Nash Equilibria in Concurrent Games with Büchi Objectives

Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels

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

Abstract

We study the problem of computing pure-strategy Nash equilibria in multiplayer concurrent games with Büchi-definable objectives. First, when the objectives are Büchi conditions on the game, we prove that the existence problem can be solved in polynomial time. In a second part, we extend our technique to objectives defined by deterministic Büchi automata, and prove that the problem then becomes EXPTIME-complete. We prove PSPACE-completeness for the case where the Büchi automata are 1-weak.

1998 ACM Subject Classification F.1.1; F.1.2; F.2.2

Keywords and phrases Concurrent games, Nash equilibria, Büchi objectives

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2011.375

1 Introduction

Game theory (especially games played on graphs) is used in computer science as a powerful framework for modelling interactions in embedded systems [18, 13]. Until recently, more focus had been put on purely antagonistic situations, where the system should fulfil its specification however the environment behaves. This situation can be modelled as a two-player game (one player for the system, and one for the environment), and a winning strategy for the first player is a good controller for the system. In this purely antagonistic view, the objectives of both players are opposite: the aim of the second player is to prevent the first player from achieving her own objective; such games are called zero-sum.

In many cases, however, games are non-zero-sum, especially when they involve more than two players. Such games appear e.g. in various problems in telecommunications, where several agents try to send data on a network [10]. Focusing only on winning strategies in this setting may then be too narrow: winning strategies must be winning against any behaviour of the other agents, and do not consider the fact that the other agents also have their own objectives. In the non-zero-sum setting, each player can have a different payoff associated with an outcome of the game; it is then more interesting to look for *equilibria*. For instance, a Nash equilibrium is a behaviour of the agents in which they play rationally, in the sense that no agent can get a better payoff by unilaterally switching to another strategy [15]. This corresponds to stable states of the game. Note that Nash equilibria need not exist and are not necessarily *optimal*: several equilibria can coexist, possibly with different payoffs.

Our contribution. We focus here on qualitative objectives for the players: such objectives are ω-regular properties over infinite plays, and a player receives payoff 1 if the property is fulfilled and 0 otherwise. Our aim is to decide the existence of pure-strategy Nash equilibria in nondeterministic concurrent games. Being concurrent (instead of the more classical *turn-based*

31st Int'l Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011).
Editors: Supratik Chakraborty, Amit Kumar; pp. 375–386
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

games) and nondeterministic are two important properties of *timed games* (which are games played on timed automata [2, 9]), to which we ultimately want to apply our algorithms.

In a first part, we focus on internal Büchi conditions (defined on the game directly) and show that we can decide the existence of equilibria in polynomial time, which has to be compared with the NP-completeness of the problem in the case of reachability objectives [8, 3]. This relies on an iterated version of a *repellor operator* [3]. Roughly speaking, the repellor is to the computation of Nash equilibria in non-zero-sum games what the attractor is to the computation of winning states in zero-sum games. The repellor operator we use for Büchi objectives is a generalisation of the one defined in [3] for reachability objectives, and the proof techniques are more involved.

Then, using a simulation lemma, we show how to compute Nash equilibria in case the objectives of the players are given by deterministic Büchi automata. This encompasses many winning conditions (among which reachability, Büchi, safety, etc.), and we show that deciding the existence of Nash equilibria with constraints on the payoff is EXPTIME-complete. Under a certain restriction on the automata (1-weakness), we prove that the complexity reduces to PSPACE, and we prove PSPACE-hardness for the special case of safety objectives. When the game is *deterministic* and the 1-weak Büchi automata defining the winning conditions have bounded size (this includes safety and reachability objectives), we show that the constrained existence problem becomes NP-complete. The simulation lemma can also be used to lift our results to timed games, for which all our problems are EXPTIME-complete.

Related work. Concurrent and, more generally, stochastic games go back to Shapley [17]. However, most research in game theory and economics has focused on games with rewards, which are either averaged or discounted along an infinite path. In particular, Fink [11] proved that every discounted stochastic game has a Nash equilibrium in pure strategies, and Vielle [22] proved the existence of *ϵ-equilibria* in randomised strategies for two-player stochastic games under the average-reward criterion. For two-player concurrent games with Büchi objectives, the existence of ϵ-equilibria (in randomised strategies) was proved by Chatterjee [5]. However, exact Nash equilibria need not exist. An important subclass where even Nash equilibria in pure strategies exist are turn-based games with Büchi objectives [8].

The complexity of Nash equilibria in games played on graphs was first addressed in [8, 19]. In particular, it was shown in [19] that the existence of a Nash equilibrium with a constraint on its payoff can be decided in polynomial time for turn-based games with Büchi objectives. In this paper, we extend this result to concurrent games. It was also shown in [19] that the same problem is NP-hard for turn-based games with co-Büchi conditions, which implies hardness for concurrent games with this kind of objectives. For concurrent games with ω-regular objectives, the decidability of the constrained existence problem w.r.t. pure strategies was established by Fisman et al. [12], but their algorithm runs in doubly exponential time, whereas our algorithm for Büchi games runs in polynomial time. Finally, Ummels and Wojtczak [21] proved that the existence of a Nash equilibrium in pure or randomised strategies is undecidable for *stochastic* games with reachability or Büchi objectives, which justifies our restriction to concurrent games without probabilistic transitions (see [20] for a similar undecidability result for randomised Nash equilibria in non-stochastic games).

2 Preliminaries

2.1 Concurrent 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. In a transition system $\mathcal{S}$, a *path* π is a non-empty sequence $(s_i)_{0 \le i < n}$ (where $n \in \mathbb{N} \cup \{\infty\}$) of states 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*) of $\mathcal{S}$ is denoted by $\text{Hist}_{\mathcal{S}}$, 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 $\pi = (s_i)_{0 \le i < n}$ and an integer $j < n$, the *j-th prefix* (resp. *j-th suffix*, *j-th state*) of π, denoted by $\pi_{\le j}$ (resp. $\pi_{\ge j}$, $\pi_{=j}$), is the finite path $(s_i)_{0 \le i < j+1}$ (resp. $(s_i)_{j \le i < n}$, state s_j). If $\pi = (s_i)_{0 \le i < n}$ is a history, we write $\text{last}(\pi) = s_{|\pi|}$.

