}_1 and \mathcal{T}_2 are the two turn-based (deterministic) games constructed in this section. Then, there is a pseudo Nash equilibrium in \mathcal{G} from $(\ell_0, \mathbf{0})$ with best play ρ iff there are two Nash equilibria in \mathcal{T}_1 and \mathcal{T}_2 from $((\ell_0, \mathbf{0}), [m]_{(\ell_0, \mathbf{0})})$ with best plays ν_1 and ν_2 respectively, where m is a minimal action w.r.t. $<_{(\ell_0, \mathbf{0})}$, such that $\psi_1(\nu_1) = \psi_2(\nu_2) = \pi_{\mathfrak{R}}(\rho)$. Furthermore this equivalence is constructive.*

Remark 1. The three-player game on the right has several Nash equilibria, for instance player 1 (plain arrows) chooses her transition at time 0.6, player 2 (dotted arrows) chooses her transition at time 0.7, and player 3 (dashed arrows) chooses her transition at time 0.8. If we build the region abstraction, each player will have a single possible move (play her transition in the region $0 < x < 1$), and the game will proceed by selecting non-deterministically one of them. There would be several ways to extend the method developed in this paper to three players: have a copy of the game for each player, assuming she plays against a coalition of the other players, or have a copy of the game for each priority order given to the players. It is not hard to be convinced that none of these choices will be correct on this example.



4 Decidability results

4.1 Some general decidability results

We first need a representation for the preference relations (which must be region-uniform) of both players. Let $\mathcal{G} = \langle \text{Loc}, \text{Cl}, \text{Inv}, \text{Trans}, \text{Owner}, (\preceq_1, \preceq_2) \rangle$ be a two-player timed game. We assume the preference relation for player i is given by a (possibly infinite) sequence of linear-time objectives $(\Omega_j^i)_{j \geq 1}$ where it is better for a run to satisfy Ω_j^i than Ω_k^i as soon as $k > j$ (w.l.o.g. we assume that Ω_{j+1}^i implies $\neg \Omega_j^i$ for all $l \leq j$). In other terms, the aim of player i is to minimize the index j for which the play belongs to Ω_j^i . These objectives include ω -regular or LTL-definable objectives, and also more *quantitative* objectives (for instance, given a distinguished goal state $\text{Goal}_i \in \text{LOC}$ for player i , by defining Ω_j^i to be the set of traces visiting Goal_i in less than j steps).

We first need to (be able to) transfer objectives (and preference relations) to the two turn-based games \mathcal{T}_1 and \mathcal{T}_2 : a linear-time objective Ω in \mathcal{G} is said to be *transferable to game \mathcal{T}_i* whenever we can construct an objective $\widehat{\Omega}$ such that for every run ν in \mathcal{T}_i , $\nu \models \widehat{\Omega}$ iff for all ρ with $\pi_{\mathfrak{R}}(\rho) = \psi_i(\nu)$, $\rho \models \Omega$. It is said *transferable* whenever it is transferable to both \mathcal{T}_1 and \mathcal{T}_2 . For example, notice that (sequences of) stutter-free region-uniform objectives are transferable.

Nash equilibria in game \mathcal{T}_i will be rather easy to characterize since player $3 - i$ will never be inclined to deviate from her strategy (all runs are equivalent for her preference relation). We assume all objectives Ω_j^i are transferable, and we write $W_i^{3-i}(j)$ for the set of winning states in game \mathcal{T}_i for player $3 - i$ with the objective $\bigwedge_{1 \leq k < j} (\neg \widehat{\Omega}_k^i)$. Those sets are computable for many classes of objectives. Then:

Theorem 13. *Let \mathcal{G} be a timed game with preference relations given as transferable, region-uniform, prefix-independent sequences $(\Omega_j^i)_j$ of objectives. There is a pseudo Nash equilibrium in \mathcal{G} with payoff (Ω_1^j, Ω_2^k) iff there are two runs ν_1 in \mathcal{T}_1 and ν_2 in \mathcal{T}_2 s.t.⁹ (i) $\nu_1 \models (\mathbf{G} W_1^2(j)) \wedge \widehat{\Omega}_j^1$, (ii) $\nu_2 \models (\mathbf{G} W_2^1(k)) \wedge \widehat{\Omega}_k^2$, and (iii) $\psi_1(\nu_1) = \psi_2(\nu_2)$.*

Notice that this allows to handle ω -regular (and LTL-definable) objectives by considering the product of the game with a suitable (deterministic) automaton.

The sequence of states $(W_i^{3-i}(j))_{j \geq 1}$ in game \mathcal{T}_i is non-increasing and hence stationary (because \mathcal{T}_i is finite-state). Hence there exist indices $h_0 = 1 < h_1 < h_2 < \dots < h_l$ such that the function $j \mapsto W_i^{3-i}(j)$ is constant on all intervals $[h_p, h_{p+1})$ and on $[h_l, +\infty)$. Those indices can be computed together with the corresponding sets of winning states. Then the only possible equilibria in \mathcal{T}_i are those such that there is a run satisfying $\widehat{\Omega}_j^i$ that stays furthermore within the set $W_i^{3-i}(h_p)$ if $h_p \leq j < h_{p+1}$, or within $W_i^{3-i}(h_l)$ if $j \geq h_l$. This can be done for instance if each player is given a goal state Goal_i , and Ω_j^i is “reach Goal_i in j steps”. In that case, $W_i^{3-i}(h_l)$ is the set of states from which player $3 - i$ can avoid Goal_i . Hence we can compute Nash equilibria in two-player timed games where each player tries to minimize the number of steps to the goal state. This allows to recover part of the results of [8] for two-player games.

