Two-player Nonzero-sum ω-regular Games

Krishnendu Chatterjee

Dept. of EECS, University of California, Berkeley
`c_krish@eecs.berkeley.edu`

Abstract. We study infinite stochastic games played by two-players on a finite graph with goals specified by sets of infinite traces. The games are *concurrent* (each player simultaneously and independently chooses an action at each round), *stochastic* (the next state is determined by a probability distribution depending on the current state and the chosen actions), *infinite* (the game continues for an infinite number of rounds), *nonzero-sum* (the players' goals are not necessarily conflicting), and undiscounted. We show that if each player has an ω-regular objective expressed as a parity objective, then there exists an ε-Nash equilibrium, for every $\varepsilon > 0$. However, exact Nash equilibria need not exist. We study the complexity of finding values (payoff profile) of an ε-Nash equilibrium. We show that the values of an ε-Nash equilibrium in nonzero-sum concurrent parity games can be computed by solving the following two simpler problems: computing the values of zero-sum (the goals of the players are strictly conflicting) concurrent parity games and computing ε-Nash equilibrium values of nonzero-sum concurrent games with reachability objectives. As a consequence we establish that values of an ε-Nash equilibrium can be computed in TFNP (total functional NP), and hence in EXPTIME.

1 Introduction

Stochastic games. Non-cooperative games provide a natural framework to model interactions between agents [12, 14]. The simplest class of non-cooperative games consists of the "one-step" games — games with single interaction between the agents after which the game ends and the payoffs are decided (e.g., matrix games). However, a wide class of games progress over time and in stateful manner, and the current game depends on the history of interactions. Infinite *stochastic games* [16, 8] are a natural model for such games. A stochastic game is played over a finite *state space* and is played in rounds. In concurrent games, in each round, each player chooses an action from a finite set of available actions, simultaneously and independently of other players. The game proceeds to a new state according to a probabilistic transition relation (stochastic transition matrix) based on the current state and the joint actions of the players. Concurrent games subsume the simpler class of *turn-based games*, where at every state at most one player can choose between multiple actions. In verification and control of finite state reactive systems such games proceed for infinite rounds, generating an infinite sequence of states, called the *outcome* of the game. The players receive a payoff based on a payoff function that maps every outcome to a real number.

Objectives. Payoffs are generally Borel measurable functions [11]. The payoff set for each player is a Borel set B_i in the Cantor topology on S^ω (where S is the set of states), and player i gets payoff 1 if the outcome of the game is in B_i, and 0 otherwise. In verification, payoff functions are usually index sets of *ω-regular languages*. The ω-regular languages generalize the classical regular languages to infinite strings, they occur in low levels of the Borel hierarchy (they are in $\Sigma_3 \cap \Pi_3$), and they form a robust and expressive language for determining payoffs for commonly used specifications. The simplest ω-regular objectives correspond to safety ("closed sets") and reachability ("open sets") objectives.

Zero-sum games. Games may be *zero-sum*, where two players have directly conflicting objectives and the payoff of one player is one minus the payoff of the other, or *nonzero-sum*, where each player has a prescribed payoff function based on the outcome of the game. The fundamental question for games is the existence of equilibrium values. For zero-sum games, this involves showing a *determinacy* theorem that states that the expected optimum value obtained by player 1 is exactly one minus the expected optimum value obtained by player 2. For one-step zero-sum games, this is von Neumann's minmax theorem [21]. For infinite games, the existence of such equilibria is not obvious, in fact, by using the axiom of choice, one can construct games for which determinacy does not hold. However, a remarkable result by Martin [11] shows that all stochastic zero-sum games with Borel payoffs are determined.

Nonzero-sum games. For nonzero-sum games, the fundamental equilibrium concept is a *Nash equilibrium* [10], that is, a strategy profile such that no player can gain by deviating from the profile, assuming the other player continues playing the strategy in the profile. Again, for one-step games, the existence of such equilibria is guaranteed by Nash's theorem [10]. However, the existence of Nash equilibria in infinite games is not immediate: Nash's theorem holds for finite bimatrix games, but in case of stochastic games, the strategy space is not compact. The existence of Nash equilibria is known only in very special cases of stochastic games. In fact, Nash equilibria may not exist, and the best one can hope for is an ε-Nash equilibrium for all $\varepsilon > 0$, where an ε-Nash equilibrium is a strategy profile where unilateral deviation can only increase the payoff of a player by at most ε. Exact Nash equilibria do exist in discounted stochastic games [9]. For concurrent nonzero-sum games with payoffs defined by Borel sets, surprisingly little is known. Secchi and Sudderth [15] showed that exact Nash equilibria do exist when all players have payoffs defined by closed sets ("safety objectives"). In the case of open sets ("reachability objectives"), the existence of ε-Nash equilibrium for every $\varepsilon > 0$, has been established in [5]. The above results hold even in the case of n-player games. In an important recent result in stochastic game theory, Vieille shows the existence of ε-Nash equilibrium, for every $\varepsilon > 0$, in two-player nonzero-sum concurrent games with limit-average payoff [19, 20]. The existence of ε-Nash equilibrium in two-player concurrent games with objectives in higher levels of Borel hierarchy has been an intriguing open problem.

Result and proof techniques. In this paper we show that ε-Nash equilibrium exists, for every $\varepsilon > 0$, for two-player concurrent games with ω-regular objectives. However, exact Nash equilibria need not exist. For two-player concurrent games our result extends the existence of ε-Nash equilibrium from the lowest level of Borel hierarchy (open and closed sets) to ω-regular objectives that lie in the higher levels of Borel hierarchy; and our result for ω-regular objectives parallels Vieille's result for limit-average objectives. Our result is organized as follows:

1. In Section 3 we first show the existence of ε-Nash equilibrium, for every $\varepsilon > 0$, for a sub-class of concurrent games, namely single strongly connected component (Sscc) games, with ω-regular objectives.
2. We extend the above result to all concurrent games in Section 4.

The result for Sscc games involves the following key ideas:

- We identify four sufficient conditions that ensure existence of ε-Nash equilibrium, for every $\varepsilon > 0$, in Sscc games.
- We then show that if the sufficient conditions are not satisfied, then the game can be reduced to a nonzero-sum game with reachability objectives, with some desired properties. The result is proved by generalizing a result from [2] and using a fragment of analysis of Vieille [19].
- The existence of ε-Nash equilibrium, for all $\varepsilon > 0$, in the original game is then established by the use of *punishing or spoiling strategies*.

Complexity of ε-Nash equilibrium. Computing the values of a Nash equilibria, when it exists, is another challenging problem [13, 22]. For one-step zero-sum games, equilibrium values and strategies can be computed in polynomial time (by reduction to linear programming) [12]. For one-step nonzero-sum games, no polynomial time algorithm is known to compute an exact Nash equilibrium in two-player games [13]. In case of zero-sum concurrent games with ω-regular objectives several algorithms are known to compute values with in ε-approximation [7, 2]. Since the values can be irrational, ε-approximation is the best one can achieve. From the computational aspects, a desirable property of an existence proof of Nash equilibrium is its ease of algorithmic analysis. We show that our proof for existence of ε-Nash equilibrium is completely constructive and algorithmic. Our proof shows that the computation of values of an ε-Nash equilibrium in two-player concurrent games with parity objectives can be reduced to the following two simpler problems:

1. Computing values of zero-sum concurrent games with parity objectives.
2. Computing values of some special ε-Nash equilibrium of nonzero-sum concurrent games with reachability objectives.