We consider nondeterministic concurrent games [3], which extend standard concurrent games [1] with nondeterminism.

▶ **Definition 1.** A *(nondeterministic) concurrent game* is a tuple $\mathcal{G} = \langle \text{States}, \text{Edg}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\mathcal{L}_A)_{A \in \text{Agt}} \rangle$, where $\langle \text{States}, \text{Edg} \rangle$ is a transition system, Agt is a finite set of players, Act is a (possibly uncountable) set of actions, and

- $\text{Mov}\colon \text{States} \times \text{Agt} \to 2^{\text{Act}} \setminus \{\varnothing\}$ is a mapping indicating the actions available to a given player in a given state;
- $\text{Tab}\colon \text{States} \times \text{Act}^{\text{Agt}} \to 2^{\text{Edg}} \setminus \{\varnothing\}$ associates with a state and an action profile the resulting set of edges; we require that $s = s'$ if $(s', s'') \in \text{Tab}(s, \langle m_A \rangle_{A \in \text{Agt}})$;
- $\mathcal{L}_A \subseteq \text{States}^\omega$ defines the *objective* for player $A \in \text{Agt}$; the *payoff* for player A is the function $\nu_A\colon \text{States}^\omega \to \{0, 1\}$, where $\nu_A(\pi) = 1$ if $\pi \in \mathcal{L}_A$, and $\nu_A(\pi) = 0$ otherwise; we say that player A *prefers* play π' over play π, denoted $\pi \preccurlyeq_A \pi'$, if $\nu_A(\pi) \le \nu_A(\pi')$.

We call a game $\mathcal{G}$ *finite* if its set of states is finite.

Non-determinism naturally appears in timed games, and this is our most important motivation for investigating this extension of standard concurrent games. We explain in Section 4.3 how our results apply to timed games.

We say that a *move* $\langle m_A \rangle_{A \in \text{Agt}} \in \text{Act}^{\text{Agt}}$ (which we may write m_{Agt} in the sequel) is *legal* at s if $m_A \in \text{Mov}(s, A)$ for all $A \in \text{Agt}$. A concurrent game is *deterministic* if $\text{Tab}(s, m_{\text{Agt}})$ is a singleton for each $s \in \text{States}$ and each legal move m_{Agt} (at s). A game is *turn-based* if for each state the set of allowed moves is a singleton for all but at most one player.

In a nondeterministic concurrent game, whenever we arrive at a state s, the players (simultaneously) choose a legal move m_{Agt}. Then, one of the transitions in $\text{Tab}(s, m_{\text{Agt}})$ is nondeterministically selected, which results in a new state of the game. In the sequel, we 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 2.** Let $\mathcal{G}$ be a concurrent game, and $A \in \text{Agt}$. A *strategy* for A is a mapping $\sigma_A\colon \text{Hist}_{\mathcal{G}} \to \text{Act}$ such that $\sigma_A(\pi) \in \text{Mov}(\text{last}(\pi), A)$ for all $\pi \in \text{Hist}_{\mathcal{G}}$. A strategy σ_P for a coalition P is a tuple of strategies, one for each player in P. We write $\sigma_P = \langle \sigma_A \rangle_{A \in P}$ for such a strategy. A *strategy profile* is a strategy for the coalition Agt. We write $\text{Strat}^P_{\mathcal{G}}$ for the set of strategies of coalition P (or simply $\text{Strat}^B_{\mathcal{G}}$ if $P = \{B\}$), and $\text{Prof}_{\mathcal{G}} = \text{Strat}^{\text{Agt}}_{\mathcal{G}}$.

Note that we only consider non-randomised (*pure*) strategies in this paper. Notice also that strategies are based on the sequences of visited states, and not on the actions played by the players. This is a realistic assumption for concurrent systems where different components interact with each other and each component has a set of internal actions which cannot be observed by the other components. However, it makes the computation of equilibria harder: when a deviation from the equilibrium profile occurs, the given sequence of states does not uniquely determine the player who has deviated. While the main part of the paper focuses on state-based strategies for Büchi objectives, we show in Section 6 that equilibria in the action-based setting can be computed more easily, even for parity objectives.

Let $\mathcal{G}$ be a game, P a coalition, and σ_P a strategy for P. A path $\pi = (s_j)_{0\le j\le|\pi|}$ is *compatible* with the strategy σ_P if, for all $k < |\pi|$, there exists a move m_{Agt} such that (i) m_{Agt} is legal at s_k, (ii) $m_A = \sigma_A(\pi_{\le k})$ for all $A \in P$, and (iii). $(s_k, s_{k+1}) \in \mathrm{Tab}(s_k, m_{\mathrm{Agt}})$. We write $\mathrm{Out}_{\mathcal{G}}(\sigma_P)$ for the set of paths (or *outcomes*) in $\mathcal{G}$ that are compatible with the strategy σ_P, and we write $\mathrm{Out}^{\mathrm{f}}_{\mathcal{G}}(\sigma_P)$ (resp. $\mathrm{Out}^{\infty}_{\mathcal{G}}(\sigma_P)$) for the finite (resp. infinite) outcomes, and $\mathrm{Out}_{\mathcal{G}}(s, \sigma_P)$, $\mathrm{Out}^{\mathrm{f}}_{\mathcal{G}}(s, \sigma_P)$ and $\mathrm{Out}^{\infty}_{\mathcal{G}}(s, \sigma_P)$ for the respective sets of outcomes that start in state s. In general there might be several infinite outcomes for a strategy profile from a given state. However, in the case of deterministic games, any strategy profile has a single infinite outcome from a given state.

▶ **Example 3.** Figure 1 depicts a two-player concurrent game, called the *matching-penny* game. A pair $\langle a, b\rangle$ represents a move, where Player 1 plays action a and Player 2 plays b. Starting from state s_0, if both players choose the same action, then the game proceeds to s_1; otherwise, the game proceeds to s_2. In the matching-penny game, the objective for Player 1 is to visit s_1 (which is encoded as $\mathcal{L}_1 = \mathrm{States}^* \cdot \{s_1\} \cdot \mathrm{States}^\omega$), while for Player 2 it is to visit s_2. Figure 2 shows another game (our running example), in which the objective of Player 1 is to *loop* in s_1 ($\mathcal{L}_1 = \mathrm{States}^* \cdot \{s_1\}^\omega$), whereas the objective for Player 2 is to *loop* in s_2 ($\mathcal{L}_2 = \mathrm{States}^* \cdot \{s_2\}^\omega$).

