

The Complexity of Nash Equilibria in Infinite Multiplayer Games*

Michael Ummels

Mathematische Grundlagen der Informatik, RWTH Aachen, Germany
E-Mail: ummels@logic.rwth-aachen.de

ABSTRACT. *We study the complexity of Nash equilibria in infinite (turn-based, qualitative) multiplayer games. Chatterjee & al. showed the existence of a Nash equilibrium in any such game with ω -regular winning conditions, and they devised an algorithm for computing one. We argue that in applications it is often insufficient to compute just some Nash equilibrium. Instead, we enrich the problem by allowing to put (qualitative) constraints on the payoff of the desired equilibrium. Our main result is that the resulting decision problem is NP-complete for games with co-Büchi, parity or Streett winning conditions but fixed-parameter tractable for many natural restricted classes of games with parity winning conditions. For games with Büchi winning conditions we show that the problem is, in fact, decidable in polynomial time.*

We also analyse the complexity of strategies realising a Nash equilibrium. In particular, we show that pure finite-state strategies as opposed to arbitrary mixed strategies suffice to realise any Nash equilibrium of a game with ω -regular winning conditions with a qualitative constraint on the payoff.

1 Introduction

We study *infinite games of perfect information* [10] played by multiple players on a finite directed graph. Intuitively, a *play* of such a game evolves by moving a token along edges of the graph. Every vertex of the graph is controlled by precisely one player. Whenever the token arrives at some vertex, the player who controls this vertex must move the token to a successor vertex. Thus a play of such a game is an infinite path through the graph. Plays are mapped to *payoffs*, one for each player. In the simplest case, which we discuss here, payoffs are just 0 and 1, i.e. each player either wins or loses a given play of the game. In this case, the payoff function of each player can be described by the set of plays where she receives payoff 1, her *winning condition*.

Infinite games have been successfully applied in the verification and synthesis of reactive systems. Such a system is usually modelled as a game between the system and its environment where the environment's objective

*This research has been supported by the DFG Research Training Group "Algorithmic Synthesis of Reactive and Discrete-Continuous Systems" (ALCoSyn).

is the complement of the system's objective, so the environment is considered hostile. Therefore, traditionally, the research in this area has mostly looked at two-player games where each play is won by precisely one of the two players, so-called *two-player zero-sum games*. However, motivated by the modelling of distributed systems, interest in the general case has increased in recent years [3, 4].

The most common interpretation of rational behaviour in multiplayer games is captured by the notion of a *Nash equilibrium*. In a Nash equilibrium, no player can improve her payoff by unilaterally switching to a different strategy. Chatterjee & al. [4] showed that any infinite multiplayer game with ω -regular winning conditions has a Nash equilibrium in pure strategies, and they also gave an algorithm for computing one. We argue that this is not satisfactory. Indeed, it can be shown that their algorithm may compute an equilibrium where all players lose when there exist other equilibria where all players win.

In applications, one might look for an equilibrium where as many players win as possible or where it is guaranteed that certain players win while certain others lose. Formulated as a decision problem, we want to know, given a k -player game \mathcal{G} with initial vertex v_0 and two payoff thresholds $\bar{x}, \bar{y} \in \{0, 1\}^k$, whether (\mathcal{G}, v_0) has a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$.

When restricted to two-player zero-sum games, this problem, which we call NE for short, is nothing else than the classical game problem of deciding the winner. For parity games, the latter problem is known to be in $UP \cap co-UP$ [12] and therefore unlikely to be NP-hard. Moreover, if e.g. the number of priorities is bounded (as it is the case for Büchi and co-Büchi winning conditions), the problem is known to be decidable in polynomial time.

Our main result is that NE is NP-complete for games with Streett, parity or co-Büchi winning conditions. However, if each player has a Büchi winning condition, the problem becomes decidable in polynomial time.

For the proof of NP-hardness, it is essential that the number of players is unbounded. In applications, the number of players may be much smaller than the size of the game graph. So it is interesting to know the complexity of the problem if the number of players is small. A more fine-grained parameter is the *Hamming distance* of the two thresholds, i.e. the number of non-matching bits. Clearly, if the number of players is bounded, then the Hamming distance of the thresholds can be bounded as well, but even if the number of players is not bounded, the distance of the thresholds may be.

We show that, for games with parity winning conditions, NE retains the complexity of the classical parity game problem when restricted to payoff thresholds of bounded Hamming distance: In general, the complexity goes down to $UP \cap co-UP$, but for restricted classes of games we get membership in P. It follows that there exists an FPT algorithm for NE on these restricted

classes if one considers the Hamming distance of the payoff thresholds as the parameter.

Another interesting question is about the complexity of strategies realising a Nash equilibrium. We show that not only pure strategies (as opposed to mixed strategies, which allow for randomisation between actions) suffice, but also ones that require only a finite amount of memory. In particular, we show that in a multiplayer Streett game if there exists an arbitrary Nash equilibrium with a payoff between two (qualitative) thresholds, then there exists one that can be implemented by a finite automaton with $O(kdn)$ states, where k is the number of players, d is the maximal number of Streett pairs for each player and n is the size of the arena.

Related Work. Determining the complexity of Nash Equilibria has attracted much interest in recent years. Daskalakis & al. [7] showed that the problem of computing *some* Nash equilibrium of a game in strategic form is complete for the complexity class PPAD (a class of function problems which lies between FP and TFNP), and Chen and Deng [5] showed that this remains true for two-player games. More in the spirit of our work, Conitzer and Sandholm [6] showed that deciding whether there exists a Nash equilibrium in a two-player game in strategic form where player 1 receives payoff at least x and related decision problems are all NP-hard.

For infinite games, not much is known. Chatterjee & al. [4] showed that one can compute a Nash equilibrium of a multiplayer parity game in nondeterministic polynomial time. They also showed that it is NP-hard to determine whether a multiplayer game with reachability conditions has a Nash equilibrium where all players win.¹ However, translating a multiplayer reachability game into a game with (co-)Büchi or parity conditions typically entails an exponential blowup, so their results do not transfer to games with these winning conditions. In fact, it follows from our results that the problem is unlikely to be NP-hard in the case of multiplayer parity games, since we show that the problem is in $UP \cap \text{co-UP}$ for multiplayer parity games and even in P for multiplayer Büchi or co-Büchi games.

2 Preliminaries

The definition of an infinite (two-player zero-sum) game played on a finite directed graph easily generalises to the multiplayer setting. Formally, we define an *infinite multiplayer game* as a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, E, (\text{Win}_i)_{i \in \Pi})$ where

- Π is a finite set of *players*;
- (V, E) is a finite directed graph;
- $(V_i)_{i \in \Pi}$ is a partition of V ;

¹In fact, they consider concurrent games, but it is easy to see that their reduction can be modified to work for the kind of games we consider here.

- Win_i is a Borel set over V^ω for all $i \in \Pi$.

The structure $G = (V, (V_i)_{i \in \Pi}, E)$ is called the *arena* of \mathcal{G} , and Win_i is called the *winning condition* of player $i \in \Pi$. For the sake of simplicity, we assume that $vE := \{w \in V : (v, w) \in E\} \neq \emptyset$ for all $v \in V$, i.e. each vertex of G has at least one outgoing edge. We call \mathcal{G} a *zero-sum game* if the sets Win_i define a partition of V^ω . Thus if \mathcal{G} is an infinite two-player zero-sum game with players 0 and 1 it suffices to define V_0 and Win_0 , and we just write $\mathcal{G} = (V, V_0, E, \text{Win}_0)$.