Since zero-sum games are special cases of nonzero-sum games, computing ε-Nash equilibrium in nonzero-sum games are at least as hard as solving the optimum values in zero-sum games. Our result shows that the extra cost of computing ε-Nash equilibrium for ω-regular objectives is no more than solving some special ε-Nash equilibrium of games with reachability objectives. We then prove that the equilibrium values of an ε-Nash equilibrium can be computed in TFNP (total functional NP) and hence in EXPTIME. Our result matches the best known complexity bound for the simpler case of turn-based games [5].

2 Definitions

Notation. For a countable set A, a *probability distribution* on A is a function $\delta : A \mapsto [0,1]$ such that $\sum_{a \in A} \delta(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. Given a distribution $\delta \in \mathcal{D}(A)$, we denote by $\mathrm{Supp}(\delta) = \{x \in A \mid \delta(x) > 0\}$ the *support* of δ.

Definition 1 (Concurrent games). *A (two-player)* concurrent game structure *$\mathcal{G} = \langle S, Moves, \Gamma_1, \Gamma_2, \delta \rangle$ consists of the following components:*

- *A finite state space S and a finite set Moves of moves.*
- *Two move assignments $\Gamma_1, \Gamma_2 : S \mapsto 2^{Moves} \setminus \emptyset$. For $i \in \{1,2\}$, assignment Γ_i associates with each state $s \in S$ the non-empty set $\Gamma_i(s) \subseteq Moves$ of moves available to player i at state s.*
- *A probabilistic transition function $\delta : S \times Moves \times Moves \to \mathcal{D}(S)$, that gives the probability $\delta(s, a_1, a_2)(t)$ of a transition from s to t when player 1 plays a_1 and player 2 plays a_2, for all $s, t \in S$ and $a_1 \in \Gamma_1(s)$, $a_2 \in \Gamma_2(s)$.* ∎

A special class of concurrent game structures are Markov decision processes (MDPs). A concurrent game structure is a Markov decision process (MDP) if there exists an $i \in \{1,2\}$ such that at every state s, $|\Gamma_i(s)| = 1$. In other words, MDPs are one-player stochastic games: only one player has a non-trivial choice of moves and for the other player the choice of the moves are fixed.

We define the size of the game structure $\mathcal{G}$ to be equal to the size of the transition function δ; specifically, $|\mathcal{G}| = \sum_{s \in S} \sum_{a \in \Gamma_1(s)} \sum_{b \in \Gamma_2(s)} \sum_{t \in S} |\delta(s,a,b)(t)|$, where $|\delta(s,a,b)(t)|$ denotes the space to specify the probability distribution. At every state $s \in S$, player 1 chooses a move $a_1 \in \Gamma_1(s)$, and simultaneously and independently player 2 chooses a move $a_2 \in \Gamma_2(s)$. The game then proceeds to the successor state t with probability $\delta(s, a_1, a_2)(t)$, for all $t \in S$. A state s is called an *absorbing state* if for all $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$ we have $\delta(s, a_1, a_2)(s) = 1$. In other words, at s for all choices of moves of the players the next state is always s. Each player chooses her strategy independently and secretly from the other player, and is only interested in maximizing her own payoff. For all states $s \in S$ and moves $a_1 \in \Gamma_1(s)$ and $a_2 \in \Gamma_2(s)$, we indicate by $\mathrm{Dest}(s, a_1, a_2) = \mathrm{Supp}(\delta(s, a_1, a_2))$ the set of possible successors of s when moves a_1, a_2 are selected.

A *path* or a *play* ω of $\mathcal{G}$ is an infinite sequence $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ of states in S such that for all $k \geq 0$, there are moves $a_1^k \in \Gamma_1(s_k)$ and $a_2^k \in \Gamma_2(s_k)$ with $\delta(s_k, a_1^k, a_2^k)(s_{k+1}) > 0$. We denote by Ω the set of all paths and by Ω_s the set of all paths $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ such that $s_0 = s$, i.e., the set of plays starting from state s.

Randomized strategies. A *selector* ξ for player $i \in \{ 1, 2 \}$ is a function $\xi : S \mapsto \mathcal{D}(\mathit{Moves})$ such that for all $s \in S$ and $a \in \mathit{Moves}$, if $\xi(s)(a) > 0$ then $a \in \Gamma_i(s)$. We denote by Λ_i the set of all selectors for player $i \in \{1, 2\}$. A *strategy* for player 1 is a function $\sigma : S^+ \to \Lambda_1$ that associates with every finite non-empty sequence of states, representing the history of the play so far, a selector. Similarly we define strategies π for player 2. A memoryless strategy is independent of the history of the play and depends only on the current state. Memoryless strategies coincide with selectors, and we often write σ for the selector corresponding to a memoryless strategy σ. We denote by Σ and Π the set of all strategies for player 1 and player 2, respectively.

Once the starting state s and the strategies σ and π for the two players have been chosen, the game is reduced to an ordinary stochastic process. Hence, the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega_s$ is a measurable set of paths. For an event $\mathcal{A} \subseteq \Omega_s$, we denote by $\Pr_s^{\sigma,\pi}(\mathcal{A})$ the probability that a path belongs to $\mathcal{A}$ when the game starts from s and the players follow the strategies σ and π.

Objectives. An *objective* for a player in a game $\mathcal{G}$ is a set $\mathcal{W} \subseteq \Omega$ of infinite paths. We consider the following objectives.

- *Reachability objective.* For a set $R \subseteq S$ of *target* states, the Reachability objective is defined as $\mathrm{Reach}(R) = \{ \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid \exists k \in \mathbb{N}.\ s_k \in R \}$.
- *Safety objective.* For a set $F \subseteq S$ of *safe* states, the Safety objective is defined as $\mathrm{Safe}(F) = \{ \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid \forall k \in \mathbb{N}.\ s_k \in F \}$. Note that $\Omega \setminus \mathrm{Reach}(R) = \mathrm{Safe}(S \setminus R)$. Hence the reachability objective with target set R is complementary to the safety objective with safe set $S \setminus R$.
- *Parity objective.* Given $d \in \mathbb{N}$, we write $[d]$ for the set $\{ 0, 1, 2, \ldots, d \}$ and $[d]_+$ for the set $\{ 1, 2, \ldots, d \}$. Let $p : S \mapsto [d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. For an infinite path $\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega$, we define $\mathrm{Inf}(\omega) = \{ i \in [d] \mid p(s_k) = i \text{ for infinitely many } k \geq 0 \}$. The *parity objective* is defined as $\mathrm{Parity}(p) = \{ \omega \in \Omega \mid \min\big(\mathrm{Inf}(\omega)\big) \text{ is even} \}$. Informally we say that a path ω satisfies the parity objective, $\mathrm{Parity}(p)$, if $\omega \in \mathrm{Parity}(p)$.