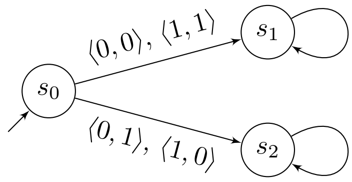

Figure 1 The matching-penny game

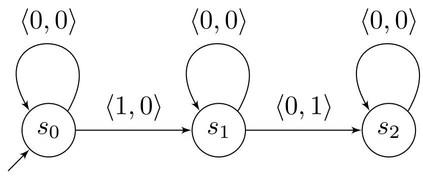

Figure 2 A game with Büchi objectives

2.2 Pseudo-Nash Equilibria

Given a move m_{Agt} and an action m' for some player B, we write $m_{\mathrm{Agt}}[B \mapsto m']$ for the move n_{Agt} with $n_A = m_A$ if $A \ne B$ and $n_B = m'$. This is extended to strategies in the natural way. For non-zero-sum games, several notions of equilibria have been defined, e.g. Nash equilibria [15], subgame-perfect equilibria [16], and secure equilibria [6]. None of these notions apply to nondeterministic games. Bouyer et al. have therefore proposed the notion of *pseudo-Nash equilibria* [3], which extend standard Nash equilibria to nondeterministic games.

▶ **Definition 4.** Let $\mathcal{G}$ be a nondeterministic concurrent game with objectives $(\mathcal{L}_A)_{A\in\mathrm{Agt}}$, and let s be a state of $\mathcal{G}$. A *pseudo-Nash equilibrium* of $\mathcal{G}$ from s is a pair $(\sigma_{\mathrm{Agt}}, \pi)$ of a strategy profile $\sigma_{\mathrm{Agt}} \in \mathrm{Prof}_{\mathcal{G}}$ and a play $\pi \in \mathrm{Out}_{\mathcal{G}}(s, \sigma_{\mathrm{Agt}})$ such that $\pi' \preccurlyeq_B \pi$ for all players $B \in \mathrm{Agt}$, all strategies $\sigma' \in \mathrm{Strat}^B$, and all plays $\pi' \in \mathrm{Out}_{\mathcal{G}}(s, \sigma_{\mathrm{Agt}}[B \mapsto \sigma'])$. The outcome π is then called an *optimal play* for the strategy profile σ_{Agt}.

For deterministic games, the play π is uniquely determined by σ_{Agt}, so that pseudo-Nash equilibria coincide with *Nash equilibria* [15]: these are strategy profiles where no player has an incentive to unilaterally deviate from her strategy.

In the case of nondeterministic games, a strategy profile for an equilibrium may give rise to several outcomes. The outcome π is then chosen cooperatively by all players: once a strategy profile is fixed, nondeterminism is resolved by all players choosing one of the possible outcomes in such a way that each player has no incentive to unilaterally changing her choice (nor her strategy). To the best of our knowledge, this cannot be encoded by adding an extra player for resolving the nondeterminism.

▶ **Example 5.** Clearly, there is no pure-strategy Nash equilibrium in the game of Figure 1 since a losing player can always improve her payoff by switching her choice. In the game of Figure 2 (where Player i wants to visit s_i infinitely often), there are several Nash equilibria: one with payoff $(0, 1)$, in which both players play action 1 when it is available; another one with payoff $(0, 0)$, in which Player 1 always plays 0, and Player 2 plays 1 when available.

In this paper, we study several decision problems related to the existence of pseudo-Nash equilibria. The *existence problem* consists in deciding the existence of a pseudo-Nash equilibrium in a given state of a game. Since several pseudo-Nash equilibria may coexist, it is also interesting to decide whether there is one with a given payoff (0 or 1) for some of the players; this is the *constrained existence problem*. Finally, the *verification problem* asks whether a given payoff function (totally defined over Agt) is the payoff of some pseudo-Nash equilibrium. Notice that the first and third problems are trivially logspace-reducible to the second one.

3 Internal Büchi Objectives

In this section, we fix a nondeterministic concurrent game $\mathcal{G} = \langle \text{States}, \text{Edg}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\mathcal{L}_A)_{A \in \text{Agt}} \rangle$, where the objectives are *internal* Büchi conditions given by a set $\Omega_A \subseteq \text{States}$ of target states for each player $A \in \text{Agt}$. The corresponding objective for player A is the set $\mathcal{L}_A = \{\pi \in \text{States}^\omega \mid \pi_{=j} \in \Omega_A \text{ for infinitely many } j \in \mathbb{N}\}$.

3.1 Characterising Equilibria Using Fixpoints

In [3], pseudo-Nash equilibria are characterised for qualitative reachability objectives using a fixpoint computation called the *repellor*. This was the counter-part of the attractor in non-zero-sum games for computing equilibria. In this section, we extend the repellor to handle internal Büchi objectives.

Suspect players. Let $e = (s, s')$ be an edge. Given a move m_{Agt}, we define the set of suspect players for e as the set

$$\text{Susp}(e, m_{\text{Agt}}) = \{B \in \text{Agt} \mid \exists m' \in \text{Mov}(s, B) \text{ such that } e \in \text{Tab}(s, m_{\text{Agt}}[B \mapsto m'])\}.$$

Intuitively, Player $B \in \text{Agt}$ is a suspect for edge e and if she can unilaterally change her action to trigger edge e. Notice that if $e \in \text{Tab}(s, m_{\text{Agt}})$, then $\text{Susp}(e, m_{\text{Agt}}) = \text{Agt}$.