A *play* or *history* of \mathcal{G} is an infinite or finite path in G , respectively. We say that a play π is *won* by player $i \in \Pi$ if $\pi \in \text{Win}_i$. The *payoff* of a play π of \mathcal{G} is the vector $\text{pay}(\pi) \in \{0, 1\}^\Pi$ defined by $\text{pay}(\pi)(i) = 1$ if π is won by player i .

A (*mixed*) *strategy* of player i in \mathcal{G} is a total function $\sigma : V^*V_i \rightarrow \mathcal{D}(V)$ assigning to each nonempty sequence xv of vertices ending in a vertex v of player i a (discrete) probability distribution over V such that $\sigma(xv)(w) > 0$ only if $(v, w) \in E$. We say that a play π of \mathcal{G} is *consistent* with a strategy σ of player i if $\sigma(\pi(0) \dots \pi(k))(\pi(k+1)) > 0$ for all $k < \omega$ with $\pi(k) \in V_i$. A (*mixed*) *strategy profile* of \mathcal{G} is a tuple $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ where σ_i is a strategy of player i in \mathcal{G} . Note that a strategy profile can be identified with a function $V^+ \rightarrow \mathcal{D}(V)$. Given a strategy profile $\bar{\sigma} = (\sigma_j)_{j \in \Pi}$ and a strategy τ of player i , we denote by $(\bar{\sigma}_{-i}, \tau)$ the strategy profile resulting from $\bar{\sigma}$ by replacing σ_i with τ .

A strategy σ of player i is called *pure* if for each $xv \in V^*V_i$ there exists $w \in vE$ with $\sigma(xv)(w) = 1$. Note that a pure strategy of player i can be identified with a function $\sigma : V^*V_i \rightarrow V$. Finally, a strategy profile $\bar{\sigma}$ is called *pure* if each one of the strategies σ_i is pure. A strategy σ of player i in \mathcal{G} is called *positional* if σ depends only on the current vertex, i.e. if $\sigma(xv) = \sigma(v)$ for all $xv \in V^*V_i$. A strategy profile $\bar{\sigma}$ of \mathcal{G} is called *positional* if each σ_i is positional.

It is sometimes convenient to designate an initial vertex $v_0 \in V$ of the game. We call the tuple (\mathcal{G}, v_0) an *initialised game*. A *play (history)* of (\mathcal{G}, v_0) is a play (history) of \mathcal{G} starting with v_0 . A strategy (strategy profile) of (\mathcal{G}, v_0) is just a strategy (strategy profile) of \mathcal{G} . Note that if $\bar{\sigma}$ is a pure strategy profile then there is a unique play of (\mathcal{G}, v_0) consistent with each σ_i , which we denote by $\langle \bar{\sigma} \rangle$.

Given a strategy profile $\bar{\sigma}$ and an initial vertex v_0 , the *probability* of a basic open set $v_0v_1 \dots v_k \cdot V^\omega$ is defined as the product of the probabilities $\bar{\sigma}(v_0 \dots v_{j-1})(v_j)$ for $j = 1, \dots, k$. It is a classical result of measure theory that this extends to a unique probability measure over Borel sets, which we denote by $\text{Pr}_{\bar{\sigma}}$. The *payoff* of a strategy profile $\bar{\sigma}$ is the vector $\text{pay}(\bar{\sigma}) \in [0, 1]^\Pi$ defined by $\text{pay}(\bar{\sigma})(i) = \text{Pr}_{\bar{\sigma}}(\text{Win}_i)$.

A strategy profile $\bar{\sigma}$ is called a *Nash equilibrium* of (\mathcal{G}, v_0) if $\text{Pr}_{(\bar{\sigma}_{-i}, \tau)}(\text{Win}_i) \leq \text{Pr}_{\bar{\sigma}}(\text{Win}_i)$ for each player i and each strategy τ of player i . Thus, in a Nash equilibrium no player can improve her payoff by (unilaterally)

switching to a different strategy. The following proposition rephrases the condition for the case that $\bar{\sigma}$ is pure. In this case, it suffices to check that no player can gain from switching to another pure strategy. This follows from the fact that pure strategy suffice to win one-player games.

Proposition 1. Let $\bar{\sigma}$ be a pure strategy profile of a game (\mathcal{G}, v_0) . The profile $\bar{\sigma}$ is Nash equilibrium iff $\langle \bar{\sigma}_{-i}, \tau \rangle \in \text{Win}_i \Rightarrow \langle \bar{\sigma} \rangle \in \text{Win}_i$ for each player i and each pure strategy τ of player i .

A strategy σ of player i is called *winning* if $\Pr_{(\bar{\tau}_{-i}, \sigma)}(\text{Win}_i) = 1$ for any strategy profile $\bar{\tau}$. For a game \mathcal{G} we call the set of all vertices $v \in V$ such that player i has a winning strategy for (\mathcal{G}, v) the *winning region of player i* , and a strategy of player i that is winning in (\mathcal{G}, v) for each vertex v in the winning region of player i an *optimal* strategy.

A celebrated theorem due to Martin [14] states that any two-player zero-sum game \mathcal{G} with a Borel set as its winning condition is *determined* by pure strategies, i.e. both players have pure optimal strategies, and the union of the winning regions is the set of all vertices.

Winning Conditions

We have introduced winning conditions as abstract Borel sets of infinite sequences of vertices. In verification winning conditions are usually *ω -regular sets*. Special cases are the following well-studied winning conditions:

- *Büchi* (given by $F \subseteq V$): the set of all $\alpha \in V^\omega$ such that $\alpha(k) \in F$ for infinitely many $k < \omega$.
- *co-Büchi* (given by $F \subseteq V$): the set of all $\alpha \in V^\omega$ such that $\alpha(k) \in F$ for all but finitely many $k < \omega$.
- *Parity* (given by a *priority function* $\Omega : V \rightarrow \omega$): the set of all $\alpha \in V^\omega$ such that the least number occurring infinitely often in $\Omega(\alpha)$ is even.
- *Streett* (given by a set Ω of *pairs* (L, R) where $L, R \subseteq V$): the set of all $\alpha \in V^\omega$ such that for all pairs $(L, R) \in \Omega$ with $\alpha(k) \in L$ for infinitely many $k < \omega$ it is the case that $\alpha(k) \in R$ for infinitely many $k < \omega$.

Note that any Büchi condition is a parity conditions with two priorities and that any parity condition is a Streett condition. In fact, the intersection of any two parity conditions is a Streett condition. Moreover, the complement of a Büchi condition is a co-Büchi condition and vice versa, whereas the class of parity conditions is closed under complementation. Finally, note that any Streett condition is *prefix independent*, i.e. for every $\alpha \in V^\omega$ and $x \in V^*$ it is the case that α satisfies the condition if and only if $x\alpha$ does.

We call a game \mathcal{G} a *multiplayer ω -regular, (co-)Büchi, parity or Streett game* if the winning condition of *each* player is of the respective type. This differs somehow from the usual convention for two-player zero-sum games where a Büchi, co-Büchi or Streett game is a game where the winning condition of the *first* player is a Büchi, co-Büchi or Streett condition, respectively.

3 Characterising Nash Equilibria

In this section we aim to characterise the existence of a Nash equilibrium with a qualitative constraint on the payoff. The characterisation works for any prefix-independent Borel winning condition, though we will only use it for games with (co-)Büchi, parity or Streett winning conditions in the following sections.