The ability to solve games with parity objectives suffices for solving games with arbitrary ω-regular objectives, since every ω-regular objective can be specified as a parity objective [17].

A concurrent nonzero-sum parity game consists of a game structure $\mathcal{G}$ and two priority functions p_1 and p_2 for player 1 and player 2, respectively. The objectives of player 1 and player 2 are $\mathrm{Parity}(p_1)$ and $\mathrm{Parity}(p_2)$, respectively. We write Ψ for an arbitrary parity objective. We write the objectives of player 1 and player 2 as Ψ_1 and Ψ_2, respectively, where Ψ_1 and Ψ_2 are arbitrary ω-regular objectives formalized as parity objectives. We also use Ψ_1 to denote the set of paths $\omega \in \Omega$ such that $\omega \in \mathrm{Parity}(p_1)$. Similarly we write Ψ_2 to denote the set of paths $\mathrm{Parity}(p_2)$. Given a state s we write Ψ_{1s} to denote $\Omega_s \cap \Psi_1$ and similarly we write Ψ_{2s} to denote $\Omega_s \cap \Psi_2$. We also write Ψ_s to denote $\Omega_s \cap \Psi$. Given a parity objective Ψ, the set of paths Ψ_s is measurable for any choice of strategies for the two players [18]. Hence, the probability that a path satisfies objective Ψ starting from state $s \in S$ under strategies σ, π for the two players is $\Pr_s^{\sigma,\pi}(\Psi_s)$.

Concurrent nonzero-sum games. A concurrent nonzero-sum game consists of a concurrent game structure $\mathcal{G}$ and objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. A concurrent game is zero-sum if the objectives of the players are complementary, i.e., $\Psi_1 = \Omega \setminus \Psi_2$. The zero-sum values for the players in concurrent games with objectives Ψ_1 and Ψ_2, for player 1 and player 2, respectively, are defined as follows.

Definition 2 (Zero-sum values). *Let $\mathcal{G}$ be a concurrent game structure with objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. Given a state $s \in S$ we call the maximal probability with which player 1 can ensure that Ψ_1 holds from s against any strategy of player 2 is the* zero-sum value *of player 1 at s. The zero-sum value for player 2 is defined symmetrically. Formally, the zero-sum value for player 1 and player 2 are given by functions $\langle\langle 1\rangle\rangle_{val}(\Psi_1) : S \mapsto [0,1]$ and $\langle\langle 2\rangle\rangle_{val}(\Psi_2) : S \mapsto [0,1]$, defined for all $s \in S$ by*

$$\langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) = \sup_{\sigma\in\Sigma} \inf_{\pi\in\Pi} \Pr_s^{\sigma,\pi}(\Psi_{1s}); \qquad \langle\langle 2\rangle\rangle_{val}(\Psi_2)(s) = \sup_{\pi\in\Pi} \inf_{\sigma\in\Sigma} \Pr_s^{\sigma,\pi}(\Psi_{2s}). \blacksquare$$

Concurrent zero-sum games satisfy a *quantitative* version of determinacy [11], stating that for all parity objectives Ψ_1 and Ψ_2, such that $\Psi_1 = \Omega \setminus \Psi_2$, and all $s \in S$, we have $\langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) + \langle\langle 2\rangle\rangle_{val}(\Psi_2)(s) = 1$. A strategy σ for player 1 is *optimal* with respect to objective Ψ, if for all $s \in S$ we have $\inf_{\pi\in\Pi} \Pr_s^{\sigma,\pi}(\Psi_s) = \langle\langle 1\rangle\rangle_{val}(\Psi)(s)$. For $\varepsilon > 0$, a strategy σ for player 1 is *ε-optimal* with respect to objective Ψ, if for all $s \in S$ we have $\inf_{\pi\in\Pi} \Pr_s^{\sigma,\pi}(\Psi_s) \geq \langle\langle 1\rangle\rangle_{val}(\Psi)(s) - \varepsilon$. We define optimal and ε-optimal strategies for player 2 symmetrically. Note that the quantitative determinacy of concurrent zero-sum games is equivalent to the existence of ε-optimal strategies for both players for all $\varepsilon > 0$, at all states $s \in S$.

Definition 3 (Cooperative value). *Given a concurrent game structure $\mathcal{G}$ and an objective Ψ we define the* cooperative *value at a state $s \in S$ as the maximal probability with which player 1 and player 2 can cooperate to satisfy the objective Ψ at s. Formally, the cooperative value is given by the function $\langle\langle 1,2\rangle\rangle_{val}(\Psi) : S \mapsto [0,1]$, defined for all $s \in S$ by $\langle\langle 1,2\rangle\rangle_{val}(\Psi)(s) = \sup_{(\sigma,\pi)\in\Sigma\times\Pi} \Pr_s^{\sigma,\pi}(\Psi_s)$.* $\blacksquare$

Note that the computation of the cooperative value function $\langle\langle 1,2\rangle\rangle_{val}(\Psi)$ can be interpreted as the computation of a value function in a MDP with objective Ψ, where player 1 and player 2 cooperatively choose strategies.

Definition 4 (ε-Nash equilibrium). *Let $\mathcal{G}$ be a concurrent game structure and let the objectives for player 1 and player 2 be Ψ_1 and Ψ_2, respectively. For $\varepsilon \geq 0$, a strategy profile $(\sigma^*, \pi^*) \in \Sigma \times \Pi$ is an ε-Nash equilibrium for a state $s \in S$ iff the following two conditions hold:*

$$\sup_{\sigma\in\Sigma} \Pr_s^{\sigma,\pi^*}(\Psi_{1s}) \leq \Pr_s^{\sigma^*,\pi^*}(\Psi_{1s}) + \varepsilon; \qquad \sup_{\pi\in\Pi} \Pr_s^{\sigma^*,\pi}(\Psi_{2s}) \leq \Pr_s^{\sigma^*,\pi^*}(\Psi_{2s}) + \varepsilon.$$

An exact Nash equilibrium is an ε-Nash equilibrium with $\varepsilon = 0$. $\blacksquare$

It may be noted that in case of zero-sum concurrent games with parity objectives optimal strategies need not exist, and only existence of ε-optimal strategies can be guaranteed, for all $\varepsilon > 0$ [6]. Hence in the general case of nonzero-sum concurrent games with parity objectives Nash equilibrium need not exist, and existence of ε-Nash equilibrium, for all $\varepsilon > 0$, is the best one can achieve.