The iterated (or Büchi) repellor. For any $n \in \mathbb{N}$ and $P \subseteq \text{Agt}$, we define the n*-th repellor set* $\text{Rep}^n_{\mathcal{G}}(P)$ as follows. If $n = 0$, then $\text{Rep}^0_{\mathcal{G}}(P) = \emptyset$ for any $P \subseteq \text{Agt}$. Now fix $n \in \mathbb{N}$, and assume that repellor sets $\text{Rep}^n_{\mathcal{G}}(P)$ have been defined for any $P \subseteq \text{Agt}$. As the base case for level $n + 1$, we set $\text{Rep}^{n+1}_{\mathcal{G}}(\emptyset) = \text{States}$. Then, assuming that $\text{Rep}^{n+1}_{\mathcal{G}}(P')$ has been defined for all $P' \subsetneq P$, we let $\mathrm{Rep}_{\mathcal{G}}^{n+1}(P)$ be the largest set fulfilling the following condition: for all $s \in \mathrm{Rep}_{\mathcal{G}}^{n+1}(P)$ there exists a legal move m_{Agt} (at s) such that

1. $s' \in \mathrm{Rep}_{\mathcal{G}}^{n+1}(P \cap \mathrm{Susp}_{\mathcal{G}}((s,s'), m_{\mathrm{Agt}}))$ for all $s' \in \mathrm{States}$, and
2. if $s' \in \Omega_A$ for some player $A \in P \cap \mathrm{Susp}_{\mathcal{G}}((s,s'), m_{\mathrm{Agt}})$, then $s' \in \mathrm{Rep}_{\mathcal{G}}^{n}(P \cap \mathrm{Susp}_{\mathcal{G}}((s,s'), m_{\mathrm{Agt}}))$.

Given a state $s \in \mathrm{Rep}_{\mathcal{G}}^{n+1}(P)$, a legal move m_{Agt} that fulfils 1. and 2. is called a *secure move* (w.r.t. P and $n+1$); we write $\mathrm{Secure}_{\mathcal{G}}^{n+1}(s,P)$ for the set of these moves. Finally, we define $\mathrm{Rep}_{\mathcal{G}}^{\infty}(P) = \bigcup_{n\geq 0} \mathrm{Rep}^n(P)$. In the following, to improve readability, we will omit the index $\mathcal{G}$ in all the notions we have defined, when the game is clear from the context.

Intuitively, a state s is an element of $\mathrm{Rep}_{\mathcal{G}}^{n}(P)$ iff at s there is a legal move such that no player $A \in P$ can force to visit her set of target states at least n times by changing her action. For finite games, it follows that a state s is an element of $\mathrm{Rep}_{\mathcal{G}}^{\infty}(P)$ iff at s there is a legal move such that no player $A \in P$ can force to visit her set set of target states *infinitely often* by changing her action.

▶ Remark. The repellor defined for reachability objectives in [3] is rather similar to $\mathrm{Rep}^1(P)$; it differs only in the second condition, which was "$\mathrm{Rep}^1(P) \cap \Omega_A = \emptyset$ for all $A \in P$" in [3]. This change is required since a play that is losing w.r.t. a Büchi objective might visit a winning state a finite number of times (whereas this cannot happen for reachability objectives).

▶ **Example 6.** In the game of Figure 2, if we assume reachability objectives (state s_i for Player i), there is no equilibrium with payoff $(0,0)$, since Player 1 can enforce a visit to her winning state. If we assume Büchi objectives, we have seen in Example 5 that there is an equilibrium with payoff $(0,0)$. Table 1 displays the values of the iterated repellors in this game, for all possible sets of players. These results were obtained with our prototype implementation of our algorithms, available at `http://www.lsv.ens-cachan.fr/Software/praline/`.

Table 1 Computing the repellor sets in the game of Figure 2

P	$\mathrm{Rep}^0(P)$	$\mathrm{Rep}^1(P)$	$\mathrm{Rep}^2(P)$	$\mathrm{Rep}^\infty(P) = \mathrm{Rep}^3(P)$
$\emptyset$	$\emptyset$	$\{s_0, s_1, s_2\}$	$\{s_0, s_1, s_2\}$	$\{s_0, s_1, s_2\}$
$\{A_1\}$	$\emptyset$	$\{s_1, s_2\}$	$\{s_0, s_1, s_2\}$	$\{s_0, s_1, s_2\}$
$\{A_2\}$	$\emptyset$	$\{s_0\}$	$\{s_0\}$	$\{s_0\}$
$\{A_1, A_2\}$	$\emptyset$	$\emptyset$	$\{s_0\}$	$\{s_0\}$

▶ **Lemma 7.** *The repellor and the secure moves satisfy the following properties:*

- *If* $P' \subseteq P \subseteq \mathrm{Agt}$, *then* $\mathrm{Rep}^n(P) \subseteq \mathrm{Rep}^n(P')$ *for all* $n \in \mathbb{N}$.
- $\mathrm{Rep}^n(P) \subseteq \mathrm{Rep}^{n+1}(P)$ *for all* $P \subseteq \mathrm{Agt}$ *and* $n \in \mathbb{N}$.
- $\mathrm{Secure}^n(s,P) \subseteq \mathrm{Secure}^{n+1}(s,P)$ *for all* $P \subseteq \mathrm{Agt}$, $n \in \mathbb{N}$ *and* $s \in \mathrm{States}$.

We define the *n-th repellor transition system* $\mathcal{S}^n(P) = \langle \mathrm{States}, \mathrm{Edg}_n \rangle$ by $(s,s') \in \mathrm{Edg}_n$ iff there exists $m_{\mathrm{Agt}} \in \mathrm{Secure}^n(s,P)$ such that $(s,s') \in \mathrm{Tab}(s, m_{\mathrm{Agt}})$. Note in particular that any $s \in \mathrm{Rep}^n(P)$ has an outgoing transition in $\mathcal{S}^n(P)$. We also define the limit repellor transition system $\mathcal{S}^\infty(P) = \langle \mathrm{States}, \bigcup_{n\geq 0} \mathrm{Edg}_n \rangle$. The following lemma bounds the number of iteration steps required to reach $\mathrm{Rep}^\infty(P)$.

▶ **Lemma 8.** *Let $\mathcal{G}$ be a a finite game, $P \subseteq \mathrm{Agt}$, and let ℓ be the length of the longest acyclic path in $\mathcal{G}$. Then $\mathrm{Rep}^n(P) = \mathrm{Rep}^\infty(P)$ for all $n \geq \ell \cdot |P|$.*

The correctness of the iterated repellor for *finite* games is stated in the next proposition.