For the rest of this section, let (\mathcal{G}, v_0) be any k -player game with prefix-independent Borel winning conditions Win_i ; let W_i be the winning region of player i in \mathcal{G} , and let τ_i be a pure optimal strategy of player i . For each player $i \in \Pi$, we define $\text{Hit}_i := V^* \cdot W_i \cdot V^\omega$. Moreover, for $\bar{x}, \bar{y} \in \{0, 1\}^k$ let $\text{Plays}(\bar{x}, \bar{y})$ be the set of all plays π such that $\bar{x} \leq \text{pay}(\pi) \leq \bar{y}$.

Lemma 2. Let $\bar{\sigma}$ be a strategy profile such that $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$ for $\bar{x}, \bar{y} \in \{0, 1\}^k$. Then $\Pr_{\bar{\sigma}}(\text{Plays}(\bar{x}, \bar{y})) = 1$.

For each strategy σ of player i , we define a new strategy σ^* of player i by

$$\sigma^*(v_1 \dots v_k) = \begin{cases} \sigma(v_1 \dots v_k) & \text{if } v_j \notin W_i \text{ for all } 0 < j \leq k, \\ \tau_i(v_j \dots v_k) & \text{if } v_j \in W_i \text{ and } v_l \notin W_i \text{ for all } 0 < l < j. \end{cases}$$

Intuitively, σ^* is at least as good as σ , since σ^* behaves like σ as long as W_i is not visited and switches to τ_i after the first visit to W_i , which guarantees a win with probability 1. The exact gain in probability is given by the following lemma.

Lemma 3. For any strategy profile $\bar{\sigma}$ and each player $i \in \Pi$, $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i) = \Pr_{\bar{\sigma}}(\text{Win}_i) + \Pr_{\bar{\sigma}}(\overline{\text{Win}_i} \cap \text{Hit}_i)$.

The next lemma allows to infer the existence of a Nash equilibrium from the existence of a certain play. The proof uses so-called *threat strategies* (also known as *trigger strategies*), which are the basis of the *folk theorems* in the theory of repeated games (cf. [1] and [17, Chapter 8]).

Lemma 4. If there exists a play π such that $\pi \in \text{Win}_i \cup \overline{\text{Hit}_i}$ for each player $i \in \Pi$, then there exists a pure Nash equilibrium $\bar{\sigma}$ with $\pi = \langle \bar{\sigma} \rangle$.

Proof. Let π be a play of (\mathcal{G}, v_0) such that $\pi \in \text{Win}_i \cup \overline{\text{Hit}_i}$ for each player $i \in \Pi$. Moreover, let $\tau_{\Pi \setminus \{j\}}$ be an optimal pure strategy of the coalition $\Pi \setminus \{j\}$ in the two-player zero-sum game \mathcal{G}_j where player j plays with her winning condition against all other players in \mathcal{G} ; let $\tau_{i,j}$ be the corresponding strategy of player $i \neq j$ in \mathcal{G} (i.e. $\tau_{i,j}(xv) = \tau_{\Pi \setminus \{j\}}(xv)$ for each $v \in V_i$), and for each player i let $\tau_{i,i}$ be an arbitrary pure strategy. For each player i , we define a new pure strategy σ_i in \mathcal{G} as follows:

$$\sigma_i(xv) = \begin{cases} \pi(k+1) & \text{if } xv = \pi(0) \dots \pi(k) \prec \pi, \\ \tau_{i,j}(x_2v) & \text{otherwise,} \end{cases}$$

where, in the latter case, $x = x_1x_2$ such that x_1 is the longest prefix of x still being a prefix of π , and j is the player whose turn it was after that prefix (i.e. x_1 ends in V_j), where $j = i$ if $x_1 = \varepsilon$.

Clearly, we have $\pi = \langle \bar{\sigma} \rangle$. We claim that $\bar{\sigma}$ is a Nash equilibrium. Towards a contradiction, assume that some player $i \in \Pi$ with $\pi \notin \text{Win}_i$ can improve her payoff by playing according to some (w.l.o.g. pure) strategy τ instead of σ_i . Then there exists $k < \omega$ such that $\tau(\pi(k)) \neq \sigma_i(\pi(k))$, and consequently from this point onwards $\langle \bar{\sigma}_{-i}, \tau \rangle$ is consistent with τ_{-i} , the optimal strategy of the coalition $\Pi \setminus \{i\}$ in \mathcal{G}_i . Hence, it must be the case that τ_{-i} is not winning from $\pi(k)$. By determinacy, this implies that $\pi(k) \in W_i$, so $\pi \in \text{Hit}_i$, a contradiction to the assumption that $\pi \in \text{Win}_i \cup \overline{\text{Hit}}_i$. Q.E.D.

Now we have all the ingredients to prove the following proposition, which characterises the existence of a Nash equilibrium with a payoff between two thresholds $\bar{x}, \bar{y} \in \{0, 1\}^k$.

Proposition 5. Let $\bar{x}, \bar{y} \in \{0, 1\}^k$. The following statements are equivalent:

1. There exists a Nash equilibrium $\bar{\sigma}$ with payoff $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$;
2. There exists a strategy profile $\bar{\sigma}$ with payoff $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$ such that $\Pr_{\bar{\sigma}}(\text{Win}_i \cup \overline{\text{Hit}}_i) = 1$ for each player $i \in \Pi$;
3. There exists a play π with $\bar{x} \leq \text{pay}(\pi) \leq \bar{y}$ such that $\pi \in \text{Win}_i \cup \overline{\text{Hit}}_i$ for each player $i \in \Pi$;
4. There exists a pure Nash equilibrium $\bar{\sigma}$ with payoff $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$.

Proof. (1. \Rightarrow 2.) Let $\bar{\sigma}$ be a Nash equilibrium with $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$. We claim that we have $\Pr_{\bar{\sigma}}(\text{Win}_i \cup \overline{\text{Hit}}_i) = 1$ for each player $i \in \Pi$. Otherwise, $\Pr_{\bar{\sigma}}(\overline{\text{Win}}_i \cap \text{Hit}_i) > 0$ for some player $i \in \Pi$, and by Lemma 3 we would have $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i) > \Pr_{\bar{\sigma}}(\text{Win}_i)$, so player i could improve her payoff by switching from σ to σ^* , a contradiction to the fact that $\bar{\sigma}$ is a Nash equilibrium.

(2. \Rightarrow 3.) Assume that there exists a strategy profile $\bar{\sigma}$ with $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$ such that $\Pr_{\bar{\sigma}}(\text{Win}_i \cup \overline{\text{Hit}}_i) = 1$ for each player $i \in \Pi$. Moreover, by Lemma 2 we have $\Pr_{\bar{\sigma}}(\text{Plays}(\bar{x}, \bar{y})) = 1$. Hence, by elementary probability theory, also $\Pr_{\bar{\sigma}}(\text{Plays}(\bar{x}, \bar{y}) \cap \bigcap_{i \in \Pi} (\text{Win}_i \cup \overline{\text{Hit}}_i)) = 1$. But then, in particular, there must exist a play $\pi \in \text{Plays}(\bar{x}, \bar{y})$ such that $\pi \in \text{Win}_i \cup \overline{\text{Hit}}_i$ for each player i .

(3. \Rightarrow 4.) The claim follows from Lemma 4.

(4. \Rightarrow 1.) Trivial. Q.E.D.