Definition 5 (ε-optimal and ε-spoiling strategies). *Given a nonzero-sum concurrent game with objective Ψ_1 for player 1 and Ψ_2 for player 2, a strategy σ_ε is ε-optimal if it is ε-optimal with respect to objective Ψ_1, and a strategy $\overline{\sigma}_\varepsilon$ is ε-spoiling if it is ε-optimal with respect to objective $\overline{\Psi_2} = \Omega \setminus \Psi_2$. The ε-optimal and ε-spoiling strategies for player 2 are defined similarly. We denote by Σ_ε and Π_ε the set of ε-optimal strategies for player 1 and player 2, respectively. Similarly, we denote by $\overline{\Sigma_\varepsilon}$ and $\overline{\Pi_\varepsilon}$ the set of ε-spoiling strategies for player 1 and player 2, respectively.* ∎

The determinacy of concurrent games with parity objectives ensures that for all $\varepsilon > 0$, the sets $\Sigma_\varepsilon, \Pi_\varepsilon, \overline{\Sigma_\varepsilon}$ and $\overline{\Pi_\varepsilon}$ are non-empty.

3 Sscc Games

In this section we prove the existence of ε-Nash equilibrium for all $\varepsilon > 0$, in a subclass of concurrent games, namely, single strongly connected component games. In the next section we generalize the existence of ε-Nash equilibrium, for all $\varepsilon > 0$, to all concurrent games using the result of this section. Given a game structure $\mathcal{G}$ we define a underlying graph $G_{\mathcal{G}}$ of $\mathcal{G}$.

Definition 6 (Graph of a game structure $\mathcal{G}$). *Given a concurrent game structure $\mathcal{G} = \langle S, Moves, \Gamma_1, \Gamma_2, \delta \rangle$ the graph of the game structure $\mathcal{G}$ is a directed graph $G_{\mathcal{G}} = (S_{\mathcal{G}}, E_{\mathcal{G}})$ that is defined as follows:*

- *$S_{\mathcal{G}} = S$, i.e., the set of states of $G_{\mathcal{G}}$ is same as the state space of $\mathcal{G}$.*
- *$E_{\mathcal{G}} = \{ (s,t) \mid \exists\, a_1 \in \Gamma_1(s), \exists\, a_2 \in \Gamma_2(s).\ t \in \mathrm{Dest}(s, a_1, a_2) \}$.* ∎

Definition 7 (Single strongly connected component (Sscc) game structures). *Let $\mathcal{G}$ be a concurrent game structure with parity objectives $\Psi_1 = Parity(p_1)$ and $\Psi_2 = Parity(p_2)$ for player 1 and player 2, respectively. Let $G_{\mathcal{G}}$ be the graph of $\mathcal{G}$. We call $\mathcal{G}$ a single strongly connected component (Sscc) game structure if the graph $G_{\mathcal{G}}$ satisfy the following conditions:*

- *The state space $S_{\mathcal{G}}$ can be partitioned into three sets: C, U, T, with $T = \{ t_{00}, t_{10}, t_{01}, t_{11} \}$.*
- *C is a strongly connected component in the graph $G_{\mathcal{G}}$.*
- *The states $t_{ij} \in T$ are absorbing states, for $i, j \in \{ 0, 1 \}$. The priority function for the states in T are as follows: $p_1(t_{ij}) = i$ and $p_2(t_{ij}) = j$, for $i, j \in \{ 0, 1 \}$. Note that at state t_{00} objective of both the players are satisfied; at state t_{01} only player 1's objective is satisfied; at state t_{10} only player 2's objective is satisfied and at state t_{11} none of the players objective is satisfied.*

– *For every state $s \in U$ we have $|\Gamma_i(s)| = 1$ for $i \in \{1,2\}$ and $(\{s\} \times S_{\mathcal{G}}) \cap E_{\mathcal{G}} \subseteq \{s\} \times T$. In other words, at states in U there is no non-trivial choice of moves for the players and thus for any state s in U the game proceeds to the set T according to the probability distribution of the transition function δ at s.*
– *$C \times (S_{\mathcal{G}} \setminus C) \cap E_{\mathcal{G}} \subseteq C \times U$, i.e., the edges out of C end at a state in U. We also require that $(C \times U) \cap E_{\mathcal{G}} \neq \emptyset$.* ■

Reduction gadget. Let $\mathcal{G}$ be a Sscc game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. Suppose for all $\varepsilon > 0$, there is an ε-Nash equilibrium $(\sigma^*_\varepsilon, \pi^*_\varepsilon)$ at s, with $x_1(s) = \lim_{\varepsilon \to 0} \Pr_s^{\sigma^*_\varepsilon, \pi^*_\varepsilon}(\Psi_{1s})$ and $x_2(s) = \lim_{\varepsilon \to 0} \Pr_s^{\sigma^*_\varepsilon, \pi^*_\varepsilon}(\Psi_{2s})$. Consider the gadget $\mathrm{gad}(s)$ to replace s as follows:

– Without loss of generality let $x_1(s) \leq x_2(s)$ (when $x_2(s) \leq x_1(s)$ the gadget is symmetric). Then gadget to replace s is as follows: $\Gamma_1(s) = \{ a \}$, $\Gamma_2(s) = \{ b \}$, and

$$\begin{array}{ll} \delta(s,a,b)(t_{00}) = x_1(s), & \delta(s,a,b)(t_{10}) = x_2(s) - x_1(s), \\ \delta(s,a,b)(t_{11}) = 1 - x_2(s), & \delta(s,a,b)(t_{01}) = 0, \end{array}$$

where t_{ij} are as defined in Definition 7.

The construction ensures that at state s the set $\{ t_{00}, t_{01} \}$ of states is reached with probability $x_1(s)$, i.e., player 1's objective is satisfied with probability $x_1(s)$, and the set $\{ t_{00}, t_{10} \}$ of states is reached with probability $x_2(s)$, i.e., player 2's objective is satisfied with probability $x_2(s)$.

Proposition 1 states that if existence of ε-Nash equilibrium is established at a state s, then state s can be replaced by the gadget $\mathrm{gad}(s)$ and to prove existence of ε-Nash equilibrium in the original game it suffices to prove existence of ε-Nash equilibrium in the transformed game with the gadget $\mathrm{gad}(s)$ replacing state s.

Proposition 1. *Let $\mathcal{G}$ be a Sscc game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. Suppose for every $\varepsilon > 0$, there is an ε-Nash equilibrium $(\sigma^*_\varepsilon, \pi^*_\varepsilon)$ at s, with $x_1(s) = \lim_{\varepsilon \to 0} \Pr_s^{\sigma^*_\varepsilon, \pi^*_\varepsilon}(\Psi_{1s})$ and $x_2(s) = \lim_{\varepsilon \to 0} \Pr_s^{\sigma^*_\varepsilon, \pi^*_\varepsilon}(\Psi_{2s})$. The game structure $\mathcal{G}$ can be transformed to a game structure $\mathcal{G}'$ by replacing the state s with the gadget $\mathrm{gad}(s)$ such that if there is an ε-Nash equilibrium in the transformed game structure $\mathcal{G}'$ for all states in $\mathcal{G}'$, for all $\varepsilon > 0$, then there is an ε-Nash equilibrium in the original game structure $\mathcal{G}$ for all states in $\mathcal{G}$, for all $\varepsilon > 0$.* ■

The result follows from the observation that player 1 and player 2 can switch to strategies $(\sigma^*_\varepsilon, \pi^*_\varepsilon)$ when the game reaches s.