▶ **Proposition 9.** *Let $\mathcal{G}$ be a finite game, $P \subseteq \text{Agt}$, and let $\rho \in \text{Play}(s)$ be a play that visits $\bigcup_{B\in P} \Omega_B$ only finitely often. Then ρ is a path in $\mathcal{S}^\infty(P)$ if and only if there exists $\sigma_{\text{Agt}} \in \text{Prof}$ such that $\rho \in \text{Out}(s, \sigma_{\text{Agt}})$ and ρ' does not visit Ω_B infinitely often for all plays ρ' that can arise when some player $B \in P$ changes her strategy, i.e. $\rho' \in \text{Out}(s, \sigma_{\text{Agt}}[B \mapsto \sigma'])$ for some $B \in P$ and some $\sigma' \in \text{Strat}^B$.*

We can deduce from this proposition that if ρ is an infinite path from state s in $\mathcal{S}^\infty(P)$ that visits Ω_A infinitely often *if and only if* $A \notin P$, then there is a pseudo Nash equilibrium from s with optimal play ρ.

▶ **Corollary 10.** *Let $\mathcal{G}$ be a finite game, $s \in \text{States}$, and $\nu\colon \text{Agt} \to \{0,1\}$. There exists a pseudo-Nash equilibrium in $\mathcal{G}$ with payoff ν if and only if there exists an infinite path ρ in $\mathcal{S}^\infty(\nu^{-1}(0))$ with payoff $\nu_A(\rho) = \nu(A)$ for all $A \in \text{Agt}$.*

3.2 Application to Solving the Three Problems

We use the previous characterisation for analysing the complexity of the various decision problems that we have defined in Section 2.2.

▶ **Theorem 11.** *The verification, existence and constrained existence problems for finite games with internal Büchi objectives are* PTIME*-complete.*

The lower bounds are simple adaptations of the PTIME-hardness of the circuit value problem. We now focus on the PTIME upper bounds, and prove it for the constrained existence problem (which implies the same upper bound for the other two problems).

We first use the equivalence given in Proposition 9 to get a set-based characterisation of (pseudo-)Nash equilibria. We fix a set of winning players $W \subseteq \text{Agt}$ and a set of losing players $L \subseteq \text{Agt}$, and we fix an initial state s. Given a transition system $\langle S, E\rangle$ and a set of players P, we say that they satisfy condition (‡) if the following properties are fulfilled:

(1) $\Omega_A \cap S = \emptyset$ if and only if $A \in P$;
(2) $L \subseteq P$ and $P \cap W = \emptyset$;
(3) $\langle S, E\rangle$ is strongly connected;
(4) $\langle S, E\rangle \subseteq \mathcal{S}^\infty(P)$;
(5) S is reachable from s in $\mathcal{S}^\infty(P)$.

The following is then a corollary to Proposition 9.

▶ **Corollary 12** (Set-based characterisation)**.** *A pair $(\langle S, E\rangle, P)$ satisfies condition (‡) iff there is an infinite path ρ in $\mathcal{S}^\infty(P)$ from s that, from some point onwards, stays in $\langle S, E\rangle$. In particular, ρ is losing for all players in L. Moreover, if $(\langle S, E\rangle, P)$ satisfies (‡), then there exists an infinite path ρ in $\mathcal{S}^\infty(P)$ from s with the same property that visits* all *states of S infinitely often (and is thus winning for all players in W).*

Note that in the above characterisation, P is uniquely determined by the set S; hence we write $P(S) = \{A \in \text{Agt} \mid S \cap \Omega_A = \emptyset\}$, and we say that $\langle S, E\rangle$ satisfies condition (‡) if $(\langle S, E\rangle, P(S))$ does. Our aim is to compute in polynomial time all *maximal* pairs $\langle S, E\rangle$ that satisfy condition (‡). As a prerequisite, we assume that we can compute $\mathcal{S}^\infty(P)$ in polynomial time whenever $P \subseteq \text{Agt}$ is given. This can be proved using similar arguments as in [4]. Now, we define a recursive operator SSG (SolveSubGame) by setting $\mathsf{SSG}(\langle S, E\rangle) = \{\langle S, E\rangle\}$ if $\langle S, E\rangle \subseteq \mathcal{S}^\infty(P(S))$ and $\langle S, E\rangle$ is strongly connected, and

$$\mathsf{SSG}(\langle S, E\rangle) = \bigcup_{\langle S', E'\rangle \in \text{SCC}(\langle S, E\rangle)} \mathsf{SSG}(\langle S', E'\rangle \cap \mathcal{S}^\infty(P(S')))$$

in all other cases. Here, $\mathrm{SCC}(\langle S, E\rangle)$ denotes the set of strongly connected components of $\langle S, E\rangle$ (which can be computed in linear time). Finally, we define

$$\mathsf{Sol}(L, W) = \mathsf{SSG}\Big(\langle \mathrm{States} \setminus \bigcup_{A\in L} \Omega_A, \mathrm{Edg}\rangle\Big) \cap \{\langle S, E\rangle \mid P(S) \cap W = \emptyset\}\,.$$

▶ **Lemma 13.** *If $\langle S, E\rangle \in \mathsf{Sol}(L, W)$ and S is reachable from s in $\mathcal{S}^\infty(P(S))$, then it satisfies condition (‡). Conversely, if $\langle S, E\rangle$ satisfies condition (‡), then there exists $\langle S', E'\rangle \in \mathsf{Sol}(L, W)$ such that $\langle S, E\rangle \subseteq \langle S', E'\rangle$.*

▶ **Lemma 14.** *The set $\mathsf{Sol}(L, W)$ can be computed in polynomial time.*

The PTIME upper bound of Theorem 11 follows from the above analysis.

▶ Remark. This result may seem surprising since we know that the problem is NP-complete for reachability objectives, even in turn-based games [8, 3]. Intuitively, the problem is harder for reachability objectives because whether a play satisfies or not a reachability objective is not only determined by its behaviour in the strongly connected component in which it settles but on *all* visited states.

4 Game Simulations

Our aim is to transfer our results for internal Büchi objectives to larger classes of objectives. A useful tool is the notion of game simulation, which we develop now.

4.1 Definition and General Properties

We already gave a definition of game simulation in [3], which was tailored to games with reachability objectives; we extend this notion to games with arbitrary qualitative objectives.