4.2 Zero-sum games

Our chain of transformations also yields a new point-of-view on classical two-player timed games with zero-sum objectives. In that case the preference relation of player 1 is characterized by the sequence $(\Omega, \neg\Omega)$ whereas that of player 2 is characterized by the sequence $(\neg\Omega, \Omega)$. In that case we say that the objective of player 1 is Ω .

Theorem 14. *Let \mathcal{G} be a zero-sum timed game where player 1’s objective is Ω , and is assumed to be transferable. Then player 1 has a winning strategy in \mathcal{G} from $(\ell, \mathbf{0})$ iff player 1 has a winning strategy in game \mathcal{T}_2 from $(\ell, [\mathbf{0}])$ for the objective $\widehat{\Omega}$.*

5 Conclusion

We have proposed a series of transformations of two-player timed games into two turn-based finite games. These transformations reduce the computation of Nash equilibria in timed games for a large class of objectives (the so-called region-uniform objectives) to the computation of “twin” equilibria for related objectives in the two turn-based finite games. We give an example on how this can be used to compute Nash equilibria in timed games. In turn our transformations give a nice and new point-of-view on zero-sum timed games, which can then be interpreted as a turn-based finite game.

Our method does not extend to n players. In [6], we have developed a completely new approach that allows to compute Nash equilibria in timed games with an arbitrary number of players but only for reachability objectives. We plan to continue working on the computation of Nash equilibria in timed games with an arbitrary number of players.

⁹ \mathbf{G} is the LTL modality for “always”.

Another interesting research direction is the computation of other kinds of equilibria in timed games (secure equilibria, subgame-perfect equilibria, *etc*). We believe that the transformations that we have made in this paper are correct also for these other notions, and that we can for instance reduce the computation of subgame-perfect equilibria to the computation of subgame-perfect equilibria in the two turn-based finite games. A major difference is that Theorem 13 has to be refined. Tree automata could be the adequate tool for this problem [17].

References

1. R. Alur and D. L. Dill. A theory of timed automata. *TCS*, 126(2):183–235, 1994.
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49:672–713, 2002.
3. E. Asarin, O. Maler, A. Pnueli, and J. Sifakis. Controller synthesis for timed automata. In *SSSC’98*, p. 469–474. Elsevier, 1998.
4. G. Behrmann, A. Cougnard, A. David, E. Fleury, K. G. Larsen, and D. Lime. UPPAAL-Tiga: Time for playing games! In *CAV’07*, LNCS 4590, p. 121–125. Springer, 2007.
5. P. Bouyer, R. Brenguier, and N. Markey. Computing equilibria in two-player timed games *via* turn-based finite games. Research Report LSV-10-11, Lab. Spécification & Vérification, ENS Cachan, France, 2010.
6. P. Bouyer, R. Brenguier, and N. Markey. Nash equilibria for reachability objectives in multi-player timed games. In *CONCUR’10*, LNCS. Springer, 2010. To appear.
7. R. Brenguier. Calcul des équilibres de Nash dans les jeux temporisés. Master’s thesis, ENS Cachan, 2009.
8. Th. Brihaye, V. Bruyère, and J. De Pril. Equilibria in quantitative reachability games. In *CSR’10*, LNCS. Springer, 2010. To appear.
9. K. Chatterjee, T. A. Henzinger, and M. Jurdziński. Games with secure equilibria. In *LICS’06*, p. 160–169. IEEE Comp. Soc. Press, 2006.
10. K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In *CSL’04*, LNCS 3210, p. 26–40. Springer, 2004.
11. 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, p. 142–156. Springer, 2003.
12. M. Félegyházi, J.-P. Hubaux, and L. Buttyán. Nash equilibria of packet forwarding strategies in wireless ad hoc networks. *IEEE Trans. Mobile Computing*, 5(5):463–476, 2006.
13. T. A. Henzinger. Games in system design and verification. In *TARK’05*, p. 1–4. Nat. Univ. Singapore, 2005.
14. J. F. Nash. Equilibrium points in n -person games. *Proc. Nat. Academy of Sciences of the USA*, 36(1):48–49, 1950.
15. S. Paul and S. Simon. Nash equilibrium in generalised Muller games. In *FSTTCS’09*, LIPIcs 4, p. 335–346. Leibniz-Zentrum für Informatik, 2009.
16. W. Thomas. Infinite games and verification. In *CAV’02*, LNCS 2404, p. 58–64. Springer, 2002. Invited Tutorial.
17. M. Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *FSTTCS’06*, LNCS 4337, p. 212–223. Springer, 2006.
18. M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *FoSSaCS’08*, LNCS 4962, p. 20–34. Springer, 2008.
19. M. Ummels and D. Wojtczak. The complexity of Nash equilibria in simple stochastic multiplayer games. In *ICALP’09*, LNCS 5556, p. 297–308. Springer, 2009.
20. M. Ummels and D. Wojtczak. Decision problems for Nash equilibria in stochastic games. In *CSL’09*, LNCS 5771, p. 515–530. Springer, 2009.