Four properties (P1-P4). Let $\mathcal{G}$ be a Sscc game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. We define four properties (P1-P4) for a state $s \in C$ as follows:

(P1) $\langle\langle 1,2 \rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) = 1$; **(P2)** $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 1$;
(P3) $\langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s) = 1$; **(P4)** $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 0$ and $\langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s) = 0$.

Lemma 1. *Let $\mathcal{G}$ be a* Sscc *game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. If any of the four properties (P1-P4) hold for a state $s \in C$, then for every $\varepsilon > 0$, there is an ε-Nash equilibrium (σ^*, π^*) for the state $s \in C$.*

Proof. 1. Suppose there is a state $s \in C$ such that $\langle\langle 1, 2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) = 1$, then there is a strategy profile (σ^*, π^*) such that $\Pr_s^{\sigma^*,\pi^*}(\Psi_{1s}) = 1$ and $\Pr_s^{\sigma^*,\pi^*}(\Psi_{2s}) = 1$. Since 1 is the maximum payoff a player can achieve, clearly (σ^*, π^*) is a Nash equilibrium at s.

2. Suppose there is a state $s \in C$ such that $\langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) = 1$. Then for every $\varepsilon > 0$, there is an ε-optimal strategy σ_ε for player 1 such that $\inf_{\pi \in \Pi} \Pr_s^{\sigma_\varepsilon,\pi}(\Psi_{1s}) \geq 1 - \varepsilon$. Consider a strategy π^* such that $\Pr_s^{\sigma_\varepsilon,\pi^*}(\Psi_{2s}) \geq \sup_{\pi \in \Pi} \Pr_s^{\sigma_\varepsilon,\pi}(\Psi_{2s}) - \varepsilon$. In other words, we fix an ε-optimal strategy σ_ε for player 1 and a strategy π^* for player 2 that ensures player 2 the maximal probability to satisfy Ψ_2 against the strategy σ_ε, within ε-precision. Thus we have $\sup_{\sigma \in \Sigma} \Pr_s^{\sigma,\pi^*}(\Psi_{1s}) \leq 1 \leq \Pr_s^{\sigma_\varepsilon,\pi^*}(\Psi_{1s}) + \varepsilon$ and $\sup_{\pi \in \Pi} \Pr_s^{\sigma_\varepsilon,\pi}(\Psi_{2s}) \leq \Pr_s^{\sigma_\varepsilon,\pi^*}(\Psi_{2s}) + \varepsilon$. Hence (σ^*, π^*) is an ε-Nash equilibrium at s, where $\sigma^* = \sigma_\varepsilon$. The proof for the case when we have a state s such that $\langle\langle 2\rangle\rangle_{val}(\Psi_2)(s) = 1$ is symmetric.

3. Suppose there is a state $s \in C$ such that $\langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) = 0$ and $\langle\langle 2\rangle\rangle_{val}(\Psi_2)(s) = 0$. Then consider ε-spoiling strategy pair $(\overline{\sigma_\varepsilon}, \overline{\pi_\varepsilon}) \in \overline{\Sigma_\varepsilon} \times \overline{\Pi_\varepsilon}$. Since $\overline{\sigma_\varepsilon}$ and $\overline{\pi_\varepsilon}$ are ε-spoiling strategies it follows that $\sup_{\pi \in \Pi} \Pr_s^{\overline{\sigma_\varepsilon},\pi}(\Psi_{2s}) \leq \varepsilon$ and $\sup_{\sigma \in \Sigma} \Pr_s^{\sigma,\overline{\pi_\varepsilon}}(\Psi_{1s}) \leq \varepsilon$. Hence $(\sigma^*, \pi^*) = (\sigma_\varepsilon, \pi_\varepsilon)$ is an ε-Nash equilibrium at s. ∎

Let $W_1 = \{ t_{00}, t_{01} \}$ and $W_2 = \{ t_{00}, t_{10} \}$. We consider a nonzero-sum reachability game $\mathcal{G}_R$ on the Sscc game structure $\mathcal{G}$, such that the objectives for player 1 and player 2 are $\mathrm{Reach}(W_1)$ and $\mathrm{Reach}(W_2)$, respectively. The following key lemma states that if properties (P1-P4) do not hold for every state s in C, then there exists an ε-Nash equilibrium in the game $\mathcal{G}_R$ such that the values of the ε-Nash equilibrium is greater than the respective zero-sum values of the original game. The proof idea is as follows: with the assumption that properties (P1-P4) do not hold for every state in C we establish the existence of ε-optimal strategies σ_ε and π_ε such that $\Pr_s^{\sigma_\varepsilon,\pi_\varepsilon}(\mathrm{Reach}(U)) = 1$, for all states $s \in C$, as $\varepsilon \to 0$. The above fact and a fragment of analysis of Vieille [19] enables us to establish the following lemma. The proof of the lemma is non-trivial, and requires involved construction of *punishing* ε-optimal strategies (details available in [1]).

Lemma 2. *Let $\mathcal{G}$ be a* Sscc *game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. If properties (P1-P4) do not hold for every state $s \in C$, then for every $\varepsilon > 0$, there is an ε-Nash equilibrium (σ^*, π^*) in the nonzero-sum reachability game $\mathcal{G}_R$, and there exists $k \in \mathbb{N}$ such that*

1. *$\Pr_s^{\sigma^*,\pi^*}(Reach^k(U)) \geq 1 - \varepsilon$; where $Reach^k(U)$ denotes reachability to U in k steps, i.e., $Reach^k(U) = \{ \langle s_0, s_1, s_2, \ldots\rangle \in \Omega \mid \exists i.\ 0 \leq i \leq k.\ s_i \in U \}$;*
2. *for all plays $\omega = \langle s_0, s_1, s_2, \ldots\rangle \in \mathrm{Outcome}(s, \sigma^*, \pi^*)$, if $\omega_k = \langle s_0, s_1, \ldots, s_k\rangle$ and $s_0 = s$, then (a) $\Pr_s^{\sigma^*,\pi^*}(Reach(W_1) \mid \omega_k) \geq \langle\langle 1\rangle\rangle_{val}(\Psi_1)(s_k) - \varepsilon$; and (b) $\Pr_s^{\sigma^*,\pi^*}(Reach(W_2) \mid \omega_k) \geq \langle\langle 2\rangle\rangle_{val}(\Psi_2)(s_k) - \varepsilon$.*

Lemma 3. *Let $\mathcal{G}$ be a* Sscc *game structure with parity objective Ψ_1 for player 1 and Ψ_2 for player 2. If for every state $s \in C$ the properties (P1-P4) do not hold, then for every $\varepsilon > 0$, there is an ε-Nash equilibrium for every state $s \in C$.*