As a corollary we can infer that randomised strategies are not more powerful than pure ones as far as the existence of a Nash equilibrium with a qualitative constraint on the payoff is concerned.

Corollary 6. Let $\bar{x}, \bar{y} \in \{0, 1\}^k$. There exists a Nash equilibrium $\bar{\sigma}$ with $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$ iff there exists a pure Nash equilibrium $\bar{\sigma}$ with $\bar{x} \leq \text{pay}(\bar{\sigma}) \leq \bar{y}$.

Remark 7. Corollary 6 fails if the thresholds \bar{x} and \bar{y} are not bitvectors: One can easily construct an example of a two-player game where there is a Nash

equilibrium with payoff $(1, x)$ for each real number $x \in [0, 1]$, whereas any pure Nash equilibrium has payoff $(1, 0)$ or $(1, 1)$.

4 Computational Complexity

Previous research on algorithms for finding Nash equilibria in infinite games has focused on computing *some* Nash equilibrium [4]. However, a game may have several Nash equilibria with different payoffs, so one might not be interested in *any* Nash equilibrium but in one that fulfils certain requirements. For example, one might look for a Nash equilibrium where certain players win while certain other players lose. Or one might look for a *maximal* Nash equilibrium, i.e. a Nash equilibrium such that there is no Nash equilibrium with a higher payoff. This idea is captured by the following decision problem, which we call NE:

Given a multiplayer game (\mathcal{G}, v_0) with ω -regular winning conditions and thresholds $\bar{x}, \bar{y} \in \{0, 1\}^{\Pi}$, decide whether there exists a Nash equilibrium of (\mathcal{G}, v_0) with a payoff $\geq \bar{x}$ and $\leq \bar{y}$.

4.1 Upper Bounds

Our first result on the complexity of NE is that there is a nondeterministic polynomial-time algorithm to decide NE for multiplayer Streett games. This may surprise the reader, since the problem of deciding whether player 0 has a winning strategy in a two-player zero-sum game with a Streett winning condition is, in fact, co-NP-complete [8]. However, recall that, according to our definition, a two-player zero-sum game where one player has a Streett winning condition is, in general, not a Streett game, since we require both players to have the same type of winning condition.

Theorem 8. For multiplayer Streett games, NE is in NP.

Proof. To decide whether there exists a Nash equilibrium of a multiplayer Streett game (\mathcal{G}, v_0) with payoff $\geq \bar{x}$ and $\leq \bar{y}$, a nondeterministic algorithm can guess a payoff $\bar{z} \in \{0, 1\}^{\Pi}$ with $\bar{x} \leq \bar{z} \leq \bar{y}$ and, for each player $i \in \Pi$ with $z_i = 0$, a set Z_i together with a pure positional strategy τ_{-i} for the coalition $\Pi \setminus \{i\}$. Finally, the algorithm guesses a strongly connected set $U \subseteq \bigcap_{z_i=0} (V \setminus Z_i)$ that is reachable from v_0 inside $\bigcap_{z_i=0} (V \setminus Z_i)$. If there is no such set, the algorithm rejects immediately. In the next step, the algorithm checks for each i with $z_i = 0$ whether τ_{-i} is winning on $V \setminus Z_i$ as well as whether U fulfils the winning condition (i.e. $U \cap L \neq \emptyset \Rightarrow U \cap R \neq \emptyset$ for each Streett pair (L, R)) of each player i with $z_i = 1$ and whether U does not fulfil the winning condition of each player i with $z_i = 0$. If all checks are successful, the algorithm accepts; otherwise it rejects.

The correctness of the algorithm follows from Proposition 5: If there is a positional winning strategy of the coalition $\Pi \setminus \{i\}$ on $V \setminus Z_i$, then Z_i must

be a superset of the winning region of player i . So we can build a play with payoff \bar{z} staying outside of the winning region of each player i with $z_i = 0$ by going from v_0 to U and visiting each vertex of U again and again, which is possible because U is strongly connected. On the other hand, if there is a play π with payoff \bar{z} and staying inside $V \setminus W_i$ for each player i with $z_i = 0$, then the checks will succeed for the guesses $Z_i = W_i$, τ_i an optimal pure positional strategy of the coalition $\Pi \setminus \{i\}$ in \mathcal{G} and $U = \text{Inf}(\pi)$, since pure positional strategies suffice for player 2 to win a two-player zero-sum game with a Streett winning condition for player 1 [13].

It remains to show that the algorithm runs in polynomial time. The only critical steps are the checks whether a pure positional strategy of the coalition $\Pi \setminus \{i\}$ is winning on $V \setminus Z_i$. In order to check this, the algorithm can construct the one-player Streett game that arises from fixing the transitions taken by the positional strategy and check whether there exists a winning play for the remaining player from a vertex outside Z_i . Emerson and Lei [9] showed that there exists a polynomial-time algorithm for the latter problem. Q.E.D.

As an immediate consequence of Theorem 8, we get that for multiplayer parity games the problem NE is in NP, as well. However, in many cases, we can do better: For two payoff vectors $\bar{x}, \bar{y} \in \{0, 1\}^{\Pi}$, let $\text{dist}(\bar{x}, \bar{y})$ be the *Hamming distance* of \bar{x} and \bar{y} , i.e. the number $\sum_{i \in \Pi} |y_i - x_i|$ of nonmatching bits. We show that if $\text{dist}(\bar{x}, \bar{y})$ is bounded then NE retains the complexity of the parity game problem, which is known to be in $\text{UP} \cap \text{co-UP}$ [12].

Theorem 9. For multiplayer parity games and bounded $\text{dist}(\bar{x}, \bar{y})$, NE is in $\text{UP} \cap \text{co-UP}$.

Proof. In the following, let us assume that $\text{dist}(\bar{x}, \bar{y})$ is bounded. A UP algorithm for NE works as follows: On input (\mathcal{G}, v_0) the algorithm starts by guessing the winning region W_i of each player. Then, for each vertex v and each player i , the guess whether $v \in W_i$ or $v \notin W_i$ is verified by running the UP algorithm for the respective problem. If one guess was incorrect, the algorithm rejects immediately. Otherwise, the algorithm checks for each payoff $\bar{z} \in \{0, 1\}^{\Pi}$ with $\bar{x} \leq \bar{z} \leq \bar{y}$ whether there exists a winning play from v_0 in the one-player Streett game with winning condition $\bigcap_{z_i=1} \text{Win}_i \cap \bigcap_{z_i=0} \overline{\text{Win}_i}$ on the arena $G \upharpoonright \bigcap_{z_i=0} (V \setminus W_i)$. The algorithm accepts if this is the case for at least one such payoff. Analogously, a UP algorithm for the complement of NE accepts if there is *no* winning play in the same game from v_0 for each $\bar{x} \leq \bar{z} \leq \bar{y}$.

It is easy to see that both algorithms are indeed UP algorithm. The correctness of the two algorithms follows again from Proposition 5 with a similar reasoning as in the proof of Theorem 8. Q.E.D.

It is a major open problem whether winning regions of two-player zero-sum parity games can be computed in polynomial time, in general. This

would allow us to decide NE for multiplayer parity games and bounded $\text{dist}(\bar{x}, \bar{y})$ in polynomial time, as well.

