Solution Concepts and Algorithms for Infinite Multiplayer Games

Erich Grädel and Michael Ummels

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

Abstract. *We survey and discuss several solution concepts for infinite turn-based multiplayer games with qualitative (i.e. win-lose) objectives of the players. These games generalise in a natural way the common model of games in verification which are two-player, zero-sum games with ω-regular winning conditions. The generalisation is in two directions: our games may have more than two players, and the objectives of the players need not be completely antagonistic.*

The notion of a Nash equilibrium is the classical solution concept in game theory. However, for games that extend over time, in particular for games of infinite duration, Nash equilibria are not always satisfactory as a notion of rational behaviour. We therefore discuss variants of Nash equilibria such as subgame perfect equilibria and secure equilibria. We present criteria for the existence of Nash equilibria and subgame perfect equilibria in the case of arbitrarily many players and for the existence of secure equilibria in the two-player case. In the second part of this paper, we turn to algorithmic questions: For each of the solution concepts that we discuss, we present algorithms that decide the existence of a solution with certain requirements in a game with parity winning conditions. Since arbitrary ω-regular winning conditions can be reduced to parity conditions, our algorithms are also applicable to games with arbitrary ω-regular winning conditions.

1 Introduction

Infinite games in which two or more players take turns to move a token through a directed graph, tracing out an infinite path, have numerous applications in computer science. The fundamental mathematical questions on such games concern the existence of optimal strategies for the players, the complexity and structural properties of such strategies, and their realisation by efficient algorithms. Which games are determined, in the sense that from each position, one of the players has a winning strategy? How to compute winning positions and optimal strategies? How much knowledge on the past of a play is necessary to determine an optimal next action? Which games are determined by memoryless strategies? And so on.

The case of two-player, zero-sum games with perfect information and ω-regular winning conditions has been extensively studied, since it is the basis of a rich methodology for the synthesis and verification of reactive systems. On the other side, other models of games, and in particular the case of infinite multiplayer games, are less understood and much more complicated than the two-player case.

In this paper we discuss the advantages and disadvantages of several solution concepts for infinite multiplayer games. These are Nash equilibria, subgame perfect equilibria, and secure equilibria. We focus on turn-based games with perfect information and qualitative winning conditions, i.e. for each player, the outcome of a play is either win or lose. The games are not necessarily completely antagonistic, which means that a play may be won by several players or by none of them.

Of course, the world of infinite multiplayer games is much richer than this class of games, and includes also concurrent games, stochastic games, games with various forms of imperfect or incomplete information, and games with quantitative objectives of the players. However, many of the phenomena that we wish to illustrate appear already in the setting studied here. To which extent our ideas and solutions can be carried over to other scenarios of infinite multiplayer games is an interesting topic of current research.

The outline of this paper is as follows. After fixing our notation in Section 2, we proceed with the presentation of several solution concepts for infinite multiplayer games in Section 3. For each of the three solution concepts (Nash equilibria, subgame perfect equilibria, and secure equilibria) we discuss, we devise criteria for their existence. In particular, we will relate the existence of a solution to the determinacy of certain two-player zero-sum games.

In Section 4, we turn to algorithmic questions, where we focus on games with parity winning conditions. We are interested in deciding the existence of a solution with certain requirements on the payoff. For Nash equilibria, it turns out that the problem is NP-complete, in general. However, there exists a natural restriction of the problem where the complexity goes down to UP $\cap$ co-UP (or even P for less complex winning conditions). Unfortunately, for subgame perfect equilibria we can only give an ExpTime upper bound for the complexity of the problem. For secure equilibria, we focus on two-player games. Depending on which requirement we impose on the payoff, we show that the problem falls into one of the complexity classes UP $\cap$ co-UP, NP, or co-NP.

2 Infinite Multiplayer Games

We consider here infinite turn-based multiplayer games on graphs with perfect information and qualitative objectives for the players. The definition of such games readily generalises from the two-player case. A game is defined by an arena and by the winning conditions for the players. We usually assume that the winning condition for each player is given by a set of infinite sequences of colours (from a finite set of colours) and that the winning conditions of the players are, a priori, independent.

Definition 1. An *infinite (turn-based, qualitative) multiplayer game* is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (\mathrm{Win}_i)_{i\in\Pi})$ where Π is a finite set of *players*, (V, E) is a (finite or infinite) directed graph, $(V_i)_{i\in\Pi}$ is a partition of V into the position sets for each player, $\chi : V \to C$ is a colouring of the position by some set C, which is usually assumed to be finite, and $\mathrm{Win}_i \subseteq C^\omega$ is the winning condition for player i.

The structure $G = (V, (V_i)_{i\in\Pi}, E, \chi)$ is called the *arena* of $\mathcal{G}$. For the sake of simplicity, we assume that $uE := \{v \in V : (u, v) \in E\} \neq \emptyset$ for all $u \in V$, 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 C^ω.

A *play* of $\mathcal{G}$ is an infinite path through the graph (V, E), and a *history* is a finite initial segment of a play. We say that a play π is *won* by player $i \in \Pi$ if $\chi(\pi) \in \mathrm{Win}_i$. The *payoff* of a play π of $\mathcal{G}$ is the vector $\mathrm{pay}(\pi) \in \{0, 1\}^\Pi$ defined by $\mathrm{pay}(\pi)_i = 1$ if π is won by player i. A *(pure) strategy of player i in $\mathcal{G}$* is a function $\sigma : V^*V_i \to V$ assigning to each sequence xv of position ending in a position v of player i a next position $\sigma(xv)$ such that $(v, \sigma(xv)) \in E$. We say that a play $\pi = \pi(0)\pi(1) \ldots$ of $\mathcal{G}$ is *consistent* with a strategy σ of player i if $\pi(k + 1) = \sigma(\pi(0) \ldots \pi(k))$ for all $k < \omega$ with $\pi(k) \in V_i$. A *strategy profile of $\mathcal{G}$* is a tuple $(\sigma_i)_{i\in\Pi}$ where σ_i is a strategy of player i.

A strategy σ of player i is called *positional* if σ depends only on the current vertex, i.e. if $\sigma(xv) = \sigma(v)$ for all $x \in V^*$ and $v \in V_i$. More generally, σ is called a *finite-memory strategy* if the equivalence relation $\sim_\sigma$ on V^* defined by $x \sim_\sigma x'$ if $\sigma(xz) = \sigma(x'z)$ for all $z \in V^*V_i$ has finite index. In other words, a finite-memory strategy is a strategy that can be implemented by a finite automaton with output. A strategy profile $(\sigma_i)_{i\in\Pi}$ is called *positional* or a *finite-memory strategy profile* if each σ_i is positional or a finite-memory strategy, respectively.

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 infinite multiplayer 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 $\mathcal{G}$. A strategy σ of some player i in $(\mathcal{G}, v_0)$ is *winning* if every play of $(\mathcal{G}, v_0)$ consistent with σ is won by player i. A strategy profile $(\sigma_i)_{i\in\Pi}$ of $(\mathcal{G}, v_0)$ determines a unique play of $(\mathcal{G}, v_0)$ consistent with each σ_i, called the *outcome of* $(\sigma_i)_{i\in\Pi}$ and denoted by $\langle(\sigma_i)_{i\in\Pi}\rangle$ or, in the case that the initial vertex is not understood from the context, $\langle(\sigma_i)_{i\in\Pi}\rangle_{v_0}$. In the following, we will often use the term *game* to denote an *(initialised) infinite multiplayer game* according to Definition 1.

We have introduced winning conditions as abstract sets of infinite sequences over the set of colours. In verification the winning conditions usually are *ω-regular sets* specified by formulae of the logic S1S (monadic second-order logic on infinite words) or LTL (linear-time temporal logic) referring to unary predicates P_c indexed by the set C of colours. Special cases are the following well-studied winning conditions:

- *Büchi* (given by $F \subseteq C$): defines the set of all $\alpha \in C^\omega$ such that $\alpha(k) \in F$ for infinitely many $k < \omega$.
- *co-Büchi* (given by $F \subseteq C$): defines the set of all $\alpha \in C^\omega$ such that $\alpha(k) \in F$ for all but finitely many $k < \omega$.
- *Parity* (given by a *priority function* $\Omega : C \to \omega$): defines the set of all $\alpha \in C^\omega$ such that the least number occurring infinitely often in $\Omega(\alpha)$ is even.
- *Rabin* (given by a set Ω of *pairs* (G_i, R_i) where $G_i, R_i \subseteq C$): defines the set of all $\alpha \in C^\omega$ such that there exists an index i with $\alpha(k) \in G_i$ for infinitely many $k < \omega$ but $\alpha(k) \in R_i$ only for finitely many $k < \omega$.
- *Streett* (given by a set Ω of *pairs* (G_i, R_i) where $G_i, R_i \subseteq C$): defines the set of all $\alpha \in C^\omega$ such that for all indices i with $\alpha(k) \in R_i$ for infinitely many $k < \omega$ also $\alpha(k) \in G_i$ for infinitely many $k < \omega$.
- *Muller* (given by a family $\mathcal{F}$ of *accepting sets* $F_i \subseteq C$): defines the set of all $\alpha \in C^\omega$ such that there exists an index i with the set of colours seen infinitely often in α being precisely the set F_i.