Proof. Fix arbitrary $\varepsilon > 0$, and we show that there is an 3ε-Nash equilibrium for every state $s \in C$. Since ε is arbitrary the result follows. Let (σ^*, π^*) be an ε-Nash equilibrium of the reachability game $\mathcal{G}_R$ as specified in Lemma 2. Consider the strategy σ^*_ε for player 1 defined as follows:

$$\sigma^*_\varepsilon(s_0, s_1 \ldots, s_l) = \begin{cases} \sigma^*(s_0, s_1, \ldots, s_l) & \text{if } l < k \\ \overline{\sigma}_\varepsilon(s_0, s_1, \ldots, s_l) & \text{if } l \geq k \end{cases}$$

where k of Lemma 2 is used and $\overline{\sigma}_\varepsilon \in \overline{\Sigma_\varepsilon}$, i.e., player 1 plays σ^* for k steps and then switches to an ε-spoiling strategy $\overline{\sigma}_\varepsilon$. Similarly, we define the strategy π^*_ε for player 2. Since $\Pr_s^{\sigma^*,\pi^*}(\text{Reach}^k(U)) \geq 1 - \varepsilon$, we have that

$$\Pr_s^{\sigma^*_\varepsilon,\pi^*_\varepsilon}(\Psi_{1s}) \geq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_1)) - \varepsilon; \qquad \Pr_s^{\sigma^*_\varepsilon,\pi^*_\varepsilon}(\Psi_{2s}) \geq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_2)) - \varepsilon.$$

Recall that (σ^*, π^*) is an ε-Nash equilibrium of the reachability game such that for all plays $\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \text{Outcome}(s, \sigma^*, \pi^*)$, if $\omega_k = \langle s_0, s_1, \ldots, s_k \rangle$ and $s_0 = s$, then $\Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_1) \mid \omega_k) \geq \langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s_k) - \varepsilon$ and $\Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_2) \mid \omega_k) \geq \langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s_k) - \varepsilon$. Since the players play an ε-spoiling strategy after k-steps it follows that

$$\sup_{\sigma \in \Sigma} \Pr_s^{\sigma,\pi^*_\varepsilon}(\Psi_{1s}) \leq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_1)) + 2\varepsilon \leq \Pr_s^{\sigma^*_\varepsilon,\pi^*_\varepsilon}(\Psi_{1s}) + 3\varepsilon;$$

$$\sup_{\pi \in \Pi} \Pr_s^{\sigma^*_\varepsilon,\pi}(\Psi_{2s}) \leq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(W_2)) + 2\varepsilon \leq \Pr_s^{\sigma^*_\varepsilon,\pi^*_\varepsilon}(\Psi_{2s}) + 3\varepsilon.$$

Hence it follows that $(\sigma^*_\varepsilon, \pi^*_\varepsilon)$ is an 3ε-Nash equilibrium. ∎

Theorem 1 (ε-Nash equilibrium in Sscc **game).** *Let $\mathcal{G}$ be a* Sscc *game structure with parity objective Ψ_1 for player 1 and Ψ_2 for player 2. For every $\varepsilon > 0$, there is an ε-Nash equilibrium for every state $s \in C$.*

Proof. If for every state $s \in C$, the properties (P1-P4) do not hold, then the result follows from Lemma 3. Otherwise, there is a state $s \in C$, such that one of the properties (P1-P4) hold at s, and then by Lemma 1, for every $\varepsilon > 0$, there is an ε-Nash equilibrium at state s. By Proposition 1 we can replace s by the gadget gad(s). This breaks C into smaller strongly connected components. We can then proceed recursively on the smaller strongly connected components in a bottom-up order. The idea is as follows: consider the transformed game $\mathcal{G}'$ with s replaced by gad(s). Observe that in the graph $G_{\mathcal{G}'}$ edges out of s end in T, and after replacing s by gad(s) it belongs to the set U. Consider a lowest strongly connected component $C_1 \subset C$ in the gamegraph $\mathcal{G}'$, i.e., in the graph $G_{\mathcal{G}'}$ there is no edge from C_1 to a state in $C \setminus (C_1 \cup \{s\})$. We consider two cases.

1. If for every state $s_1 \in C_1$, the properties (P1-P4) do not hold, and then from Lemma 3 we conclude that for all $\varepsilon > 0$, ε-Nash equilibrium exists for all states $s_1 \in C_1$.
2. Else for some state $s_1 \in C_1$, one of the properties (P1-P4) hold at s_1, and then using the result of Lemma 1, s_1 can be replaced by $\mathrm{gad}(s_1)$ and we proceed recursively.

Hence we conclude (by induction on the size of the components) that for all $\varepsilon > 0$, ε-Nash equilibrium exists for every state $s_1 \in C_1$. By Proposition 1 every state $s_1 \in C_1$ can be replaced by the gadget $\mathrm{gad}(s_1)$. This gives a smaller Sscc game and we proceed in bottom-up fashion to establish the desired result. ∎

4 Existence of ε-Nash equilibrium

In this section we show that for all nonzero-sum concurrent game structures $\mathcal{G}$, with ω-regular objectives specified as parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively, for every $\varepsilon > 0$, there exists an ε-Nash equilibrium for every state s of game $\mathcal{G}$. The proof follows from an inductive argument: by induction on the size of the state space of $\mathcal{G}$ and by application of Theorem 1. We assume without loss of generality that there are four special states $\{ t_{00}, t_{01}, t_{10}, t_{11} \}$ in $\mathcal{G}$, as defined in Definition 7.

Lemma 4. *Let $\mathcal{G}$ be a concurrent game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. Let $G_{\mathcal{G}}$ be the graph of $\mathcal{G}$ and* TC *be a terminal strongly connected component in $G_{\mathcal{G}}$. Then for every $\varepsilon > 0$, there is an ε-Nash equilibrium for every state $s \in \mathrm{TC}$.*

Proof. The proof is by induction on the size of TC. It is easy to argue when $|\mathrm{TC}| = 1$, i.e., TC consists of an absorbing state. Consider the sub-game induced by the set of states TC and call the sub-game $\mathcal{G}_{\mathrm{TC}}$.

- Suppose there is a state $s \in \mathrm{TC}$ such that $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 1$. Then fix an ε-optimal strategy σ for player 1 and let π be an ε-optimal strategy for player 2 against σ. Then (σ, π) is an ε-Nash equilibrium. We can replace s by the gadget described in Proposition 1. This will break TC into (possibly many) smaller strongly connected components. By induction hypothesis, Theorem 1 and the bottom-up evaluation procedure described in Theorem 1 it follows that ε-Nash equilibrium exists at every state in TC. Similar arguments hold if there is a state $s \in \mathrm{TC}$ such that $\langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s) = 1$.
- Suppose for every state $s \in \mathrm{TC}$ we have $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) < 1$ and $\langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s) < 1$. It follows from Corollary 1 of [6] that in a zero-sum concurrent game with ω-regular objectives if for every state s we have $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) < 1$, then for every state s in the game we have $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 0$, i.e., if the zero-sum value is positive for player 1 at some state, then there exists a state s where the zero-sum value is 1. Hence it follows from the above condition that for all states $s \in \mathrm{TC}$ we have $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 0$ and $\langle\langle 2 \rangle\rangle_{val}(\Psi_2)(s) = 0$. Let $\overline{\pi}_\varepsilon$ be an ε-spoiling strategy for

player 2 and $\overline{\sigma}_\varepsilon$ be an ε-spoiling strategy for player 1. Hence we have the following inequalities: $\sup_{\sigma\in\Sigma} \Pr_s^{\sigma,\overline{\pi}_\varepsilon}(\Psi_{1s}) \leq \varepsilon$ and $\sup_{\pi\in\Pi} \Pr_s^{\overline{\sigma}_\varepsilon,\pi}(\Psi_{2s}) \leq \varepsilon$. Hence we have $(\overline{\sigma}_\varepsilon, \overline{\pi}_\varepsilon)$ is an ε-Nash equilibrium for all states $s \in$ TC. ∎