▶ **Definition 15.** Consider two games $\mathcal{G} = \langle \mathrm{States}, \mathrm{Edg}, \mathrm{Agt}, \mathrm{Act}, \mathrm{Mov}, \mathrm{Tab}, (\mathcal{L}_A)_{A\in\mathrm{Agt}}\rangle$ and $\mathcal{G}' = \langle \mathrm{States}', \mathrm{Edg}', \mathrm{Agt}, \mathrm{Act}', \mathrm{Mov}', \mathrm{Tab}', (\mathcal{L}'_A)_{A\in\mathrm{Agt}}\rangle$ with the same set Agt of players. A relation $\triangleleft \subseteq \mathrm{States} \times \mathrm{States}'$ is a *game simulation* if $s \triangleleft s'$ implies that for each move m_{Agt} in $\mathcal{G}$ there exists a move m'_{Agt} in $\mathcal{G}'$ such that

1. for each $t' \in \mathrm{States}'$ there exists $t \in \mathrm{States}$ with $t \triangleleft t'$ and $\mathrm{Susp}((s', t'), m'_{\mathrm{Agt}}) \subseteq \mathrm{Susp}((s, t), m_{\mathrm{Agt}})$, and
2. for each $(s, t) \in \mathrm{Tab}(s, m_{\mathrm{Agt}})$ there exists $(s', t') \in \mathrm{Tab}'(s', m'_{\mathrm{Agt}})$ with $t \triangleleft t'$.

If $\triangleleft$ is a game simulation, we say that $\mathcal{G}'$ *simulates* $\mathcal{G}$. Finally, a game simulation $\triangleleft$ is *winning-preserving* from $(s_0, s'_0) \in \mathrm{States} \times \mathrm{States}'$ if for all $\rho \in \mathrm{Play}_{\mathcal{G}}(s_0)$ and $\rho' \in \mathrm{Play}_{\mathcal{G}'}(s'_0)$ with $\rho \triangleleft \rho'$ (i.e. $\rho_{=i} \triangleleft \rho'_{=i}$ for all $i \in \mathbb{N}$) it holds that $\rho \in \mathcal{L}_A$ iff $\rho' \in \mathcal{L}'_A$ for all $A \in \mathrm{Agt}$.

▶ **Proposition 16.** *Game simulation is transitive.*

▶ **Proposition 17.** *Assume $\mathcal{G}$ and $\mathcal{G}'$ are games. Fix two states s and s' in $\mathcal{G}$ and $\mathcal{G}'$ respectively, and assume that $\triangleleft$ is a winning-preserving game simulation from (s, s'). If there exists a pseudo-Nash equilibrium $(\sigma_{\mathrm{Agt}}, \rho)$ of $\mathcal{G}$ from s, then there exists a pseudo-Nash equilibrium $(\sigma'_{\mathrm{Agt}}, \rho')$ of $\mathcal{G}'$ from s' with $\rho \triangleleft \rho'$. In particular, ρ and ρ' have the same payoff.*

4.2 Product of a Game with Deterministic Büchi Automata

We use the results on game simulation to study objectives that are defined by deterministic Büchi automata. A *deterministic Büchi automaton* $\mathcal{A}$ over alphabet Σ is a tuple $\langle Q, \Sigma, \delta, q_0, R\rangle$, where Q is a finite set of *states*, Σ is the *input alphabet*, $\delta\colon Q \times \Sigma \to Q$ is the *transition function*, $q_0 \in Q$ is the *initial* state, and $R \subseteq Q$ is the set of *repeated* states. We assume that the reader is familiar with Büchi automata, and we write $\mathrm{L}(\mathcal{A}) \subseteq \Sigma^\omega$ for the language accepted by $\mathcal{A}$.

Fix a game $\mathcal{G} = \langle \mathrm{States}, \mathrm{Edg}, \mathrm{Agt}, \mathrm{Act}, \mathrm{Mov}, \mathrm{Tab}, (\mathcal{L}_B)_{B\in\mathrm{Agt}}\rangle$ and a player $A \in \mathrm{Agt}$, and assume that $\mathcal{L}_A = \mathrm{L}(\mathcal{A})$ for a deterministic Büchi automaton $\mathcal{A} = \langle Q, \mathrm{States}, \delta, q_0, R\rangle$ over States. We show how to compute pseudo-Nash equilibria in $\mathcal{G}$ by building a product of $\mathcal{G}$ with $\mathcal{A}$ and computing the pseudo-Nash equilibria in the resulting game.

We define the product of the game $\mathcal{G}$ with the automaton $\mathcal{A}$ as the game $\mathcal{G} \ltimes \mathcal{A} = \langle \mathrm{States}', \mathrm{Edg}', \mathrm{Agt}, \mathrm{Act}, \mathrm{Mov}', \mathrm{Tab}', (\mathcal{L}'_B)_{B\in\mathrm{Agt}}\rangle$, where:

- $\mathrm{States}' = \mathrm{States} \times Q$;
- $\mathrm{Edg}' = \{((s,q),(s',q')) \mid (s,s') \in \mathrm{Edg} \text{ and } \delta(q,s) = q'\}$;
- $\mathrm{Mov}'((s,q),A_i) = \mathrm{Mov}(s,A_i)$ for every $A_i \in \mathrm{Agt}$;
- $\mathrm{Tab}'((s,q),m_{\mathrm{Agt}}) = \{((s,q),(s',q')) \mid (s,s') \in \mathrm{Tab}(s,m_{\mathrm{Agt}}) \text{ and } \delta(q,s) = q'\}$;
- if $B = A$, then $\mathcal{L}'_B$ is the internal Büchi objective given by the set $\Omega = \mathrm{States} \times R$; otherwise, $\mathcal{L}'_B = \pi^{-1}(\mathcal{L}_B)$, where π is the natural projection from States' to States and its extension to plays (i.e. $\pi((s_0,q_0)(s_1,q_1)\dots) = s_0 s_1 \dots$).

▶ Remark. Note that, if $\mathcal{L}_B$ is defined by an internal Büchi condition, then so is $\mathcal{L}'_B$.