Note that (co-)Büchi conditions are a special case of parity conditions with two priorities, and parity conditions are a special case of Rabin and Streett conditions, which are special cases of Muller conditions. Moreover, the complement of a Büchi or Rabin condition is a co-Büchi or Streett condition, respectively, and vice versa, whereas the class of parity conditions and the class of Muller conditions are closed under complement. Finally, any of these conditions is *prefix independent*, i.e. for every $\alpha \in C^\omega$ and $x \in C^*$ 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, Rabin, Streett, or Muller game* if the winning condition of *each* player is of the specified type. This differs somewhat from the usual convention for two-player zero-sum games where a Büchi or Rabin game is a game where the winning condition of the *first* player is a Büchi or Rabin condition, respectively.

Note that we do distinguish between colours and priorities. For two-player zero-sum parity games, one can identify them by choosing a finite subset of ω as the set C of colours and defining the parity condition directly on C, i.e. the priority function of the first player is the identity function, and the priority function of the second player is the successor function $k \mapsto k + 1$. This gives *parity games* as considered in the literature [29].

The importance of the parity condition stems from three facts: First, the condition is expressive enough to express any ω-regular objective. More precisely, for every ω-regular language of infinite words, there exists a deterministic word automaton with a parity acceptance condition that recognises this language. As demonstrated by Thomas [26], this allows to reduce a two-player zero-sum game with an arbitrary ω-regular winning condition to a parity game. (See also W. Thomas' contribution to this volume.) Second, two-player zero-sum parity games arise as the model-checking games for fixed-point logics, in particular the modal μ-calculus [11]. Third, the condition is simple enough to allow for *positional* winning strategies (see above) [8, 19], i.e. if one player has a winning strategy in a parity game she also has a positional one. It is easy to see that the first property extends to the multiplayer case: Any multiplayer game with ω-regular winning conditions can be reduced to a game with parity winning conditions [27]. Hence, in the algorithmic part of this paper, we will concentrate on multiplayer parity games.

3 Solution Concepts

So far, the infinite games used in verification mostly are two-player games with win-lose conditions, i.e. each play is won by one player and lost by the other. The key concept for such games is *determinacy*: a game is determined if, from each initial position, one of the players has a winning strategy.

While it is well-known that, on the basis of (a weak form of) the Axiom of Choice, non-determined games exist, the two-player win-lose games usually encountered in computer science, in particular all ω-regular games, are determined. Indeed, this is true for much more general games where the winning conditions are arbitrary (quasi-)Borel sets [17, 18].

In the case of a determined game, solving the game means to compute the winning regions and winning strategies for the two players. A famous result due to Büchi and Landweber [3] says that in the case of games on finite graphs and with ω-regular winning conditions, we can effectively compute winning strategies that are realisable by finite automata.

When we move to multiplayer games and/or non-zero sum games, other solution concepts are needed. We will explain some of these concepts, in particular Nash equilibria, subgame perfect equilibria, and secure equilibria, and relate the existence of these equilibria (for the kind of infinite games studied here) to the determinacy of certain associated two-player games.

3.1 Nash Equilibria

The most popular solution concept in classical game theory is the concept of a *Nash equilibrium*. Informally, a Nash equilibrium is a strategy profile from which no player has an incentive to deviate, if the other players stick to their strategies. A celebrated theorem by John Nash [21] says that in any game where each player only has a finite collection of strategies there is at least one Nash equilibrium provided that the players can randomise over their strategies, i.e. choose *mixed strategies* rather than only pure ones. For turn-based (non-stochastic) games with qualitative winning conditions, mixed strategies play no relevant role. We define Nash equilibria just in the form needed here.

Definition 2. A strategy profile $(\sigma_i)_{i\in\Pi}$ of a game $(\mathcal{G}, v_0)$ is called a *Nash equilibrium* if for every player $i \in \Pi$ and all her possible strategies σ' in $(\mathcal{G}, v_0)$ the play $\langle \sigma', (\sigma_j)_{j\in\Pi\setminus\{i\}}\rangle$ is won by player i only if the play $\langle (\sigma_j)_{j\in\Pi}\rangle$ is also won by her.

It has been shown by Chatterjee & al. [6] that every multiplayer game with Borel winning conditions has a Nash equilibrium. We will prove a more general result below.

Despite the importance and popularity of Nash equilibria, there are several problems with this solution concept, in particular for games that extend over time. This is due to the fact that Nash equilibria do not take into account the sequential nature of these games and its consequences. After any initial segment of a play, the players face a new situation and may change their strategies. Choices made because of a threat by the other players may no longer be rational, because the opponents have lost their power of retaliation in the remaining play.

Example 3. Consider a two-player Büchi game with its arena depicted in Figure 1; round vertices are controlled by player 1; boxed vertices are controlled by player 2; each of the two players wins if and only if vertex 3 is visited (infinitely often); the initial vertex is 1. Intuitively, the only rational outcome of this game should be the play 123^ω. However, the game has two Nash equilibria:

1. Player 1 moves from vertex 1 to vertex 2, and player 2 moves from vertex 2 to vertex 3. Hence, both players win.
2. Player 1 moves from vertex 1 to vertex 4, and player 2 moves from vertex 2 to vertex 5. Hence, both players lose.

The second equilibrium certainly does not describe rational behaviour. Indeed both players move according to a strategy that is always losing (whatever the other player does), and once player 1 has moved from vertex 1 to vertex 2, then the rational behaviour of player 2 would be to change her strategy and move to vertex 3 instead of vertex 5 as this is then the only way for her to win.

This example can be modified in many ways. Indeed we can construct games with Nash equilibria in which every player moves infinitely often according to a losing strategy, and only has a chance to win if she deviates from the equilibrium strategy. The following is an instructive example with quantitative objectives.

Example 4. Let $\mathcal{G}_n$ be an n-player game with positions $0, \dots, n$. Position n is the initial position, and position 0 is the terminal position. Player i moves at position i and has two options. Either she loops at position i (and stays

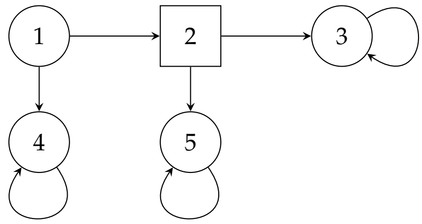

Figure 1. A two-player Büchi game.

in control) or moves to position $i-1$ (handing control to the next player). For each player, the value of a play π is $(n+1)/|\pi|$. Hence, for all players, the shortest possible play has value 1, and all infinite plays have value 0. Obviously, the rational behaviour for each player i is to move from i to $i-1$. This strategy profile, which is of course a Nash equilibrium, gives value 1 to all players. However, the 'most stupid' strategy profile, where each player loops forever at his position, i.e. moves forever according to a losing strategy, is also a Nash equilibrium.

3.2 Subgame Perfect Equilibria

An equilibrium concept that respects the possibility of a player to change her strategy during a play is the notion of a subgame perfect equilibrium [25]. For being a subgame perfect equilibrium, a choice of strategies is not only required to be optimal for the initial vertex but for every possible initial history of the game (including histories not reachable in the equilibrium play).

To define subgame perfect equilibria formally, we need the notion of a subgame: For a game $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (\mathrm{Win}_i)_{i\in\Pi})$ and a history h of $\mathcal{G}$, let the game $\mathcal{G}|_h = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (\mathrm{Win}_i|_h)_{i\in\Pi})$ be defined by $\mathrm{Win}_i|_h = \{\alpha \in C^\omega : \chi(h)\cdot\alpha \in \mathrm{Win}_i\}$. For an initialised game $(\mathcal{G}, v_0)$ and a history hv of $(\mathcal{G}, v_0)$, we call the initialised game $(\mathcal{G}|_h, v)$ the *subgame of* $(\mathcal{G}, v_0)$ *with history* hv. For a strategy σ of player $i \in \Pi$ in $\mathcal{G}$, let $\sigma|_h : V^*V_i \to V$ be defined by $\sigma|_h(xv) = \sigma(hxv)$. Obviously, $\sigma|_h$ is a strategy of player i in $\mathcal{G}|_h$.