Theorem 2 (ε-Nash equilibrium). *Let $\mathcal{G}$ be a concurrent game structure with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively. For every $\varepsilon > 0$, there is an ε-Nash equilibrium for every state $s \in S$.*

Proof. Let $G_{\mathcal{G}}$ be the graph of $\mathcal{G}$. It follows from Lemma 4 that for every state s in a terminal strongly connected component of $G_{\mathcal{G}}$ there is an ε-Nash equilibrium, for all $\varepsilon > 0$. By Proposition 1 we can replace every state s of a terminal strongly connected component by the gadget gad(s). For the rest of the strongly connected components we proceed in a bottom-up order as follows: consider a strongly connected component C when all the strongly connected component below it are replaced by the gadgets of Proposition 1. The sub-game induced by C and the gadgets of the strongly connected components below C form a Sscc game. By Theorem 1 we have there is an ε-Nash equilibrium for every state $s \in C$, for all $\varepsilon > 0$. ∎

5 Computational Complexity

In this section we show how to compute the values of an ε-Nash equilibrium of Sscc games within ε-precision. We prove that every case of the existence proof of ε-Nash equilibrium is constructive and computable. It may be noted that even in the case of zero-sum concurrent games with parity objectives the values can be irrational (for an example see [7]). Hence, one can only achieve ε-approximation of the values in the general case of nonzero-sum concurrent parity games. It follows from the inductive argument of Theorem 2 that the values of an ε-Nash equilibrium for nonzero-sum concurrent games with parity objectives can be computed by $|S|$-iterations of a procedure to compute ε-Nash equilibrium values for Sscc games.

Complexity of ε-Nash equilibrium. To analyze the complexity of computing values of an ε-Nash equilibrium in Sscc games we consider the following cases:

1. Case 1. Compute the values of ε-Nash equilibrium when the property P1 is satisfied for some state s.
2. Case 2. Compute the values of ε-Nash equilibrium when the property P4 is satisfied for some state s.
3. Case 3. Compute the values of ε-Nash equilibrium when the property P2 or P3 is satisfied for some state s.
4. Case 4. Compute the values of some special ε-Nash equilibrium of Sscc games with reachability objectives.

We analyze the above cases below.

1. Case 1. Given Ψ_1 and Ψ_2 are parity objectives, the objective $\Psi_1 \cap \Psi_2$ is a Streett objective [17]. To analyze the computation of

$\sup_{(\sigma,\pi)\in\Sigma\times\Pi} \Pr_s^{\sigma,\pi}(\Psi_1 \cap \Psi_2)$, observe that this is equivalent to the computation of values of MDPs where player 1 and player 2 cooperate to achieve the objective $\Psi_1 \cap \Psi_2$. Hence the computation reduces to computing values in a MDP with Streett objective, which can be done in polynomial time [3].
2. Case 2. After the computation of the zero-sum values $\langle\langle 1\rangle\rangle_{val}(\Psi_1)(\cdot)$ and $\langle\langle 2\rangle\rangle_{val}(\Psi_2)(\cdot)$, it is easy to determine if there is a state s such that $\langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) = 0$ and $\langle\langle 2\rangle\rangle_{val}(\Psi_2)(s) = 0$. Hence Case 2 can be solved by computing the zero-sum values for player 1 and player 2.
3. Case 3. Given the zero-sum values for player 1 and player 2, we describe a polynomial time procedure to determine the values of an ε-Nash equilibrium when property P2 or P3 is satisfied. We prove the result for the case when property P2 is satisfied and the result for the case when property P3 is satisfied is symmetric. Consider the set $W = \{ s \mid \langle\langle 1\rangle\rangle_{val}(\Psi_1)(s) = 1 \}$ of states that have zero-sum value 1 for player 1. Since property P2 is satisfied, we have $W \cap C \neq \emptyset$. Given a state $s \in W$, consider the set $\mathrm{SafeAct}(s) = \{a \in \Gamma_1(s) \mid \forall\, b \in \Gamma_2(s).\ \mathrm{Dest}(s,a,b) \subseteq W\}$ of moves for player 1 that ensures that the set W is never left. Consider a reduced sub-game $\mathcal{G}'$ induced by W such that at every state $s \in W$ the available moves for player 1 is $\mathrm{SafeAct}(s)$. Let Σ' be the set of strategies such that player 1 plays only moves in $\mathrm{SafeAct}(s)$ for every state $s \in W$, i.e., the set of strategies in $\mathcal{G}'$. We compute the values $\langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) = \sup_{(\sigma,\pi)\in\Sigma'\times\Pi} \Pr_s^{\sigma,\pi}(\Psi_1 \cap \Psi_2)$. It may be noted that there exists ε-optimal strategy σ_ε in the original game such that for every strategy $\pi \in \Pi$ we have $\Pr_s^{\sigma_\varepsilon,\pi}(\Psi_1 \cap \mathrm{Safe}(W)) \geq 1-\varepsilon$, for all states $s \in W$. Hence it follows that

$$\begin{aligned}\Pr_s^{\sigma_\varepsilon,\pi}(\Psi_2) &\leq \Pr_s^{\sigma_\varepsilon,\pi}(\Psi_1 \cap \Psi_2 \cap \mathrm{Safe}(W)) + \varepsilon \\ &\leq \sup_{(\sigma,\pi)\in\Sigma'\times\Pi} \Pr_s^{\sigma,\pi}(\Psi_1 \cap \Psi_2) + \varepsilon. \end{aligned} \tag{1}$$

   - If for some state $s \in W \cap C$ we have $\langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) = 1$, then property P1 is satisfied and then Case 1 is followed.
   - Else for every state $s \in W \cap C$ we have $\langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) < 1$. It follows from property of MDPs that for any ω-regular objective Ψ, the maximum probability to satisfy Ψ is equal to the maximum probability of reaching the set of states where the value is 1. Hence we have

$$\sup_{(\sigma,\pi)\in\Sigma'\times\Pi} \Pr_s^{\sigma,\pi}(\Psi_1 \cap \Psi_2) = \sup_{(\sigma,\pi)\in\Sigma'\times\Pi} \Pr_s^{\sigma,\pi}(\mathrm{Reach}(t_{00})) \tag{2}$$