Though it could not yet be shown that the parity game problem is in P, researchers have succeeded in identifying structural subclasses of parity games that allow to solve the game in polynomial time. First of all, it is well known that the parity game problem is in polynomial time for games with a bounded number of priorities (cf. [19]). Other classes of parity games with this property were shown to be any class of games played on graphs of bounded DAG- or Kelly width [2, 15, 11] (and thus also on graphs of bounded tree width or bounded entanglement), and on graphs of bounded clique width [16]. We give a general theorem that allows to transfer these results to the multiplayer case. Formally, for any class \mathcal{C} of two-player zero-sum parity games, let \mathcal{C}^* be the class of all multiplayer parity games of the form $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, E, (\Omega_i)_{i \in \Pi})$ where, for each player $i \in \Pi$, the two-player zero-sum game $\mathcal{G}_i = (V, V_i, E, \Omega_i)$ is in \mathcal{C} .

Theorem 10. Let \mathcal{C} be a class of two-player zero-sum parity games such that the parity game problem is decidable in P for games in \mathcal{C} . Then NE is in P for games in \mathcal{C}^* and bounded $\text{dist}(\bar{x}, \bar{y})$.

Proof. Consider the algorithm given in the proof of Theorem 9. The winning region W_i of player i in \mathcal{G} is precisely the winning region of player 0 in the game \mathcal{G}_i . So, if $\mathcal{G}_i \in \mathcal{C}$, we can compute the set W_i in polynomial time, and there is no need to guess this set. So we can make the algorithm deterministic while keeping its running time polynomial. Q.E.D.

As a corollary, we can deduce that there exists an FPT algorithm for NE when restricted to games in one of the aforementioned classes w.r.t. the parameter $\text{dist}(\bar{x}, \bar{y})$. In particular, if the general parity game problem is in P, then NE for multiplayer parity games is fixed-parameter tractable.

Corollary 11. Let \mathcal{C} be a class of two-player zero-sum parity games such that the parity game problem is decidable in P for games in \mathcal{C} . Then NE for games in \mathcal{C}^* admits an FPT algorithm w.r.t. the parameter $\text{dist}(\bar{x}, \bar{y})$.

Proof. To decide whether there is a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$, it suffices to check for each of the $2^{\text{dist}(\bar{x}, \bar{y})}$ many payoffs $\bar{z} \in \{0, 1\}^{\Pi}$ with $\bar{x} \leq \bar{z} \leq \bar{y}$ whether there exists a Nash equilibrium with payoff \bar{z} . By Theorem 10, each of these checks can be done in polynomial time for games in \mathcal{C}^* . Q.E.D.

The natural question at this point is whether NE is actually decidable in polynomial time. In the next section, we will see that this is probably not the case. However, we claim that there exists a polynomial-time algorithm for NE if all winning conditions are Büchi winning conditions. Towards this, we describe a polynomial-time algorithm that computes, given a multiplayer

Büchi game \mathcal{G} and payoff thresholds $\bar{x}, \bar{y} \in \{0, 1\}^\Pi$, the set of vertices from where there exists a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$.

The algorithm is similar to the algorithm by Emerson and Lei [9] for deciding one-player Streett games and works as follows: By Proposition 5, the game (\mathcal{G}, v) has a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$ if and only if there exists a play staying outside the winning region W_i of each player i with $\text{pay}(\pi)(i) = 0$. Clearly, this is the case if and only if there exists a payoff $\bar{z} \in \{0, 1\}^\Pi$ and a strongly connected set $U \subseteq \bigcap_{z_i=0} (V \setminus W_i)$ reachable from v inside $\bigcap_{z_i=0} (V \setminus W_i)$ such that $U \cap F_i \neq \emptyset$ iff $z_i = 1$. The essential part of the algorithm is the procedure SolveSubgame. On input X its task is to find any such set contained in X .

As we are only interested in strongly connected sets that do not intersect with F_i for each player i such that $y_i = 0$, SolveSubgame is firstly called for the subgraph of G that results from G by removing all these sets F_i .

Algorithm 1. Computing the set of vertices from where there is a Nash equilibrium in a multiplayer Büchi game with a payoff $\geq x$ and $\leq y$.

input multiplayer Büchi game $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, E, (F_i)_{i \in \Pi})$, $\bar{x}, \bar{y} \in \{0, 1\}^\Pi$
for each $i \in \Pi$ **do**

Compute the winning region W_i of player i in \mathcal{G}

end for

$X := \bigcap_{y_i=0} (V \setminus F_i)$

return SolveSubgame(X)

procedure SolveSubgame(X)

$Z := \emptyset$

Decompose $G \upharpoonright X$ into strongly connected components (SCC's)

for each non-trivial SCC C of $G \upharpoonright X$ **do**

$L := \{i \in \Pi : C \cap F_i = \emptyset\}$

if $i \notin L$ for all i with $x_i = 1$ **then**

$Y := C \cap \bigcap_{i \in L} (V \setminus W_i)$

if $Y = C$ **then**

$Z := Z \cup \{v \in V : C \text{ reachable from } v \text{ in } G \upharpoonright \bigcap_{i \in L} (V \setminus W_i)\}$

else

$Z := Z \cup \text{SolveSubgame}(Y)$

end if

end if

end for

return Z

end procedure

Theorem 12. NE is decidable in polynomial time for multiplayer Büchi games.

Proof. We claim that Algorithm 1 computes precisely the set of vertices from where there is a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$. Note that the

procedure `SolveSubgame` calls itself at most $|X| - 1$ times on input X . So the procedure is called at most $|V|$ times in total. As furthermore any graph can be decomposed into its SCCs in linear time and winning regions of Büchi games can be computed in polynomial time, this implies that the algorithm runs in polynomial time. Q.E.D.

Remark 13. In fact, by combining the proofs of Theorem 12 and Theorem 10, one can show that NE is decidable in polynomial time for games with an arbitrary number of Büchi winning conditions and a bounded number of co-Büchi winning conditions (or even a bounded number of parity winning conditions with a bounded number of priorities).

4.2 Lower Bounds

The question remains whether, in general, NE is NP-hard. We answer this question affirmatively by showing that NE is not only NP-hard for two-player Streett games, but also for multiplayer co-Büchi games with an unbounded number of players even if we only want to know whether there is an equilibrium where a certain player wins. Note that, in the light of Theorem 9, it is very unlikely that NE is NP-hard for parity games when restricted to a bounded number of players.

Theorem 14. NE is NP-hard for two-player Streett games.

Proof. The proof is a variant of the proof for NP-hardness of the problem of deciding whether player 2 has a winning strategy in a two-player zero-sum game with a Streett winning condition for player 1 [8] and by a reduction from SAT.

Given a Boolean formula φ in conjunctive normal form, we construct a two-player Streett game \mathcal{G}_φ as follows: For each clause C the game \mathcal{G}_φ has a vertex C , which is controlled by player 1, and for each literal X or $\neg X$ occurring in φ there is a vertex X or $\neg X$, respectively, which is controlled by player 2. There are edges from a clause to each literal that occurs in this clause, and from a literal to each clause occurring in φ . Player 1 wins every play of the game whereas player 2 wins if for each variable X either neither X nor $\neg X$ or both X and $\neg X$ have been visited infinitely often (clearly a Streett condition).

Obviously, \mathcal{G}_φ can be constructed from φ in polynomial time. We claim that φ is satisfiable if and only if (\mathcal{G}_φ, C) has a Nash equilibrium where player 2 loses (with C being an arbitrary clause). Q.E.D.

Since $\text{dist}(\bar{x}, \bar{y})$ is always bounded by the number of players, it follows from Theorem 14 that NE is not fixed-parameter tractable for Streett games w.r.t. this parameter.