Definition 5. A strategy profile $(\sigma_i)_{i\in\Pi}$ of a game $(\mathcal{G}, v_0)$ is called a *subgame perfect equilibrium (SPE)* if $(\sigma_i|_h)_{i\in\Pi}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$.

Example 6. Consider again the game described in Example 3. The Nash equilibrium where player 1 moves from vertex 1 to vertex 4 and player 2 moves from vertex 2 to vertex 5 is not a subgame perfect equilibrium since moving from vertex 2 to vertex 5 is not optimal for player 2 after the play has reached vertex 2. On the other hand, the Nash equilibrium where player 1 moves from vertex 1 to vertex 2 and player 2 moves from vertex 2 to vertex 3 is also a subgame perfect equilibrium.

It is a classical result due to Kuhn [16] that every *finite* game (i.e. every game played on a finite tree with payoffs attached to leaves) has a subgame perfect equilibrium. The first step in the analysis of subgame perfect equilibria for *infinite* duration games is the notion of subgame-perfect determinacy. While the notion of subgame perfect equilibrium makes sense for more general classes of infinite games, the notion of subgame-perfect determinacy applies only to games with qualitative winning conditions (which is tacitly assumed from now on).

Definition 7. A game $(\mathcal{G}, v_0)$ is *subgame-perfect determined* if there exists a strategy profile $(\sigma_i)_{i\in\Pi}$ such that for each history hv of the game one of the strategies $\sigma_i|_h$ is a winning strategy in $(\mathcal{G}|_h, v)$.

Proposition 8. Let $(\mathcal{G}, v_0)$ be a qualitative zero-sum game such that every subgame is determined. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Proof. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that, for every history hv there exists a strategy σ_i^h for some player i that is winning in $(\mathcal{G}|_h, v)$. (Note that we can assume that σ_i^h is independent of v.) We have to combine these strategies in an appropriate way to strategies σ_i. (Let us point out that the trivial combination, namely $\sigma_i(hv) := \sigma_i^h(v)$ does not work in general.) We say that a decomposition $h = h_1 \cdot h_2$ is *good* for player i w.r.t. vertex v if $\sigma_i^{h_1}|_{h_2}$ is winning in $(\mathcal{G}|_h, v)$. If the strategy σ_i^h is winning in $(\mathcal{G}|_h, v)$, then the decomposition $h = h \cdot \varepsilon$ is good w.r.t. v, so a good decomposition exists.

For each history hv, if σ_i^h is winning in $(\mathcal{G}|_h, v)$, we choose the good (w.r.t. vertex v) decomposition $h = h_1h_2$ with minimal h_1, and put

$$\sigma_i(hv) := \sigma_i^{h_1}(h_2v) .$$

Otherwise, we set

$$\sigma_i(hv) := \sigma_i^h(v) .$$

It remains to show that for each history hv of $(\mathcal{G}, v_0)$ the strategy $\sigma_i|_h$ is winning in $(\mathcal{G}|_h, v)$ whenever the strategy σ_i^h is. Hence, assume that σ_i^h is winning in $(\mathcal{G}|_h, v)$, and let $\pi = \pi(0)\pi(1)\ldots$ be a play starting in $\pi(0) = v$ and consistent with $\sigma_i|_h$. We need to show that π is won by player i in $(\mathcal{G}|_h, v)$.

First, we claim that for each $k < \omega$ there exists a decomposition of the form $h\pi(0)\ldots\pi(k-1) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-1))$ that is good for player i w.r.t. $\pi(k)$. This is obviously true for $k = 0$. Now, for $k > 0$, assume that there exists a decomposition $h\pi(0)\ldots\pi(k-2) = h_1 \cdot (h_2\pi(0)\ldots\pi(k-2))$ that is good for player i w.r.t. $\pi(k-1)$ and with h_1 being minimal. Then $\pi(k) = \sigma_i(h\pi(0)\ldots\pi(k-1)) = \sigma^{h_1}(h_2\pi(0)\ldots\pi(k-1))$, and $h\pi(0)\ldots\pi(k-1) = h_1(h_2\pi(0)\ldots\pi(k-1))$ is a decomposition that is good w.r.t. $\pi(k)$.

Now consider the sequence $h_1^0, h_1^1, \ldots$ of prefixes of the good decompositions $h\pi(0)\ldots\pi(k-1) = h_1^k h_2^k \pi(0)\ldots\pi(k-1)$ (w.r.t. $\pi(k)$) with each h_1^k being minimal. Then we have $h_1^0 \succeq h_1^1 \succeq \ldots$, since for each $k > 0$ the decomposition $h\pi(0)\ldots\pi(k-1) = h_1^{k-1}h_2^{k-1}\pi(0)\ldots\pi(k-1)$ is also good for player i w.r.t. $\pi(k)$. As $\prec$ is well-founded, there must exist $k < \omega$ such that $h_1 := h_1^k = h_1^l$ and $h_2 := h_2^k = h_2^l$ for each $k \le l < \omega$. Hence, we have that the play $\pi(k)\pi(k+1)\ldots$ is consistent with $\sigma_i^{h_1}|_{h_2\pi(0)\ldots\pi(k-1)}$, which is a winning strategy in $(\mathcal{G}|_{h\pi(0)\ldots\pi(k-1)}, \pi(k))$. So the play $h\pi$ is won by player i in $(\mathcal{G}, v_0)$, which implies that the play π is won by player i in $(\mathcal{G}|_h, v)$. Q.E.D.

We say that a class of winning conditions is closed under taking subgames, if for every condition $X \subseteq C^\omega$ in the class, and every $h \in C^*$, also $X|_h := \{x \in C^\omega : hx \in X\}$ belongs to the class. Since Borel winning conditions are closed under taking subgames, it follows that any two-player zero-sum game with Borel winning condition is subgame-perfect determined.

Corollary 9. Let $(\mathcal{G}, v_0)$ be a two-player zero-sum Borel game. Then $(\mathcal{G}, v_0)$ is subgame-perfect determined.

Multiplayer games are usually not zero-sum games. Indeed when we have many players the assumption that the winning conditions of the players form a partition of the set of plays is very restrictive and unnatural. We now drop this assumption and establish general conditions under which a multiplayer game admits a subgame perfect equilibrium. In fact we will relate the existence of subgame perfect equilibria to the determinacy of associated two-player games. In particular, it will follow that every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

In the rest of this subsection, we are only concerned with the *existence* of equilibria, not with their complexity. Thus, without loss of generality, we assume that the arena of the game under consideration is a tree or a forest with the initial vertex as one of its roots. The justification for this assumption is that we can always replace the arena of an arbitrary game by its unravelling from the initial vertex, ending up in an equivalent game.

Definition 10. Let $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (\text{Win}_i)_{i\in\Pi})$ be a multiplayer game (played on a forest), with winning conditions $\text{Win}_i \subseteq C^\omega$. The associated class $\text{ZeroSum}(\mathcal{G})$ of two-player zero-sum games is obtained as follows:

1. For each player i, $\text{ZeroSum}(\mathcal{G})$ contains the game $\mathcal{G}_i$ where player i plays $\mathcal{G}$, with his winning condition Win_i, against the coalition of all other players, with winning condition $C^\omega \setminus \text{Win}_i$.
2. Close the class under taking subgames (i.e. consider plays after initial histories).
3. Close the class under taking subgraphs (i.e. admit deletion of positions and moves).

Note that the order in which the operations 1, 2 and 3 are applied has no effect on the class $\text{ZeroSum}(\mathcal{G})$.

Theorem 11. Let $(\mathcal{G}, v_0)$ be a multiplayer game such that every game in $\text{ZeroSum}(\mathcal{G})$ is determined. Then $(\mathcal{G}, v_0)$ has a subgame perfect equilibrium.