▶ **Lemma 18.** *$\mathcal{G} \ltimes \mathcal{A}$ simulates $\mathcal{G}$, and vice versa. Furthermore, in both cases we can exhibit a game-simulation that is winning-preserving from $(s,(s,q_0))$ for all $s \in \mathrm{States}$.*

Assume that for each player $A_i \in \mathrm{Agt}$ the objective in $\mathcal{G}$ is given by a deterministic Büchi automaton $\mathcal{A}_i$. We use the transitivity of game simulation to build a product of $\mathcal{G}$ with each of the automata $\mathcal{A}_i$, namely $\mathcal{G}' = \mathcal{G} \ltimes \mathcal{A}_1 \ltimes \dots \ltimes \mathcal{A}_n$ (we assume that $\ltimes$ is left-associative). Each player $A_i \in \mathrm{Agt}$ has an internal Büchi objective in $\mathcal{G}'$, which we denote by Ω_i.

▶ **Corollary 19.** *Let $s \in \mathrm{States}$ and $\nu\colon \mathrm{Agt} \to \{0,1\}$. There is a pseudo-Nash equilibrium in $\mathcal{G}$ from s with payoff ν if and only if there is a pseudo-Nash equilibrium in $\mathcal{G}'$ from $(s, q_{01}, \dots, q_{0n})$ with payoff ν, where q_{0i} is the initial state of $\mathcal{A}_i$.*

4.3 Application to Timed Games

We now apply the game-simulation approach to the computation of pseudo-Nash equilibria in timed games. Given a timed game $\mathcal{G}$ with internal Büchi objectives, the corresponding (exponential-size) region game $\mathcal{R}_\mathcal{G}$ as defined in [4] simulates $\mathcal{G}$ and is simulated by $\mathcal{G}$ while preserving winning conditions (the proof for reachability objectives in [4] can be easily extended to our framework). Pseudo-Nash equilibria of $\mathcal{G}$ can thus be computed on the finite game $\mathcal{R}_\mathcal{G}$. If the objectives of the players are defined by deterministic Büchi automata $(\mathcal{A}_i)_{A_i\in\mathrm{Agt}}$, we can compute the product $\mathcal{R}_\mathcal{G} \ltimes \mathcal{A}_1 \ltimes \dots \ltimes \mathcal{A}_n$ with corresponding internal Büchi objectives $(\Omega_i)_{A_i\in\mathrm{Agt}}$, as defined in the previous subsection. This product has size exponential in the size of $\mathcal{G}$ and in the number of players. We can then apply the algorithm developed in Section 3.2, yielding an EXPTIME upper bound for deciding the verification, existence, and constrained existence problems in timed games. Finally we get EXPTIME-hardness for internal Büchi objectives by applying the reduction in [4, Prop. 20] (replace all accepting sink states by repeated states).

▶ **Theorem 20.** *The verification, existence, and constrained existence problems are* EXPTIME-*complete both for timed games with internal Büchi objectives and for timed games with objectives defined by deterministic Büchi automata.*

5 Büchi-Definable Objectives

The characterisation of Corollary 19 gives a procedure to compute pseudo-Nash equilibria in games with objectives defined by deterministic Büchi automata (one automaton per player). The algorithm runs in time exponential in the number of players since we have to build the product with all the deterministic Büchi automata defining the objective of a player. We prove that our problems are EXPTIME-hard by encoding two-player countdown games [14] into multiplayer games. Each bit of the countdown will be managed by a different player, who is in charge of checking that this bit is correctly updated at each transition.

▶ **Theorem 21.** *The verification, existence, and constrained existence problems for finite games with objectives defined by deterministic Büchi automata are* EXPTIME-*complete.*

We now prove that when the deterministic Büchi automata defining the objectives are 1-*weak* (i.e. when all strongly connected components of the transition graph consist of just one state), all our three problems can be solved in PSPACE. In particular, this result applies to safety (and reachability) objectives, which can be defined by 1-weak automata. Our algorithm is based on a procedure that, given parameters (P, n, q), computes the set of states s such that in the product game $(s, q) \in \mathrm{Rep}^n(P)$. The procedure computes the repellor as a fixpoint, calling itself recursively on instances (P', n', q'), where either $P' \subsetneq P$, $n' < n$, or q' is a successor of q. The maximal number of nested calls is $|P| + n + \sum_{i \in \mathrm{Agt}} \ell_i$, where ℓ_i is the length of the longest acyclic path in $\mathcal{A}_i$. According to Lemma 8, n can be bounded by a polynomial. The whole computation thus runs in polynomial space.

▶ **Theorem 22.** *The verification, existence, and constrained existence problems are in* PSPACE *for finite games with objectives defined by* 1*-weak deterministic Büchi automata.*

The matching lower bound holds already for the special case of safety objectives.

▶ **Proposition 23.** *The verification, existence, and constrained existence problems are* PSPACE*-hard for finite games with safety objectives.*

Proof. The hardness proof for the verification (and for the constrained existence) problem is by a reduction from QSAT: for every closed quantified Boolean formula ϕ in conjunctive prenex normal form, we construct a game $\mathcal{G}(\phi)$ with initial state s_1 and safety objectives such that $\mathcal{G}(\phi)$ has a pseudo-Nash equilibrium with payoff $(0, \ldots, 0)$ from s_1 iff the formula is true. Let $\phi = \exists x_1 \forall x_2 \ldots Q_n x_n \; C_1 \wedge C_2 \wedge \cdots \wedge C_m$, where each clause C_j is a disjunction of literals over the variables $x_1, \ldots, x_n$; we identify C_j with the set of literals occurring in it. Then the game $\mathcal{G}(\phi)$ is played by players $0, 1, \ldots, m$. The set of states is $\{s_1, x_1, \neg x_1, \ldots, s_n, x_n, \neg x_n, s_{n+1}\}$, and there are transitions from s_i to x_i and $\neg x_i$, and from x_i and $\neg x_i$ to s_{i+1}; additionally, there is a transition from s_{n+1} back to s_{n+1}. If i is odd, then the state s_i is controlled by player 0; otherwise, the game proceeds nondeterministically from s_i to either x_i or $\neg x_i$ (see Figure 3). Player 0 loses every play of the game, i.e. $\mathcal{L}_0 = \emptyset$, whereas for $j > 0$ player j's objective is to avoid the set of literals occurring in the clause C_j, i.e. $\mathcal{L}_j = (\mathrm{States} \setminus C_j)^\omega$. It is easy to see that ϕ is true iff there is a strategy for player 0 such that all outcomes are losing for all players.