Theorem 15. NE is NP-hard for multiplayer co-Büchi games, even with the thresholds $\bar{x} = (1, 0, \dots, 0)$ and $\bar{y} = (1, \dots, 1)$.

Proof. Again, the proof is by a reduction from SAT. Given a Boolean formula $\varphi = C_1 \wedge \dots \wedge C_m$ in CNF over variables X_1, \dots, X_n , we build a game \mathcal{G}_φ played by players $0, 1, \dots, n$ as follows. \mathcal{G}_φ has vertices C_1, \dots, C_m controlled by player 0, and for each clause C and each literal X_i or $\neg X_i$ that occurs in C , a vertex (C, X_i) or $(C, \neg X_i)$, respectively, controlled by player i . Additionally, there is a sink vertex \perp . There are edges from a clause C_j to each vertex (C_j, L) such that L occurs as a literal in C_j and from there to $C_{(j \bmod m)+1}$. Additionally, there is an edge from each vertex $(C, \neg X_i)$ to the sink vertex \perp . As \perp is a sink vertex, the only edge leaving \perp leads to \perp itself. The arena of \mathcal{G}_φ is schematically depicted in Figure 1. The co-Büchi winning conditions are as follows:

- Player 0 wins if the sink vertex is visited only finitely often (i.e. never);
- Player $i \in \{1, \dots, n\}$ wins if each vertex (C, X_i) is visited only finitely often.

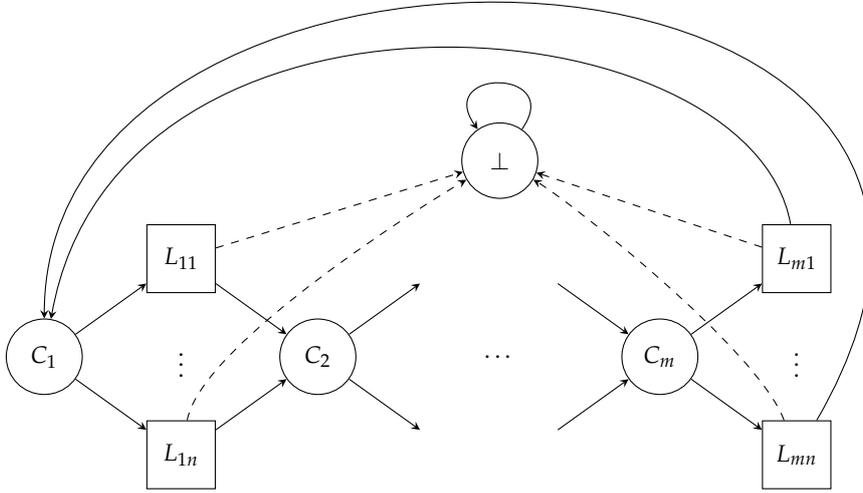


Figure 1. The arena of the game \mathcal{G}_φ .

Clearly, \mathcal{G}_φ can be constructed from φ in polynomial time. We claim that φ is satisfiable if and only if $(\mathcal{G}_\varphi, C_1)$ has a Nash equilibrium where player 0 wins. Q.E.D.

5 Strategy Complexity

Recall that we have already shown that pure strategies suffice to realise any Nash equilibrium with a qualitative constraint on the payoff (Corollary 6). In this section we aim to analyse the memory requirements of these pure strategies.

A *finite-state transducer* is a tuple $\mathcal{A} = (Q, \Sigma, \Gamma, q_0, \delta, \tau)$ where Q is a finite set of states, Σ and Γ finite sets of letters, $q_0 \in Q$ the initial state, $\delta : Q \times \Sigma \rightarrow Q$ the transition function, and $\tau : Q \times \Sigma \rightarrow \Gamma$ the output

function. The function δ is naturally extended to a function $\delta^* : \Sigma^* \rightarrow Q$ by setting $\delta^*(\varepsilon) = q_0$ and $\delta^*(xa) = \delta(\delta^*(x), a)$. The transducer \mathcal{A} computes the function $f : \Sigma^+ \rightarrow \Gamma$ defined by $f(xa) = \tau(\delta^*(x), a)$. So, if $\Sigma = \Gamma = V$ and $\tau(q, v) \in vE$ for each $q \in Q$ and $v \in V$ we can interpret f as a pure strategy profile of a game played on $G = (V, (V_i)_{i \in \Pi}, E)$. We call a pure strategy profile $\bar{\sigma}$ a *finite-state strategy profile* if there exists a finite-state transducer that computes $\bar{\sigma}$.

Theorem 16. Let (\mathcal{G}, v_0) be a multiplayer Streett game with k players, at most d Streett pairs for each player and n vertices, and let $\bar{x}, \bar{y} \in \{0, 1\}^k$. If there exists a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$, then there exists a pure Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$ that can be computed by a finite automaton with $O(kdn)$ (or, alternatively, $O(n^2)$) states.

Proof. Assume that there exists a Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$. By Proposition 5, this implies that there exists a play π with payoff $\bar{x} \leq \text{pay}(\pi) \leq \bar{y}$ such that π avoids the winning region W_i of each player i with $\pi \notin \text{Win}_i$. Now consider the set U of vertices visited infinitely often in π . In fact, it suffices to attract the token to U (and thereby avoiding each W_i with $\pi \notin \text{Win}_i$) and then, while staying inside U , to visit again and again for each player i and for each Streett pair (L, R) in the condition for player i a vertex $u \in U \cap L$ if $U \cap L \neq \emptyset$ and a vertex $u' \in U \cap R$ if $U \cap R \neq \emptyset$. The resulting play π' has the same payoff as π and will still satisfy the requirement that π' avoids each W_i with $\pi' \notin \text{Win}_i$. It is easy to see that there exists a pure strategy profile $\bar{\tau}$ with $\pi' = \langle \bar{\tau} \rangle$ that can be computed by an automaton with $O(kd)$ states. Alternatively, using $O(n)$ states, one can guarantee to visit every vertex in U (and only vertices in U) again and again.

Now consider the equilibrium $\bar{\sigma}$ constructed from π' in the proof of Lemma 4. Since player 1 has an optimal pure positional strategy in any two-player zero-sum game where player 0 has a Streett winning condition [13], the strategies $\tau_{\Pi \setminus \{j\}}$ defined there can be assumed to be positional. For each one of the states needed to realise $\bar{\tau}$, an automaton that implements $\bar{\sigma}$ needs one state for each vertex to detect a deviation. Additionally, it needs $\min(k, n)$ more states to remember which one of the strategies $\tau_{\Pi \setminus \{j\}}$ it executes after a deviation. So $\bar{\sigma}$ can be computed by an automaton with $O(kdn)$ (or, alternatively, $O(n^2)$) states. Q.E.D.

With a small modification in the proof, one can show that for multiplayer parity games one can get rid of the factor d , so in this case $O(kn)$ states suffice. Finally, since any game with ω -regular winning conditions can be reduced to a multiplayer parity game using finite memory [18], we can conclude that pure finite-state strategies suffice to realise any Nash equilibrium with a qualitative constraint on the payoff in these games.

Corollary 17. Let (\mathcal{G}, v_0) be a game with ω -regular winning conditions, and let $\bar{x}, \bar{y} \in \{0, 1\}^{\Pi}$. If there exists a Nash equilibrium with a payoff $\geq \bar{x}$ and

$\leq \bar{y}$, then there exists a pure finite-state Nash equilibrium with a payoff $\geq \bar{x}$ and $\leq \bar{y}$.