Proof. Let $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \chi, (\text{Win}_i)_{i\in\Pi})$ be a multiplayer game such that every game in ZeroSum$(\mathcal{G})$ is determined. For each ordinal α we define a set $E^\alpha \subseteq E$ beginning with $E^0 = E$ and

$$E^\lambda = \bigcap_{\alpha<\lambda} E^\alpha$$

for limit ordinals λ. To define $E^{\alpha+1}$ from E^α, we consider for each player $i \in \Pi$ the two-player zero-sum game $\mathcal{G}_i^\alpha = (V, V_i, E^\alpha, \chi, \text{Win}_i)$ where player i plays with his winning condition Win_i against the coalition of all other players (with winning condition $C^\omega \setminus \text{Win}_i$). Every subgame of $\mathcal{G}_i^\alpha$ belongs to ZeroSum$(\mathcal{G})$ and is therefore determined. Hence we can use Proposition 8 to fix a subgame perfect equilibrium $(\sigma_i^\alpha, \sigma_{-i}^\alpha)$ of $(\mathcal{G}_i^\alpha, v_0)$ where σ_i^α is a strategy of player i and σ_{-i}^α is a strategy of the coalition. Moreover, as the arena of $\mathcal{G}^\alpha$ is a forest, these strategies can be assumed to be positional. Let X_i^α be the set of all $v \in V$ such that σ_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for the unique maximal history h of $\mathcal{G}$ leading to v. For vertices $v \in V_i \cap X_i^\alpha$ we delete all outgoing edges except the one taken by the strategy σ_i^α, i.e. we define

$$E^{\alpha+1} = E^\alpha \setminus \bigcup_{i\in\Pi} \{(u,v) \in E : u \in V_i \cap X_i^\alpha \text{ and } v \neq \sigma_i^\alpha(u)\}\,.$$

Obviously, the sequence $(E^\alpha)_{\alpha\in\text{On}}$ is nonincreasing. Thus we can fix the least ordinal ξ with $E^\xi = E^{\xi+1}$ and define $\sigma_i = \sigma_i^\xi$ and $\sigma_{-i} = \sigma_{-i}^\xi$. Moreover, for each player $j \neq i$ let $\sigma_{j,i}$ be the positional strategy of player j in $\mathcal{G}$ that is induced by σ_{-i}.

Intuitively, Player i's equilibrium strategy τ_i is as follows: Player i plays σ_i as long as no other player deviates. Whenever some player $j \neq i$ deviates from her equilibrium strategy σ_j, player i switches to $\sigma_{i,j}$. Formally, define for each vertex $v \in V$ the player $p(v)$ who has to be "punished" at vertex v where $p(v) = \bot$ if nobody has to be punished. If the game has just started, no player should be punished. Thus we let

$$p(v) = \bot \text{ if } v \text{ is a root.}$$

At vertex v with predecessor u, the same player has to be punished as at vertex u as long as the player whose turn it was at vertex u did not deviate from her prescribed strategy. Thus for $u \in V_i$ and $v \in uE$ we let

$$p(v) = \begin{cases} \bot & \text{if } p(u) = \bot \text{ and } v = \sigma_i(u), \\ p(u) & \text{if } p(u) \neq i,\ p(u) \neq \bot \text{ and } v = \sigma_{i,p(u)}(u), \\ i & \text{otherwise.} \end{cases}$$

Now, for each player $i \in \Pi$ we can define the equilibrium strategy τ_i by setting

$$\tau_i(v) = \begin{cases} \sigma_i(v) & \text{if } p(v) = \bot \text{ or } p(v) = i, \\ \sigma_{i,p(v)}(v) & \text{otherwise} \end{cases}$$

for each $v \in V$.

It remains to show that $(\tau_i)_{i\in\Pi}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. First note that σ_i is winning in $(\mathcal{G}_i^\xi|_h, v)$ if σ_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ for some ordinal α because if σ_i^α is winning in $(\mathcal{G}_i^\alpha|_h, v)$ every play of $(\mathcal{G}_i^{\alpha+1}|_h, v)$ is consistent with σ_i^α and therefore won by player i. As $E^\xi \subseteq E^{\alpha+1}$, this also holds for every play of $(\mathcal{G}_i^\xi|_h, v)$. Now let v be any vertex of $\mathcal{G}$ with h the unique maximal history of $\mathcal{G}$ leading to v. We claim that $(\tau_j)_{j\in\Pi}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$. Towards this, let τ' be any strategy of any player $i \in \Pi$ in $\mathcal{G}$; let $\pi = \langle(\tau_j)_{j\in\Pi}\rangle_v$, and let $\pi' = \langle\tau', (\tau_j)_{j\in\Pi\setminus\{i\}}\rangle_v$. We need to show that $h\pi$ is won by player i or that $h\pi'$ is not won by player i. The claim is trivial if $\pi = \pi'$. Thus assume that $\pi \neq \pi'$ and fix the least $k < \omega$ such that $\pi(k+1) \neq \pi'(k+1)$. Clearly, $\pi(k) \in V_i$ and $\tau'(\pi(k)) \neq \tau_i(\pi(k))$. Without loss of generality, let $k = 0$. We distinguish the following two cases:

- σ_i is winning in $(\mathcal{G}_i^\xi|_h, v)$. By the definition of each τ_j, π is a play of $(\mathcal{G}_i^\xi|_h, v)$. We claim that π is consistent with σ_i, which implies that $h\pi$ is won by player i. Otherwise fix the least $l < \omega$ such that $\pi(l) \in V_i$ and $\sigma_i(\pi(l)) \neq \pi(l+1)$. As σ_i is winning in $(\mathcal{G}_i^\xi|_h, v)$, σ_i is also winning in $(\mathcal{G}_i^\xi|_{h\pi(0)\dots\pi(l-1)}, \pi(l))$. But then $(\pi(l), \pi(l+1)) \in E^\xi \setminus E^{\xi+1}$, a contradiction to $E^\xi = E^{\xi+1}$.
- σ_i is not winning in $(\mathcal{G}_i^\xi|_h, v)$. Hence σ_{-i} is winning in $(\mathcal{G}_i^\xi|_h, v)$. As $\tau'(v) \neq \tau_i(v)$, player i has deviated, and it is the case that $\pi' = \langle\tau', (\sigma_{j,i})_{j\in\Pi\setminus\{i\}}\rangle_v$. We claim that π' is a play of $(\mathcal{G}_i^\xi|_h, v)$. As σ_{-i} is winning in $(\mathcal{G}_i^\xi|_h, v)$, this implies that $h\pi'$ is not won by player i. Otherwise fix the least $l < \omega$ such that $(\pi'(l), \pi'(l+1)) \notin E^\xi$ together with the ordinal α such that $(\pi'(l), \pi'(l+1)) \in E^\alpha \setminus E^{\alpha+1}$. Clearly, $\pi'(l) \in V_i$. Thus σ_i^α is winning in $(\mathcal{G}_i^\alpha|_{h\pi'(0)\dots\pi'(l-1)}, \pi'(l))$, which implies that σ_i is winning in $(\mathcal{G}_i^\xi|_{h\pi'(0)\dots\pi'(l-1)}, \pi'(l))$. As π' is consistent with σ_{-i}, this means that σ_{-i} is not winning in $(\mathcal{G}_i^\xi|_h, v)$, a contradiction.

It follows that $(\tau_j)_{j\in\Pi} = (\tau_j|_h)_{j\in\Pi}$ is a Nash equilibrium of $(\mathcal{G}|_h, v)$ for every history hv of $(\mathcal{G}, v_0)$. Hence, $(\tau_j)_{j\in\Pi}$ is a subgame perfect equilibrium of $(\mathcal{G}, v_0)$. Q.E.D.

Corollary 12 ([27])**.** Every multiplayer game with Borel winning conditions has a subgame perfect equilibrium.

This generalises the result by Chatterjee & al. [6] that every multiplayer game with Borel winning conditions has a Nash equilibrium. Indeed, for the existence of Nash equilibria, a slightly weaker condition than the one in Theorem 11 suffices. Let $\text{ZeroSum}(\mathcal{G})_{\text{Nash}}$ be defined in the same way as $\text{ZeroSum}(\mathcal{G})$ but without closure under subgraphs.

Corollary 13. If every game in $\text{ZeroSum}(\mathcal{G})_{\text{Nash}}$ is determined, then $\mathcal{G}$ has a Nash equilibrium.

3.3 *Secure Equilibria*

The notion of a *secure equilibrium* introduced by Chatterjee & al. [4] tries to overcome another deficiency of Nash equilibria: one game may have many Nash equilibria with different payoffs and even several maximal ones w.r.t. to the componentwise partial ordering on payoffs. Hence, for the players it is not obvious which equilibrium to play. The idea of a secure equilibrium is that any rational deviation (i.e. a deviation that does not decrease the payoff of the player who deviates) will not only not increase the payoff of the player who deviates but it will also not decrease the payoff of any other player. Secure equilibria model rational behaviour if players not only attempt to maximise their own payoff but, as a secondary objective, also attempt to minimise their opponents' payoffs.

Definition 14. A strategy profile $(\sigma_i)_{i\in\Pi}$ of a game $(\mathcal{G}, v_0)$ is called *secure* if for all players $i \neq j$ and for each strategy σ' of player j it is the case that

$$\begin{aligned} &\langle(\sigma_i)_{i\in\Pi}\rangle \notin \mathrm{Win}_j \text{ or } \langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \sigma'\rangle \in \mathrm{Win}_j \\ \Rightarrow &\langle(\sigma_i)_{i\in\Pi}\rangle \notin \mathrm{Win}_i \text{ or } \langle(\sigma_i)_{i\in\Pi\setminus\{j\}}, \sigma'\rangle \in \mathrm{Win}_i \ . \end{aligned}$$

A strategy profile $(\sigma_i)_{i\in\Pi}$ is a *secure equilibrium* if it is both a Nash equilibrium and secure.