To prove hardness of the existence problem, it suffices to add two states s_0 and s'_0 to the game $\mathcal{G}(\phi)$: from s_0, the game proceeds nondeterministically to either s'_0 or s_1, and we add

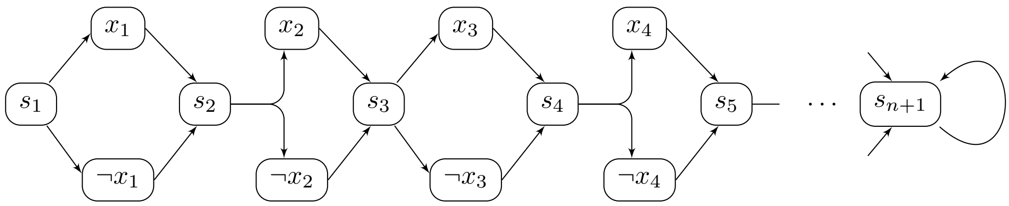

Figure 3 Reducing from QSAT

a transition from s'_0 back to s'_0. Finally, the objective of player 0 is the set $\mathcal{L}_0 = \{s_0, s'_0\}^\omega$. It follows that there is a pseudo-Nash equilibrium from s_0 (with the optimal play leading to s'_0) iff there is a pseudo-Nash equilibrium from s_1 with payoff $(0, \dots, 0)$. ◀

Note that the hardness proof requires nondeterminism. For deterministic games we can solve the problem by guessing the set of losing players and an (ultimately-periodic) path in the corresponding repellor transition system. We then have to check that all possible deviations fall in some repellor set. The best algorithm we could get for this check runs in time $O(|\text{States}|^2 \cdot |\text{Agt}| \cdot |\text{Tab}|^{\log(\max_i |Q_i|)})$, which is only polynomial when the size of the Büchi automata is bounded.

▶ **Theorem 24.** *The verification, existence, and constrained existence problems are in* NP *for finite deterministic games with objectives defined by* 1*-weak deterministic Büchi automata of bounded size.*

The matching lower bound holds again for the special case of safety objectives since no nondeterministic transitions arise in the construction used for proving Proposition 23 when we reduce from SAT (except for the existence problem, where we require a different construction).

▶ **Proposition 25.** *The verification, existence and constrained existence problems are* NP*-hard for finite deterministic games with safety objectives.*

6 Discussion

In this paper we focused on Büchi objectives. The natural next step is to go to parity objectives, which can encode arbitrary ω-regular objectives. In the turn-based case, the constrained existence problem becomes NP-complete for parity (or even co-Büchi) objectives [19]. In fact, we can get the same upper bound in deterministic concurrent games under the assumption that strategies *can* observe actions.

▶ **Theorem 26.** *The constrained existence problem is in* NP *for finite deterministic concurrent games with parity objectives if we assume that strategies can observe actions.*

In Section 2, we mentioned that making actions unobservable by players is a relevant modelling assumption, but that it makes the computation of equilibria harder. This claim is justified by the following result, which is obtained by a reduction from the strategy problem for *generalised parity games* [7].

▶ **Proposition 27.** *The verification, existence and constrained existence problems are* coNP*-hard for finite deterministic concurrent games with parity objectives. In particular, unless* NP = coNP*, these problems do not belong to* NP*.*

A natural question is whether the repellor techniques that we develop can be used to handle imperfect information in a more general sense than just state-based vs. action-based strategies. This is one of our directions for future work.

References

1 R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

2 E. Asarin, O. Maler, A. Pnueli, and J. Sifakis. Controller synthesis for timed automata. In SSC'98, p. 469–474. Elsevier Science, 1998.

3 P. Bouyer, R. Brenguier, and N. Markey. Nash equilibria for reachability objectives in multi-player timed games. In CONCUR'10, LNCS 6269, p. 192–206. Springer, 2010.

4 P. Bouyer, R. Brenguier, and N. Markey. Nash equilibria for reachability objectives in multiplayer timed games. Research Report LSV-10-12, ENS Cachan, France, 2010.

5 K. Chatterjee. Two-player nonzero-sum ω-regular games. In CONCUR'05, LNCS 3653, p. 413–427. Springer, 2005.

6 K. Chatterjee, T. A. Henzinger, and M. Jurdziński. Games with secure equilibria. In LICS'04, p. 160–169. IEEE Comp. Soc. Press, 2004.

7 K. Chatterjee, T. A. Henzinger, and N. Piterman. Generalized parity games. In FoSSaCS'07, LNCS 4423, p. 153–167. Springer, 2007.

8 K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In CSL'04, LNCS 3210, p. 26–40. Springer, 2004.

9 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.

10 M. Félegyházi, J.-P. Hubaux, and L. Buttyán. Nash equilibria of packet forwarding strategies in wireless ad hoc networks. *IEEE Trans. Mobile Comput.*, 5(5):463–476, 2006.

11 A. M. Fink. Equilibrium in a stochastic n-person game. *J. Science in Hiroshima University*, 28(1):89–93, 1964.

12 D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In TACAS'10, LNCS 6015, p. 190–204. Springer, 2010.

13 T. A. Henzinger. Games in system design and verification. In TARK'05, p. 1–4, 2005.

14 M. Jurdziński, F. Laroussinie, and J. Sproston. Model checking probabilistic timed automata with one or two clocks. In TACAS'07, LNCS 4424, p. 170–184. Springer, 2007.

15 J. F. Nash, Jr. Equilibrium points in n-person games. *Proc. National Academy of Sciences of the USA*, 36(1):48–49, 1950.

16 R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301–324 and 667–689, 1965.

17 L. S. Shapley. Stochastic games. *Proc. National Academy of Sciences of the USA*, 39:1095–1100, 1953.

18 W. Thomas. Infinite games and verification (extended abstract of a tutorial). In CAV'02, LNCS 2404, p. 58–64. Springer, 2002.

19 M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In FoSSaCS'08, LNCS 4962, p. 20–34. Springer, 2008.

20 M. Ummels and D. Wojtczak. The complexity of Nash equilibria in limit-average games. In CONCUR'11, LNCS 6901, p. 482–496. Springer, 2011.

21 M. Ummels and D. Wojtczak. The complexity of Nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.

22 N. Vielle. Two-player stochastic games I–II. *Israel J. of Mathematics*, 119(1):55–126, 2000.