6 Conclusion

We have analysed the complexity of Nash equilibria in infinite multiplayer games with (co-)Büchi, parity and Streett objectives. We remark that with the same idea as in the proof of Theorem 8 one can show that NE is in P^{NP} (in fact $P^{NP[\log]}$) for games with only Rabin or mixed Rabin and Streett winning conditions and in PSPACE for games with Muller winning conditions.

Apart from studying other winning conditions, three obvious directions for further research come to mind: Firstly, it would be interesting to know the complexity of other solution concepts in the context of infinite games. For *subgame perfect equilibria*, preliminary results were obtained in [18]. Secondly, one could consider more general game models like *stochastic* or *concurrent* games. Thirdly, there is the *quantitative* version of NE, where the thresholds contain arbitrary probabilities rather than just 0 and 1. In fact, this problem arises naturally in the context of stochastic games.

Acknowledgements. I would like to thank an anonymous reviewer for clarifying the notion of a threat/trigger strategy and for pointing out [1].

References

1. R. J. Aumann. Survey of repeated games. In *Essays in Game Theory and Mathematical Economics in Honor of Oskar Morgenstern*, pp. 11–42. Bibliographisches Institut Mannheim/Wien/Zürich, 1981.
2. D. Berwanger, A. Dawar, P. Hunter & S. Kreutzer. DAG-width and parity games. In *Proceedings of the 23rd Symposium on Theoretical Aspects of Computer Science, STACS 2006*, vol. 3884 of LNCS, pp. 524–436. Springer-Verlag, 2006.
3. K. Chatterjee, T. A. Henzinger & M. Jurdziński. Games with secure equilibria. In *Proceedings of the 19th Annual Symposium on Logic in Computer Science, LICS 2004*, pp. 160–169. IEEE Computer Society Press, 2004.
4. K. Chatterjee, R. Majumdar & M. Jurdziński. On Nash equilibria in stochastic games. In *Proceedings of the 13th Annual Conference of the European Association for Computer Science Logic, CSL 2004*, vol. 3210 of LNCS, pp. 26–40. Springer-Verlag, 2004.
5. X. Chen & X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2006*, pp. 261–272. IEEE Computer Society Press, 2006.
6. V. Conitzer & T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence, IJCAI 2003*, pp. 765–771. Morgan Kaufmann, 2003.
7. C. Daskalakis, P. W. Goldberg & C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, STOC 2006*, pp. 71–78. ACM Press, 2006.

8. E. A. Emerson & C. S. Jutla. The complexity of tree automata and logics of programs (extended abstract). In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science, FoCS '88*, pp. 328–337. IEEE Computer Society Press, 1988.
9. E. A. Emerson & C.-L. Lei. Modalities for model checking: Branching time strikes back. In *Conference Record of the 12th Annual ACM Symposium on Principles of Programming Languages, POPL '85*, pp. 84–96. ACM Press, 1985.
10. D. Gale & F. M. Stewart. Infinite games with perfect information. In *Contributions to the Theory of Games II*, vol. 28 of *Annals of Mathematical Studies*, pp. 245–266. Princeton University Press, 1953.
11. P. Hunter & S. Kreutzer. Digraph measures: Kelly decompositions, games, and orderings. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms, SODA 2007*, pp. 637–644. ACM Press, 2007.
12. M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
13. N. Klarlund. Progress measures, immediate determinacy, and a subset construction for tree automata. In *Proceedings of the 7th Annual IEEE Symposium on Logic in Computer Science, LICS '92*, pp. 382–393. IEEE Computer Society Press, 1992.
14. D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
15. J. Obdržálek. DAG-width – connectivity measure for directed graphs. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms, SODA 2006*, pp. 814–821. ACM Press, 2006.
16. J. Obdržálek. Clique-width and parity games. In *Proceedings of the 16th Annual Conference of the European Association for Computer Science Logic, CSL 2007*, vol. 4646 of *LNCS*, pp. 54–68. Springer-Verlag, 2007.
17. M. J. Osborne & A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
18. M. Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2006*, vol. 4337 of *LNCS*, pp. 212–223. Springer-Verlag, 2006.
19. W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1–2):135–183, 1998.

Appendix

Proof of Proposition 1

(\Rightarrow) Assume that $\bar{\sigma}$ is a Nash equilibrium, and let τ be any pure strategy of player i with $\langle \bar{\sigma}_{-i}, \tau \rangle \in \text{Win}_i$. Since $\bar{\sigma}$ is a Nash equilibrium, we have $1 = \Pr_{(\bar{\sigma}_{-i}, \tau)}(\text{Win}_i) \leq \Pr_{\bar{\sigma}}(\text{Win}_i)$. Hence, also $\Pr_{\bar{\sigma}}(\text{Win}_i) = 1$, i.e. $\langle \bar{\sigma} \rangle \in \text{Win}_i$.

(\Leftarrow) Assume that $\bar{\sigma}$ is not a Nash equilibrium. So there exists a strategy τ' of some player i such that $\Pr_{(\bar{\sigma}_{-i}, \tau')}(\text{Win}_i) > \Pr_{\bar{\sigma}}(\text{Win}_i) = 0$. In particular, there must exist a play $\pi \in \text{Win}_i$ that is consistent with each σ_j for $j \neq i$. But then we can find a pure strategy τ of player i with $\pi = \langle \bar{\sigma}_{-i}, \tau \rangle$. So there exists a pure strategy τ of player i such that $\langle \bar{\sigma}_{-i}, \tau \rangle \in \text{Win}_i$ but $\langle \bar{\sigma} \rangle \notin \text{Win}_i$. Q.E.D.

Proof of Lemma 2

We show that we have $\Pr_{\bar{\sigma}}(\{\pi \in V^\omega : x_i \leq \text{pay}(\pi)(i) \leq y_i\}) = 1$ for each player i . By elementary probability theory, this implies that $\Pr_{\bar{\sigma}}(\text{Plays}(\bar{x}, \bar{y})) = 1$. Hence, let $i \in \Pi$. If $x_i = 0$ and $y_i = 1$, the claim is trivial. If $x_i = y_i = 1$, then $\Pr_{\bar{\sigma}}(\{\pi \in V^\omega : x_i \leq \text{pay}(\pi)(i) \leq y_i\}) = \Pr_{\bar{\sigma}}(\text{Win}_i) = 1$. Finally, if $x_i = y_i = 0$, then $\Pr_{\bar{\sigma}}(\{\pi \in V^\omega : x_i \leq \text{pay}(\pi)(i) \leq y_i\}) = \Pr_{\bar{\sigma}}(\overline{\text{Win}_i}) = 1$. Q.E.D.

Proof of Lemma 3

To prove the lemma, we need to establish the following two claims, which relate the probability measures $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}$ and $\Pr_{\bar{\sigma}}$.

Claim. $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i \cap \text{Hit}_i) = \Pr_{\bar{\sigma}}(\text{Hit}_i)$ for any strategy profile $\bar{\sigma}$ and each player $i \in \Pi$.

Proof. First, note that $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i \cap \text{Hit}_i) = \Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Hit}_i)$, since every play in Hit_i that is consistent with σ_i^* is won by player i . But $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Hit}_i) = \Pr_{\bar{\sigma}}(\text{Hit}_i)$, since before hitting W_i , the strategies σ_i and σ_i^* behave exactly the same.