Example 15 ([4]). Consider another Büchi game played on the game graph depicted in Figure 1 by the two players 1 and 2 where, again, round vertices are controlled by player 1 and square vertices are controlled by player 2. This time player 1 wins if vertex 3 is visited (infinitely often), and player 2 wins if vertex 3 or vertex 5 is visited (infinitely often). Again, the initial vertex is 1.

Up to equivalence, there are two different strategies for each player: Player 1 can choose to go from 1 to either 2 or 4 while player 2 can choose to go from 2 to either 3 or 5. Except for the strategy profile where player 1 moves to 4 and player 2 moves to 3, all of the resulting profiles are Nash equilibria. However, the strategy profile where player 1 moves to 2 and player 2 moves to 3 is not secure: Player 2 can decrease player 1's payoff by moving to 5 instead while her payoff remains the same (namely 1). Similarly, the strategy profile where player 1 moves to 2 and player 2 moves to 5 is not secure: Player 1 can decrease player 2's payoff by moving to 4 instead while her payoff remains the same (namely 0). Hence, the strategy profile where player 1 moves to 4 and player 2 moves to 5 is the only secure equilibrium of the game.

It is an open question whether secure equilibria exist in arbitrary multi-player games with well-behaved winning conditions. However, for the case of only two players, it is not only known that there always exists a secure equilibrium for games with well-behaved winning conditions, but a unique maximal secure equilibrium payoff w.r.t. the componentwise ordering $\leq$ on payoffs, i.e. there exists a secure equilibrium (σ, τ) such that $\mathrm{pay}(\langle\sigma', \tau'\rangle) \leq \mathrm{pay}(\langle\sigma, \tau\rangle)$

for every secure equilibrium (σ', τ') of $(\mathcal{G}, v_0)$. Clearly, such an equilibrium is preferable for both players.

For two winning conditions $\text{Win}_1, \text{Win}_2 \subseteq V^\omega$, we say that the pair $(\text{Win}_1, \text{Win}_2)$ is *determined* if any Boolean combination of Win_1 and Win_2 is determined, i.e. any two-player zero-sum game that has a Boolean combination of Win_1 and Win_2 as its winning condition is determined.

Definition 16. A strategy σ of player 1 (player 2) in a two-player game $(\mathcal{G}, v_0)$ is *strongly winning* if it ensures a play with payoff $(1, 0)$ (payoff $(0, 1)$) against any strategy τ of player 2 (player 1).

The strategy σ is *retaliating* if it ensures a play with payoff $(0, 0)$, $(1, 0)$, or $(1, 1)$ against any strategy τ of player 2 (player 1).

Note that if $(\mathcal{G}, v_0)$ is a game with a determined pair $(\text{Win}_1, \text{Win}_2)$ of winning conditions, then player 1 or 2 has a strongly winning strategy if and only if the other player does not have a retaliating strategy.

Proposition 17. Let $(\mathcal{G}, v_0)$ be a two-player game with a determined pair $(\text{Win}_1, \text{Win}_2)$ of winning conditions. Then precisely one of the following four cases holds:

1. Player 1 has a strongly winning strategy;
2. Player 2 has a strongly winning strategy;
3. There is a pair of retaliating strategies with payoff $(1, 1)$;
4. There is a pair of retaliating strategies, and all pairs of retaliating strategies have payoff $(0, 0)$.

Proof. Note that if one player has a strongly winning strategy, then the other player neither has a strongly winning strategy nor a retaliating strategy. Vice versa, if one player has a retaliating strategy, then the other player cannot have a strongly winning strategy. Moreover, cases 3 and 4 exclude each other by definition. Hence, at most one of the four cases holds.

Now, assume that neither of the cases 1–3 holds. In particular, no player has a strongly winning strategy. By determinacy, this implies that both players have retaliating strategies. Let (σ, τ) be any pair of retaliating strategies. As case 3 does not hold, at least one of the two players receives payoff 0. But as both players play retaliating strategies, this implies that both players receive payoff 0, so we are in case 4. Q.E.D.

Theorem 18. Let $(\mathcal{G}, v_0)$ be a two-player game with a determined pair $(\text{Win}_1, \text{Win}_2)$ of winning conditions. Then there exists a unique maximal secure equilibrium payoff for $(\mathcal{G}, v_0)$.

Proof. We show that the claim holds in any of the four cases stated in Proposition 17:

1. In the first case, player 1 has a strongly winning strategy σ. Then, for any strategy τ of player 2, the strategy profile (σ, τ) is a secure equilibrium

with payoff $(1,0)$. We claim that $(1,0)$ is the unique maximal secure equilibrium payoff. Otherwise, there would exist a secure equilibrium with payoff 1 for player 2. But player 1 could decrease player 2's payoff while not decreasing her own payoff by playing σ, a contradiction.

2. The case that player 2 has a strongly winning strategy is analogous to the first case.
3. In the third case, there is a pair (σ, τ) of retaliating strategies with payoff $(1,1)$. But then (σ, τ) is a secure equilibrium, and $(1,1)$ is the unique maximal secure equilibrium payoff.
4. In the fourth case, there is a pair of retaliating strategies, and any pair of retaliating strategies has payoff $(0,0)$. Then there exists a strategy σ of player 1 that guarantees payoff 0 for player 2, since otherwise by determinacy there would exists a strategy for player 2 that guarantees payoff 1 for player 2. This would be a retaliating strategy that guarantees payoff 1 for player 2, a contradiction to the assumption that all pairs of retaliating strategies have payoff $(0,0)$. Symmetrically, there exists a strategy τ of player 2 that guarantees payoff 0 for player 1. By the definition of σ and τ, the strategy profile (σ, τ) is a Nash equilibrium. But it is also secure, since it gives each player the least possible payoff. Hence, (σ, τ) is a secure equilibrium. Now assume there exists a secure equilibrium (σ', τ') with payoff $(1,0)$. Then also (σ', τ) would give payoff 1 to player 1, a contradiction to the fact that (σ, τ) is a Nash equilibrium. Symmetrically, there cannot exists a secure equilibrium (σ', τ') with payoff $(0,1)$. Hence, either $(0,0)$ or $(1,1)$ is the unique maximal secure equilibrium payoff. Q.E.D.

Since Borel winning conditions are closed under Boolean combinations, as a corollary we get the result by Chatterjee & al. that any two-player game with Borel winning conditions has a unique maximal secure equilibrium payoff.

Corollary 19 ([4])**.** Let $(\mathcal{G}, v_0)$ be two-player game with Borel winning conditions. Then there exists a unique maximal secure equilibrium payoff for $(\mathcal{G}, v_0)$.

4 Algorithmic Problems

Previous research on algorithms for multiplayer games has focused on computing *some* solution of the game, e.g. *some* Nash equilibrium [6]. However, as we have seen, a game may not have a unique solution, so one might be interested not in *any* solution, but in a solution that fulfils certain requirements. For example, one might look for a solution where certain players win while certain other players lose. Or one might look for a *maximal* solution, i.e. a solution such that there does not exist another solution with a higher payoff. In the context of games with parity winning conditions, this motivation leads us to the following decision problem, which can be defined for any solution concept $\mathcal{S}$:

> Given a multiplayer parity game $(\mathcal{G}, v_0)$ played on a finite arena and thresholds $\bar{x}, \bar{y} \in \{0,1\}^k$, decide whether $(\mathcal{G}, v_0)$ has a solution $(\sigma_i)_{i\in\Pi} \in \mathcal{S}(\mathcal{G}, v_0)$ such that $\bar{x} \leq \text{pay}(\langle(\sigma_i)_{i\in\Pi}\rangle) \leq \bar{y}$.

In particular, the solution concepts of Nash equilibria, subgame perfect equilibria, and secure equilibria give rise to the decision problems NE, SPE and SE, respectively. In the following three sections, we analyse the complexity of these three problems.

4.1 Nash Equilibria

Let $(\mathcal{G}, v_0)$ be a game with prefix-independent, determined winning conditions. Assume we have found a Nash equilibrium $(\sigma_i)_{i\in\Pi}$ of $(\mathcal{G}, v_0)$ with payoff $\bar{x}$. Clearly, the play $\langle(\sigma_i)_{i\in\Pi}\rangle$ never hits the winning region W_i of some player i with $x_i = 0$ because otherwise player i can improve her payoff by waiting until the token hits W_i and then apply her winning strategy. The crucial observation is that this condition is also sufficient for a play to be induced by a Nash equilibrium, i.e. $(\mathcal{G}, v_0)$ has a Nash equilibrium with payoff $\bar{x}$ if and only if there exists a play in $(\mathcal{G}, v_0)$ with payoff $\bar{x}$ that never hits the winning region of some player i with $x_i = 0$.