   We show that for every state $s \in W \cap C$, the profile $(1, \langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s))$ is the values of an ε-Nash equilibrium profile, for all $\varepsilon > 0$. Let $(\widehat{\sigma}, \widehat{\pi})$ be a memoryless strategy profile such that $\Pr_s^{\widehat{\sigma},\widehat{\pi}}(\mathrm{Reach}(t_{00})) = \langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s)$, for all $s \in W \cap C$. The existence of such a memoryless strategy profile follows from [4]. For any $\varepsilon > 0$, let $k \in \mathbb{N}$ be such that $\Pr_s^{\widehat{\sigma},\widehat{\pi}}(\mathrm{Reach}^k(t_{00})) \geq \langle\langle 1,2\rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s) - \varepsilon$. The strategy profile (σ^*, π^*) is described as follows:

$$\sigma^*(s_0, s_1, \ldots, s_l) = \begin{cases} \widehat{\sigma}(s_0, s_1, \ldots, s_k) & l < k \\ \sigma_\varepsilon(s_0, s_1, \ldots, s_k) & l \geq k \end{cases}$$

where $\sigma_\varepsilon \in \Sigma_\varepsilon$ and $\pi^* = \widehat{\pi}$. Given strategy σ^*, for any strategy π the play never leaves W within k steps, since $\widehat{\sigma} \in \Sigma'$. Since $\sigma_\varepsilon \in \Sigma_\varepsilon$ and for every state $s \in W$ we have $\langle\langle 1 \rangle\rangle_{val}(\Psi_1)(s) = 1$ it follows that $\Pr_s^{\sigma^*,\pi^*}(\Psi_1) \geq 1 - \varepsilon$. Since σ^* follows $\widehat{\sigma}$ for k steps, it follows that $\Pr_s^{\sigma^*,\pi^*}(\Psi_2) \geq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(t_{00})) - \varepsilon$. It follows from equation 1 and 2 that $\sup_{\pi \in \Pi} \Pr_s^{\sigma^*,\pi}(\Psi_2) \leq \Pr_s^{\sigma^*,\pi^*}(\text{Reach}(t_{00})) + \varepsilon$. Hence it follows that $(1, \langle\langle 1,2 \rangle\rangle_{val}(\Psi_1 \cap \Psi_2)(s))$ is an ε-Nash equilibrium value profile for all states $s \in W \cap C$, for all $\varepsilon > 0$.

It follows from above that the values of an ε-Nash equilibrium of states $s \in C$ can be computed by a polynomial procedure and solving the zero-sum values for player 1 and player 2 when Case 1, Case 2 or Case 3 is satisfied. The analysis of Case 4 involves solving some special ε-Nash equilibrium values of the game $\mathcal{G}_R$ with reachability objectives. The existence of polynomial witness and polynomial time verification procedure for Case 4 follows from the results similar to [5] (details in [1]).

Let $\mathbf{ZS}(\mathcal{G}, \Psi, \varepsilon)$ denote the time complexity of an algorithm to compute the zero-sum values of a concurrent game structure $\mathcal{G}$ within ε-precision, for a parity objective Ψ. Let $\mathbf{NZReach}(\mathcal{G}, \varepsilon, \Psi_1^R, \Psi_2^R)$ denote the time complexity of an algorithm to compute the values of an ε-Nash equilibrium, greater than some specified value, of a concurrent game structure $\mathcal{G}$ with reachability objectives Ψ_1^R and Ψ_2^R for player 1 and player 2, respectively. It follows from [2] and [5] that there exist $\mathbf{ZS}(\mathcal{G}, \varepsilon, \Psi)$ and $\mathbf{NZReach}(\mathcal{G}, \varepsilon, \Psi_1^R, \Psi_2^R)$ that are in the complexity class TFNP, for all constants $\varepsilon > 0$. The above analysis yields the next Theorem.

Theorem 3 (Complexity of ε-Nash equilibrium).

1. *The values of an ε-Nash equilibrium of a nonzero-sum concurrent game structure $\mathcal{G}$ with parity objectives Ψ_1 and Ψ_2 for player 1 and player 2, respectively, can be computed in time*

$$O\big(n \cdot (\mathbf{ZS}(\mathcal{G}, \varepsilon, \Psi_1) + \mathbf{ZS}(\mathcal{G}, \varepsilon, \Psi_2) + \mathbf{NZReach}(\mathcal{G}, \varepsilon, \Psi_1^R, \Psi_2^R))\big) + O\big(p(|\mathcal{G}|)\big)$$

   where p is a polynomial function and $n = |S|$ is the size of the state space.
2. *For all constants $\varepsilon > 0$, the values of an ε-Nash equilibrium of nonzero-sum concurrent games with parity objectives can be computed in TFNP; and hence in EXPTIME.*

The existence of ε-Nash equilibrium, for all $\varepsilon > 0$, for higher levels of Borel hierarchy than ω-regular objectives, and for ω-regular objectives for more than two-players are interesting open problems.

Acknowledgments. I am grateful to Tom Henzinger and Luca de Alfaro for several insights on concurrent ω-regular games. I thank Rupak Majumdar for interesting discussions. I am deeply indebted to Nicollas Vieille for several key insights of his results that he gave me. This research was supported in part by the ONR grant N00014-02-1-0671, the AFOSR MURI grant F49620-00-1-0327, and the NSF grant CCR-0225610.

References

1. K. Chatterjee. Two-player nonzero-sum ω-regular games. 2004. Technical Report: UCB/CSD-04-1364.
2. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. The complexity of quantitative concurrent parity games. 2004. Technical Report: UCB/CSD-04-1354.
3. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. The complexity of stochastic Rabin and Streett games. 2004. Technical Report: UCB/CSD-04-1355.
4. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. Trading memory for randomness. In *QEST 04*. IEEE Computer Society Press, 2004.
5. K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In *CSL 04*, pages 26–40. LNCS, 2004.
6. L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *LICS 00*, pages 141–154. IEEE Computer Society Press, 2000.
7. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *STOC 01*, pages 675–683. ACM Press, 2001.
8. J. Filar and K. Vrieze. *Competitive Markov Decision Processes.* Springer-Verlag, 1997.
9. A.M. Fink. Equilibrium in a stochastic n-person game. *Journal of Science of Hiroshima University*, 28:89–93, 1964.
10. J.F. Nash Jr. Equilibrium points in n-person games. *Proceedings of the National Academny of Sciences USA*, 36:48–49, 1950.
11. D.A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
12. G. Owen. *Game Theory.* Academic Press, 1995.
13. C.H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *JCSS*, 48(3):498–532, 1994.
14. C.H. Papadimitriou. Algorithms, games, and the internet. In *STOC 01*, pages 749–753. ACM Press, 2001.
15. P. Secchi and W.D. Sudderth. Stay-in-a-set games. *International Journal of Game Theory*, 30:479–490, 2001.
16. L.S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. USA*, 39:1095–1100, 1953.
17. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
18. M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *STOC 85*, pages 327–338. IEEE Computer Society Press, 1985.
19. N. Vieille. Two player stochastic games I: a reduction. *Israel Journal of Mathematics*, 119:55–91, 2000.
20. N. Vieille. Two player stochastic games II: the case of recursive games. *Israel Journal of Mathematics*, 119:93–126, 2000.
21. J. von Neumann and O. Morgenstern. *Theory of games and economic behavior.* Princeton University Press, 1947.
22. B. von Stengel. Computing equilibria for two-person games. *Chapter 45, Handbook of Game Theory*, 3:1723–1759, 2002.