Claim. $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i \cap \overline{\text{Hit}_i}) = \Pr_{\bar{\sigma}}(\text{Win}_i \cap \overline{\text{Hit}_i})$ for any strategy profile $\bar{\sigma}$ and each player $i \in \Pi$.

Proof. Note that $\bar{\sigma}$ and $(\bar{\sigma}_{-i}, \sigma_i^*)$ are equivalent on the set $(V \setminus W_i)^*$, which is prefix-closed. Since every prefix of $\text{Win}_i \cap \overline{\text{Hit}_i}$ lies in $(V \setminus W_i)^*$, this implies the claim.

Now, using the two previous claims, we can setup the following chain of equations:

$$\begin{aligned}
 & \Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i) \\
 &= \Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i \cap \text{Hit}_i) + \Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i \cap \overline{\text{Hit}_i}) \\
 &= \Pr_{\bar{\sigma}}(\text{Hit}_i) + \Pr_{\bar{\sigma}}(\text{Win}_i \cap \overline{\text{Hit}_i}) \\
 &= \Pr_{\bar{\sigma}}(\text{Hit}_i) + \Pr_{\bar{\sigma}}(\text{Win}_i) - \Pr_{\bar{\sigma}}(\text{Win}_i \cap \text{Hit}_i)
 \end{aligned}$$

$$= \Pr_{\bar{\sigma}}(\text{Win}_i) + \Pr_{\bar{\sigma}}(\overline{\text{Win}_i} \cap \text{Hit}_i)$$

So we have $\Pr_{(\bar{\sigma}_{-i}, \sigma_i^*)}(\text{Win}_i) = \Pr_{\bar{\sigma}}(\text{Win}_i) + \Pr_{\bar{\sigma}}(\overline{\text{Win}_i} \cap \text{Hit}_i)$, which proves the lemma. Q.E.D.

Proof of Theorem 12 (continued)

It remains to show that Algorithm 1 is correct. Let Z be the set of vertices output by the algorithm. We claim that Z is precisely the set of vertices from where there exists a Nash equilibrium with payoff $\geq \bar{x}$ and $\leq \bar{y}$.

(\subseteq) Assume $v \in Z$. So, in some call of the procedure `SolveSubgame`, say with vertex set X , there is a non-trivial SCC C of $G \upharpoonright X$ reachable from v and contained in $\bigcap_{C \cap F_i = \emptyset} (V \setminus W_i)$, where $C \cap F_i \neq \emptyset$ for each $i \in \Pi$ with $x_i = 1$. Let $\bar{z} \in \{0, 1\}^{\Pi}$ be defined by $z_i = 1$ iff $C \cap F_i \neq \emptyset$. So C is, in fact, reachable from v and contained in $\bigcap_{z_i=0} (V \setminus W_i)$. By Proposition 5, this implies that (\mathcal{G}, v) has a Nash equilibrium with payoff \bar{z} . Moreover, since the algorithm maintains the invariant that $X \cap F_i = \emptyset$ for each $i \in \Pi$ with $y_i = 0$ and $C \cap F_i \neq \emptyset$ for each $i \in \Pi$ with $x_i = 1$, we have $\bar{x} \leq \bar{z} \leq \bar{y}$.

(\supseteq) Suppose there exists a Nash equilibrium of (\mathcal{G}, v) with payoff $\bar{x} \leq \bar{z} \leq \bar{y}$. So, by Proposition 5, there exists a non-trivial strongly connected set U reachable from v and contained in $\bigcap_{z_i=0} (V \setminus W_i)$ such that $U \cap F_i \neq \emptyset$ iff $z_i = 1$. Let X be as in the first call of the procedure `SolveSubgame`. Clearly, U is contained in a non-trivial SCC C of $G \upharpoonright X$. Let $L = \{i \in \Pi : C \cap F_i = \emptyset\}$. If C is, in fact, contained in $\bigcap_{i \in L} (V \setminus W_i)$, then the algorithm will add v to Z immediately. Otherwise, $U \subseteq C \cap \bigcap_{i \in L} (V \setminus W_i) = Y$, and U is contained in a non-trivial SCC of $G \upharpoonright Y$. Repeating this argument, the procedure will eventually find a non-trivial SCC $C \supseteq Y$ of $G \upharpoonright X$ such that C is contained in $\bigcap_{i \in L} (V \setminus W_i)$. So $L \subseteq \{i \in \Pi : z_i = 0\}$, and C is reachable from v in $G \upharpoonright \bigcap_{i \in L} (V \setminus W_i)$. Hence, v is added to Z . Q.E.D.

Proof of Theorem 14 (continued)

It remains to show that φ is satisfiable if and only if (\mathcal{G}_φ, C) has a Nash equilibrium where player 2 loses (with C being an arbitrary clause).

(\Rightarrow) Assume that φ is satisfiable. Consider player 1's positional strategy of playing from a clause to a literal that satisfies this clause. This ensures that for each variable X at most one of the literals X or $\neg X$ is visited infinitely often. So, for any strategy τ of player 2, the strategy profile (σ, τ) is a Nash equilibrium where player 2 loses.

(\Leftarrow) Assume that φ is not satisfiable. We claim that player 2 has a winning strategy for the game (\mathcal{G}_φ, C) . So there can be no Nash equilibrium where she loses. Otherwise, by [13], player 1 would not only have a strategy to ensure that for some variable X either X or $\neg X$ is visited infinitely often but a pure positional one. But a pure positional strategy σ of player 1 chooses for each clause a literal that is contained in this clause. Since φ is unsatisfiable

there must exist a variable X and clauses C and C' such that $\sigma(C) = X$ and $\sigma(C') = \neg X$. Player 2's counter strategy is to play from X to C' and from any other literal to C . So the strategy σ is not winning, a contradiction. Q.E.D.

Proof of Theorem 15 (continued)

It remains to show that φ is satisfiable if and only if (\mathcal{G}_φ, C) has a Nash equilibrium where player 2 loses (with C being an arbitrary clause)

(\Rightarrow) Assume that φ is satisfiable. We show that the pure positional strategy profile where at any time player 0 plays from a clause C to a (fixed) literal that satisfies this clause and all other players never play to the sink vertex is a Nash equilibrium where player 0 wins. As the induced play never reaches the sink, it is certainly true that player 0 wins. Now consider any player i that loses the induced play, which can only happen if a vertex (C, X_i) is visited infinitely often. But, as player 0 plays according to the satisfying assignment, this means that no vertex $(C', \neg X_i)$ is ever visited, hence player i has no chance to improve her payoff by playing to the sink vertex.

(\Leftarrow) Assume that $(\mathcal{G}_\varphi, C_1)$ has a (w.l.o.g. pure) Nash equilibrium where player 0 wins, hence the sink vertex is not reached in the induced play. Consider the variable assignment that maps X_i to true if some vertex (C, X_i) is visited infinitely often. We claim that this assignment satisfies the formula. To see this, consider any clause C_j . By the construction of \mathcal{G}_φ , there exists a literal X_i or $\neg X_i$ in C_j such that the vertex (C_j, X_i) or $(C_j, \neg X_i)$, respectively, is visited infinitely often. Now assume that both a vertex (C, X_i) and a vertex $(C', \neg X_i)$ are visited infinitely often. Then player i would lose, but could improve her payoff by playing from $(C', \neg X_i)$ to the sink vertex. Hence, in any case the defined interpretation maps the literal to true thus satisfying the clause. Q.E.D.