Lemma 20. Let $(\mathcal{G}, v_0)$ be a k-player game with prefix-independent, determined winning conditions, and let W_i be the winning region of player i in $\mathcal{G}$. There exists a Nash equilibrium of $(\mathcal{G}, v_0)$ with payoff $\bar{x} \in \{0,1\}^k$ if and only if there exists a play π of $(\mathcal{G}, v_0)$ with payoff $\bar{x}$ such that $\{\pi(k) : k < \omega\} \cap W_i = \emptyset$ for each player i with $x_i = 0$.

Proof. ($\Rightarrow$) This direction follows from the argumentation above.

($\Leftarrow$) Let π be a play with payoff $\bar{x}$ such that $\{\pi(k) : k < \omega\} \cap W_i = \emptyset$ for each player i with $x_i = 0$. Moreover, let τ_{-j} be an optimal strategy of the coalition $\Pi \setminus \{j\}$ in the two-player zero-sum game $\mathcal{G}_j$ where player j plays against all other players in $\mathcal{G}$, and let $\tau_{i,j}$ be the corresponding strategy of player i in $\mathcal{G}$ (where $\tau_{i,i}$ is an arbitrary strategy). For each player $i \in \Pi$, we define a strategy σ_i in $\mathcal{G}$ as follows:

$$\sigma_i(hv) = \begin{cases} \pi(k+1) & \text{if } hv = \pi(0)\dots\pi(k) \prec \pi, \\ \tau_{i,j}(h_2 v) & \text{otherwise,} \end{cases}$$

where, in the latter case, $h = h_1 h_2$ such that h_1 is the longest prefix of h still being a prefix of π, and j is the player whose turn it was after that prefix (i.e. h_1 ends in V_j), where $j = i$ if $h_1 = \varepsilon$.

Let us show that $(\sigma_i)_{i\in\Pi}$ is a Nash equilibrium of $(\mathcal{G}, v_0)$ with payoff $\bar{x}$. First observe that $\langle(\sigma_i)_{i\in\Pi}\rangle = \pi$, which has payoff $\bar{x}$, thus it remains to show that $(\sigma_i)_{i\in\Pi}$ is a Nash equilibrium. So let us assume that some player $i \in \Pi$

with $x_i = 0$ can improve her payoff by playing according to some strategy σ' instead of σ_i. Then there exists $k < \omega$ such that $\sigma'(\pi(k)) \neq \sigma_i(\pi(k))$, and consequently from this point onwards $\langle (\sigma_j)_{j \in \Pi \setminus \{i\}}, \sigma' \rangle$ is consistent with τ_{-i}, the optimal strategy of the coalition $\Pi \setminus \{i\}$ in $\mathcal{G}_i$. Hence, τ_{-i} is not winning from $\pi(k)$. By determinacy, this implies that $\pi(k) \in W_i$, a contradiction. Q.E.D.

As an immediate consequence, we get that the problem NE is in NP. However, in many cases, we can do better: For two payoff vectors $\bar{x}, \bar{y} \in \{0,1\}^k$, let $\operatorname{dist}(\bar{x}, \bar{y})$ be the *Hamming distance* of $\bar{x}$ and $\bar{y}$, i.e. the number $\sum_{i=1}^{k} |y_i - x_i|$ of nonmatching bits. Jurdziński [13] showed that the problem of deciding whether a vertex is in the winning region for player 0 in a two-player zero-sum parity game is in UP $\cap$ co-UP. Recall that UP is the class of all problems decidable by a nondeterministic Turing machine that runs in polynomial time and has at most one accepting run on every input. We show that the complexity of NE goes down to UP $\cap$ co-UP if the Hamming distance of the thresholds is bounded. If additionally the number of priorities is bounded, the complexity reduces further to P.

Theorem 21 ([28]). NE is in NP. If $\operatorname{dist}(\bar{x}, \bar{y})$ is bounded, NE is in UP $\cap$ co-UP. If additionally the number of priorities is bounded for each player, the problem is in P.

Proof. An NP algorithm for NE works as follows: On input $(\mathcal{G}, v_0)$, the algorithm starts by guessing a payoff $\bar{x} \leq \bar{z} \leq \bar{y}$ and 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 whether there exists a winning play from v_0 in the one-player game arising from $\mathcal{G}$ by merging all players into one, restricting the arena to $G \upharpoonright \bigcap_{z_i=0}(V \setminus W_i)$, and imposing the winning condition $\bigwedge_{z_i=1} \Omega_i \wedge \bigwedge_{z_i=0} \neg\Omega_i$, a Streett condition. If so, the algorithm accepts. Otherwise, the algorithm rejects.

The correctness of the algorithm follows from Lemma 20. For the complexity, note that deciding whether there exists a winning play in a one-player Streett game can be done in polynomial time [10].

If $\operatorname{dist}(\bar{x}, \bar{y})$ is bounded, there is no need to guess the payoff $\bar{z}$. Instead, one can enumerate all of the constantly many payoffs $\bar{x} \leq \bar{z} \leq \bar{y}$ and check for each of them whether there exists a winning play in the respective one-player Streett game. If this is the case for some $\bar{z}$, the algorithm may accept. Otherwise it has to reject. This gives a UP algorithm for NE in the case that $\operatorname{dist}(\bar{x}, \bar{y})$ is bounded. Analogously, a UP algorithm for the complementary problem would accept if for each $\bar{z}$ there exists *no* winning play in the respective one-player Streett game.

For parity games with a bounded number of priorities, winning regions can actually be computed in polynomial time (see e.g. [29]). Thus, if additionally the number of priorities for each player is bounded, the guessing of the winning regions can be avoided as well, so we end up with a deterministic polynomial-time algorithm. Q.E.D.

It is a major open problem whether winning regions of parity games can be computed in polynomial time, in general. This would allow us to decide the problem NE in polynomial time for bounded $\mathrm{dist}(\bar{x}, \bar{y})$ even if the number of priorities is unbounded. Recently, Jurdziński & al. [14] gave a deterministic subexponential algorithm for the problem. It follows that there is a deterministic subexponential algorithm for NE if $\mathrm{dist}(\bar{x}, \bar{y})$ is bounded.

Another line of research is to identify structural properties of graphs that allow for a polynomial-time algorithm for the parity game problem. It was shown that winning regions can be computed in polynomial time for parity games played on graphs of bounded DAG-, Kelly, or clique width [1, 23, 12, 24] (and thus also for graphs of bounded tree width [22] or bounded entanglement [2]). It follows that NE can be decided in polynomial time for games on these graphs if also $\mathrm{dist}(\bar{x}, \bar{y})$ is bounded.

Having shown that NE is in NP, the natural question that arises is whether NE is NP-complete. We answer this question affirmatively. Note that it is an open question whether the parity game problem is NP-complete. In fact, this is rather unlikely, since it would imply that NP = UP = co-UP = co-NP, and hence the polynomial hierarchy would collapse to its first level. As a matter of fact, we show NP-completeness even for the case of games with co-Büchi winning conditions, a class of games known to be solvable in polynomial time in the classical two-player zero-sum case. Also, it suffices to require that only one distinguished player, say the first one, should win in the equilibrium. In essence, this shows that NE is a substantially harder problem than the problem of deciding the existence of a winning strategy for a certain player.

Theorem 22 ([28])**.** NE is NP-complete for co-Büchi games, even with the thresholds $\bar{x} = (1, 0, \ldots, 0)$ and $\bar{y} = (1, \ldots, 1)$.

Proof. By Theorem 21, the problem is in NP. To show that the problem is NP-hard, we give a polynomial-time reduction from SAT. Given a Boolean formula $\varphi = C_1 \wedge \ldots \wedge C_m$ in CNF over variables $X_1, \ldots, X_n$, we build a game $\mathcal{G}_\varphi$ played by players $0, 1, \ldots, n$ as follows. $\mathcal{G}_\varphi$ has vertices $C_1, \ldots, C_m$ controlled by player 0, and for each clause C and each literal X_i or $\neg X_i$, a vertex (C, X_i) or $(C, \neg X_i)$, respectively, controlled by player i. Additionally, there is a sink vertex $\bot$. 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 $\bot$. As $\bot$ is a sink vertex, the only edge leaving $\bot$ leads to $\bot$ itself. For example, Figure 2 shows the essential part of the arena of $\mathcal{G}_\varphi$ for the formula $\varphi = (X_1 \vee X_3 \vee \neg X_2) \wedge (X_3 \vee \neg X_1) \wedge \neg X_3$. The co-Büchi winning conditions are as follows:

- Player 0 wins if the sink vertex is visited only finitely often (or, equivalently, if it is not visited at all);
- Player $i \in \{1 \ldots, n\}$ wins if each vertex (C, X_i) is visited only finitely often.

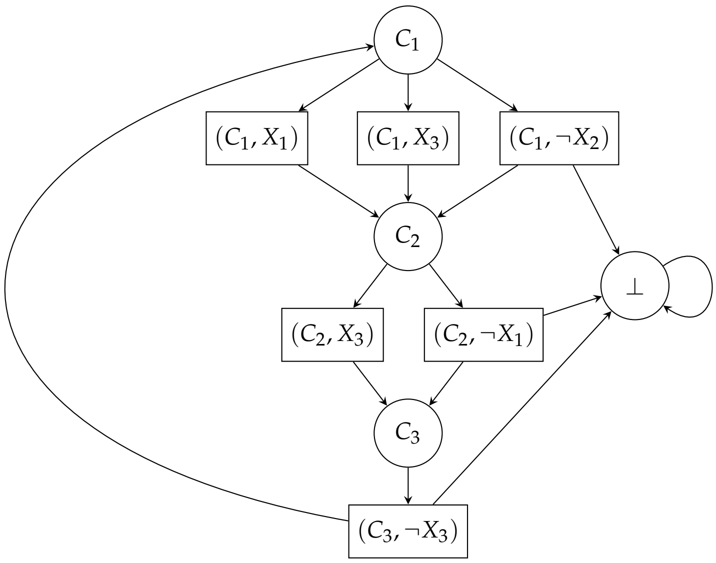

Figure 2. The game $\mathcal{G}_\varphi$ for $\varphi = (X_1 \vee X_3 \vee \neg X_2) \wedge (X_3 \vee \neg X_1) \wedge \neg X_3$

Clearly, $\mathcal{G}_\varphi$ 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.

($\Rightarrow$) Assume that φ is satisfiable. We show that the positional strategy profile where at any time player 0 plays from a clause C to a (fixed) literal that satisfies this clause and each player $j \neq 0$ plays from $\neg X_j$ to the sink if and only if the satisfying interpretation maps X_j to true is a Nash equilibrium where player 0 wins. First note that the induced play never reaches the sink and is therefore won by player 0. 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 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 $\mathcal{G}_\varphi$, 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.

4.2 Subgame Perfect Equilibria

For subgame perfect equilibria, we are not aware of a characterisation like the one in Lemma 20 for Nash equilibria. Therefore, our approach to solve SPE is entirely different from our approach to solve NE. Namely, we reduce SPE to the nonemptiness problem for tree automata (on infinite trees). However, this only gives an ExpTime upper bound for the problem as opposed to NP for the case of Nash equilibria. For the full proof of the following theorem, see [27].

Theorem 23. The problem SPE is in ExpTime. If the number of players and priorities is bounded, the problem is in P.

Proof sketch. Without loss of generality, let us assume that the input game $\mathcal{G}$ is *binary*, i.e. every vertex of $\mathcal{G}$ has at most two successors. Then we can arrange all plays of $(\mathcal{G}, v_0)$ in an infinite binary tree with labels from the vertex set V. Given a strategy profile $(\sigma_i)_{i \in \Pi}$ of $(\mathcal{G}, v_0)$, we enrich this tree with a second label component that takes the value 0 or 1 if the strategy profile prescribes going to the left or right successor, respectively.

The algorithm works as follows: We construct two *alternating parity tree automata*. The first one checks whether some arbitrary tree with labels from the alphabet $V \times \{0, 1\}$ is indeed a tree originating from a strategy profile of $(\mathcal{G}, v_0)$, and the second one checks for a tree originating from a strategy profile $(\sigma_i)_{i \in \Pi}$ of $(\mathcal{G}, v_0)$ whether $(\sigma_i)_{i \in \Pi}$ is a subgame perfect equilibrium with a payoff in between the given thresholds. The first automaton is actually a nondeterministic tree automaton with trivial acceptance (every run of the automaton is accepting) and has $\mathrm{O}(|V|)$ states. The second automaton has $\mathrm{O}(kd)$ states and $\mathrm{O}(1)$ priorities where k is the number of players and d is the maximum number of priorities in a player's parity condition. An equivalent nondeterministic parity tree automaton has $2^{\mathrm{O}(kd \log kd)}$ states and $\mathrm{O}(kd)$ priorities [20]. Finally, we construct the product automaton of the first nondeterministic parity tree automaton with the one constructed from the alternating one. As the former automaton works with trivial acceptance, the construction is straightforward and leads to a nondeterministic parity tree automaton with $\mathrm{O}(|V|) \cdot 2^{\mathrm{O}(kd \log kd)}$ states and $\mathrm{O}(kd)$ priorities. Obviously, the tree language defined by this automaton is nonempty if and only if $(\mathcal{G}, v_0)$ has a subgame perfect equilibrium with a payoff in between the given thresholds. By [9] nonemptiness for nondeterministic parity tree automata can be decided in time polynomial in the number of states and exponential in the number of priorities. Q.E.D.

The exact complexity of SPE remains an open problem. However, NP-hardness can be transferred from NE to SPE. Hence, it is unlikely that there exists a polynomial-time algorithm for SPE, in general.

Theorem 24. SPE is NP-hard for co-Büchi games, even with the thresholds $\bar{x} = (1, 0, \ldots, 0)$ and $\bar{y} = (1, \ldots, 1)$.

Proof. The proof is analogous to the proof of Theorem 22. Just note that the Nash equilibrium of $(\mathcal{G}_\varphi, C_1)$ constructed in the case that φ is satisfiable is also a subgame perfect equilibrium. Q.E.D.

4.3 Secure Equilibria

For secure equilibria we concentrate on two-player games as done by Chatterjee & al. [4], who introduced secure equilibria. If there are only two players, then there are only four possible payoffs for a secure equilibrium: $(0, 0)$, $(1, 0)$, $(0, 1)$, and $(1, 1)$. For each of these payoffs, we aim to characterise the existence of a secure equilibrium that has this payoff and analyse the complexity of deciding whether there exists such an equilibrium..

Lemma 25. Let $(\mathcal{G}, v_0)$ be a two-player game with determined winning conditions. Then $(\mathcal{G}, v_0)$ has a secure equilibrium with payoff $(0, 0)$ if and only if no player has a winning strategy.

Proof. Clearly, if (σ, τ) is a secure equilibrium with payoff $(0, 0)$, then no player can have a winning strategy, since otherwise (σ, τ) would not even be a Nash equilibrium. On the other hand, assume that no player has a winning strategy. By determinacy, there exist a strategy σ of player 1 that guarantees payoff 0 for player 2 and a strategy τ of player 2 that guarantees payoff 1 for player 1. Hence, (σ, τ) is a Nash equilibrium. But it is also secure since every player receives the lowest possible payoff. Q.E.D.

Theorem 26. The problem of deciding whether in a two-player parity game there exists a secure equilibrium with payoff $(0, 0)$ is in UP $\cap$ co-UP. If the number of priorities is bounded, the problem is decidable in polynomial time.

Proof. By Lemma 25, to decide whether there exists a secure equilibrium with payoff $(0, 0)$, one has to decide whether neither player 1 nor player 2 has a winning strategy. For each of the two players, existence (and hence also non-existence) of a winning strategy can be decided in UP $\cap$ co-UP [13]. By first checking whether player 1 does not have a winning strategy and then checking whether player 2 does not have one, we get a UP algorithm for the problem. Analogously, one can deduce that the problem is in co-UP.

If the number of priorities is bounded, deciding the existence of a winning strategy can be done in polynomial time, so we get a polynomial-time algorithm for the problem. Q.E.D.

Lemma 27. Let $(\mathcal{G}, v_0)$ be a two-player game. Then $(\mathcal{G}, v_0)$ has a secure equilibrium with payoff $(1, 0)$ or payoff $(0, 1)$ if and only if player 1 or player 2, respectively, has a strongly winning strategy.

Proof. We only show the claim for payoff $(1,0)$; the proof for payoff $(0,1)$ is completely analogous. Clearly, if σ is a strongly winning strategy for player 1, then (σ,τ) is a secure equilibrium for any strategy τ of player 2. On the other hand, if (σ,τ) is a secure equilibrium with payoff $(1,0)$, then for any strategy τ' of player 2 the strategy profile (σ,τ') has payoff $(1,0)$, hence σ is strongly winning. Q.E.D.

Theorem 28 ([4])**.** The problem of deciding whether in a two-player parity game there exists a secure equilibrium with payoff $(1,0)$, or payoff $(0,1)$, is co-NP-complete. If the number of priorities is bounded, the problem is in P.

Proof. By Lemma 27, deciding whether a two-player parity game has a secure equilibrium with payoff $(1,0)$ or $(0,1)$ amounts to deciding whether player 1 respectively player 2 has a strongly winning strategy. Assume that the game has parity winning conditions Ω_1 and Ω_2. Then player 1 or player 2 has a strongly winning strategy if and only if she has a winning strategy for the condition $\Omega_1 \wedge \neg\Omega_2$ respectively $\Omega_2 \wedge \neg\Omega_1$, a Streett condition. The existence of such a strategy can be decided in co-NP [7]. Hence, the problem of deciding whether the game has a secure equilibrium with payoff $(1,0)$, or $(0,1)$, is also in co-NP.

Chatterjee & al. [5] showed that deciding the existence of a winning strategy in a two-player zero-sum game with the conjunction of two parity conditions as its winning condition is already co-NP-hard. It follows that the problem of deciding whether a player has a strongly winning strategy in a two-player parity game is co-NP-hard.

If the number of priorities is bounded, we arrive at a Streett condition with a bounded number of pairs, for which one can decide the existence of a winning strategy in polynomial time [7], so we get a polynomial-time algorithm. Q.E.D.

Lemma 29. Let $(\mathcal{G}, v_0)$ be a two-player game with a determined pair $(\text{Win}_1, \text{Win}_2)$ of prefix-independent winning conditions. Then $(\mathcal{G}, v_0)$ has a secure equilibrium with payoff $(1,1)$ if and only if there exists a play π with payoff $(1,1)$ such that for all $k < \omega$ no player has a strongly winning strategy in $(\mathcal{G}, \pi(k))$.

Proof. Clearly, if (σ,τ) is a secure equilibrium with payoff $(1,1)$, then $\pi := \langle\sigma,\tau\rangle$ is a play with payoff $(1,1)$ such that for all $k < \omega$ no player has a strongly winning strategy in $(\mathcal{G}, \pi(k))$, since otherwise one player could decrease the other players payoff while keeping her payoff at 1 by switching to her strongly winning strategy at vertex $\pi(k)$.

Assume that there is a play π with payoff $(1,1)$ such that for all $k < \omega$ no player has a strongly winning strategy in $(\mathcal{G}, \pi(k))$. By determinacy, there exists a strategy σ_1 of player 1 and a strategy τ_1 of player 2 such that σ_1 and τ_1 are retaliating strategies in $(\mathcal{G}, \pi(k))$ for each $k < \omega$. Similarly to the proof of Lemma 20, we define a new strategy σ of player 1 for $(\mathcal{G}, v_0)$ by

$$\sigma(hv) = \begin{cases} \pi(k+1) & \text{if } hv = \pi(0)\ldots\pi(k) \prec \pi, \\ \sigma_1(h_2 v) & \text{otherwise.} \end{cases}$$

where in the latter case $h = h_1 \cdot h_2$, and h_1 is the longest prefix of h still being a prefix of π. Analogously, one can define a corresponding strategy τ of player 2 for $(\mathcal{G}, v_0)$. It follows that the strategy profile (σ, τ) has payoff $(1, 1)$, and for each strategy σ' of player 1 and each strategy τ' of player 2 the strategy profiles (σ', τ) and (σ, τ') still give payoff 1 to player 2 respectively player 1. Hence, (σ, τ) is a secure equilibrium. Q.E.D.

Theorem 30 ([4])**.** The problem of deciding whether in a two-player parity game there exists a secure equilibrium with payoff $(1, 1)$ is in NP. If the number of priorities is bounded, the problem is in P.

Proof. By Lemma 29, to decide whether there exists a secure equilibrium with payoff $(1, 1)$, one has to decide whether there exists a play that has payoff $(1, 1)$ and remains inside the set U of vertices where no player has a strongly winning strategy. By determinacy, the set U equals the set of vertices where both players have retaliating strategies. Assume that the game has parity winning conditions Ω_1 and Ω_2. Then a retaliating strategy of player 1 or player 2 corresponds to a winning strategy for the condition $\Omega_1 \vee \neg\Omega_2$ respectively $\Omega_2 \vee \neg\Omega_1$, a Rabin condition. Since positional strategies suffice to win a two-player zero-sum game with a Rabin winning condition [15], this implies that the set U also equals the set of vertices where both players have positional retaliating strategies.

An NP algorithm for deciding whether there exists a secure equilibrium with payoff $(1, 1)$ works as follows: First, the algorithm guesses a set X together with a positional strategy σ of player 1 and a positional strategy τ of player 2. Then, the algorithm checks whether σ and τ are retaliating strategies from each vertex $v \in X$. If this is the case, the algorithm checks whether there exists a play with payoff $(1, 1)$ remaining inside X. If so, the algorithm accepts, otherwise it rejects.

The correctness of the algorithm is immediate. For the complexity, note that checking whether a positional strategy of player 1 or 2 is a retaliating strategy amounts to deciding whether the other player has a winning strategy for the condition $\Omega_2 \wedge \neg\Omega_1$ respectively $\Omega_1 \wedge \neg\Omega_2$, again a Streett condition, in the one-player game where the transitions of player 1 respectively player 2 have been fixed according to her positional strategy. Also, checking whether there exists a play with payoff $(1, 1)$ remaining inside X amounts to deciding whether there exists a winning play in a one-player Streett game, namely the one derived from $\mathcal{G}$ by removing all vertices in X, merging the two players into one, and imposing the winning condition $\Omega_1 \wedge \Omega_2$. As the problem of deciding the existence of a winning play in a one-player Streett game is decidable in polynomial time, our algorithm runs in (nondeterministic) polynomial time.

If the number of priorities is bounded, we can actually compute the set U of vertices from where both players have a retaliating strategy in polynomial time, so the algorithm can be made deterministic while retaining a polynomial running time. Q.E.D.

References

1. 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.
2. D. Berwanger & E. Grädel. Entanglement – a measure for the complexity of directed graphs with applications to logic and games. In *Proceedings of the 11th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2004*, vol. 3452 of *LNCS*, pp. 209–223. Springer-Verlag, 2005.
3. J. R. Büchi & L. H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
4. 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.
5. K. Chatterjee, T. A. Henzinger & N. Piterman. Generalized parity games. In *Proceedings of the 10th International Conference on Foundations of Software Science and Computation Structures, FOSSACS 2007*, vol. 4423 of *LNCS*, pp. 153–167. Springer-Verlag, 2007.
6. 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.
7. 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.
8. E. A. Emerson & C. S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science, FoCS '91*, pp. 368–377. IEEE Computer Society Press, 1991.
9. E. A. Emerson, C. S. Jutla & A. P. Sistla. On model-checking for fragments of μ-calculus. In *Proceedings of the 5th International Conference on Computer Aided Verification, CAV '93*, vol. 697 of *LNCS*, pp. 385–396. Springer-Verlag, 1993.
10. 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.
11. E. Grädel. Finite model theory and descriptive complexity. In *Finite Model Theory and Its Applications*, pp. 125–230. Springer-Verlag, 2007.
12. 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.

13. M. Jurdziński. Deciding the winner in parity games is in UP ∩ co-UP. *Information Processing Letters*, 68(3):119–124, 1998.
14. M. Jurdziński, M. Paterson & U. Zwick. A deterministic subexponential algorithm for solving parity games. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms, SODA 2006*, pp. 117–123. ACM Press, 2006.
15. 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.
16. H. W. Kuhn. Extensive games and the problem of information. In *Contributions to the Theory of Games II*, vol. 28 of *Annals of Mathematical Studies*, pp. 193–216. Princeton University Press, 1953.
17. D. A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
18. D. A. Martin. An extension of Borel determinacy. *Annals of Pure and Applied Logic*, 49(3):279–293, 1990.
19. A. W. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
20. D. E. Muller & P. E. Schupp. Simulating alternating tree automata by nondeterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra. *Theoretical Computer Science*, 141(1–2):69–107, 1995.
21. J. F. Nash Jr. Equilibrium points in N-person games. *Proceedings of the National Academy of Sciences of the USA*, 36:48–49, 1950.
22. J. Obdržálek. Fast mu-calculus model checking when tree-width is bounded. In *Proceedings of the 15th International Conference on Computer Aided Verification, CAV 2003*, vol. 2725 of *LNCS*, pp. 80–92. Springer-Verlag, 2003.
23. 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.
24. 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.
25. R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301–324 and 667–689, 1965.
26. W. Thomas. On the synthesis of strategies in infinite games. In *Proceedings of the 12th Annual Symposium on Theoretical Aspects of Computer Science, STACS '95*, vol. 900 of *LNCS*, pp. 1–13. Springer-Verlag, 1995.
27. 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.
28. M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proceedings of the 11th International Conference on Foundations of Software Science and Computation Structures, FOSSACS 2008*, vol. 4962 of *LNCS*, pp. 20–34. Springer-Verlag, 2008.
29. W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1–2):135–183, 1998.