The Complexity of Quantitative Concurrent Parity Games

Krishnendu Chatterjee *Luca deAlfaro* *Thomas A. Henzinger*

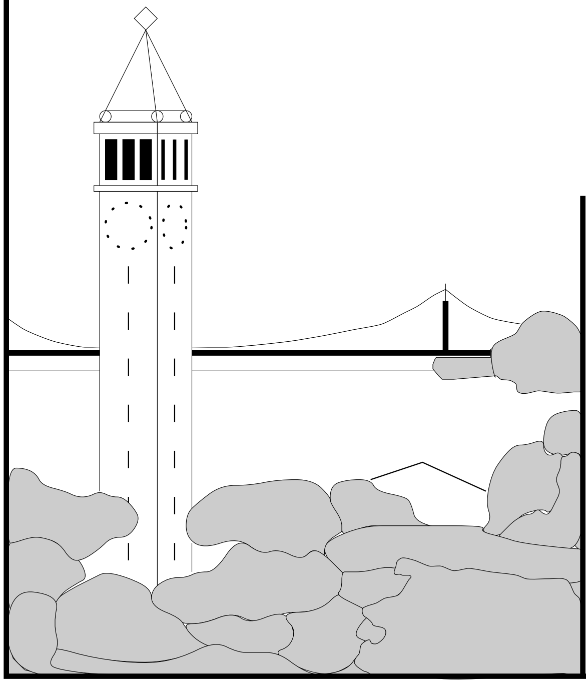

Report No. UCB/CSD-4-1354

November 2004

Computer Science Division (EECS)
University of California
Berkeley, California 94720

The Complexity of Quantitative Concurrent Parity Games*

Krishnendu Chatterjee[†] Luca de Alfaro[§] Thomas A. Henzinger[†‡]

[†] EECS, University of California, Berkeley,USA

[§] CE, University of California, Santa Cruz,USA

[‡] Computer and Communication Sciences, EPFL, Switzerland

{c_krish,tah}@eecs.berkeley.edu, luca@soe.ucsc.edu

Abstract

We consider two-player infinite games played on graphs. The games are concurrent, in that at each state the players choose their moves simultaneously and independently, and stochastic, in that the moves determine a probability distribution for the successor state. The value of a game is the maximal probability with which a player can guarantee the satisfaction of her objective. We show that the values of concurrent games with ω-regular objectives expressed as parity conditions can be computed in NP $\cap$ coNP. This result substantially improves the best known previous bound of 3EXPTIME. It also shows that the full class of concurrent parity games is no harder than the special cases of turn-based deterministic parity games (Emerson-Jutla) and of turn-based stochastic reachability games (Condon), for both of which NP $\cap$ coNP is the best known bound.

While the previous, more restricted NP $\cap$ coNP results for graph games relied on the existence of particularly simple (pure memoryless) optimal strategies, in concurrent games with parity objectives optimal strategies may not exist, and ε-optimal strategies (which achieve the value of the game within a parameter $\varepsilon > 0$) require in general both randomization and infinite memory. Hence our proof must rely on a more detailed analysis of strategies and, in addition to the main result, yields two results that are interesting on their own. First, we show

*The work was supported by the AFOSR MURI grant F49620-00-1-0327, by the ONR grant N00014-02-1-0671, and by the NSF grants CCR-0132780, CCR-0234690, CCR-9988172, and CCR-0225610, and by the NSF Career grant CCR-0132780, the NSF grant CCR-0234690, and by the ONR grant N00014-02-1-0671

that there exist ε-optimal strategies that in the limit coincide with memoryless strategies; this parallels the celebrated result of Mertens-Neyman for concurrent games with limit-average objectives. Second, we complete the characterization of the memory requirements for ε-optimal strategies for concurrent ω-regular games, by showing that memoryless strategies suffice for ε-optimality for coBüchi conditions.

1 Introduction

We consider *recursive games* played between two players over a graph [22, 10, 16]. The games proceed in an infinite number of rounds. At each round, the players choose moves; the two moves, together with the current state, determine a probability distribution for the successor state. An outcome of the game, or *play*, consists in the infinite sequence of states visited. These graph games can be broadly classified into *turn-based* and *concurrent* games. In turn-based games, in any given round only one player can choose among multiple moves: effectively, the set of states of the graph can be partitioned into the states where it is player 1's turn to play, and the states where it is player 2's turn to play. In concurrent games, both players may have multiple moves available at each state, and the players choose their moves simultaneously and independently.

An important class of winning conditions are the *ω-regular languages*. In such games, the goal of player 1 is to ensure that the play belongs to a specified ω-regular language; the goal of player 2 is to ensure that the play does not belong to the language. The games are thus *zero-sum*: the objectives of the two players are complementary. The ω-regular languages are the generalization to infinite words of the classical regular languages [24]; the properties expressible by ω-regular languages include safety, reachability, and fairness. Games with ω-regular winning conditions have been applied to system synthesis [2, 21, 19] and verification [9, 13, 7]. Of particular interest are ω regular languages that are given as *parity conditions* on game graphs; this is because every ω-regular game can be converted into a parity game [18, 25, 26].

Given a recursive game and an ω-regular language $\mathcal{L}$, the *value* $\langle\langle 1\rangle\rangle_{val}(\mathcal{L})(s)$ of the game for player 1 at a state s is equal to the maximal probability with which player 1 can ensure that the play lies in $\mathcal{L}$; the value $\langle\langle 2\rangle\rangle_{val}(\overline{\mathcal{L}})(s)$ of the game for player 2 at s is equal to the maximal probability with which player 2 can ensure that the play lies outside $\mathcal{L}$. Martin's determinacy theorem ensures that $\langle\langle 1\rangle\rangle_{val}(\mathcal{L})(s)+\langle\langle 2\rangle\rangle_{val}(\overline{\mathcal{L}})(s) = 1$ [15]. Except for the special case of turn-based games, little has been known about the *computational complexity* of finding the value for a recursive game with an ω-regular winning condition. In the turn-based case, it is known that the value of games with ω-regular conditions can be computed in NP $\cap$ coNP. This result was first obtained for turn-based *deterministic* parity games, in which each moves determines uniquely (instead of probabilistically) the successor state [9], and for turn-based stochastic *reachability* games [5]; the case of turn-based stochastic parity games was shown in [3].

Concurrent games are substantially more complex than turn-based games in several respects. To see this, consider the structure of *optimal strategies*, which are strategies that achieve the value of a given game. For turn-based stochastic ω-regular games, there always exist *pure* (deterministic) optimal strategies, which do not rely on randomized choice [3]; in the case of turn-based stochastic parity games, moreover, there are always pure *memoryless* optimal strategies, where the choice of move depends only on the current state, rather than also on the past history of the game. It is this observation that led to the NP $\cap$ coNP result for turn-based parity games.

By contrast, in concurrent games, already for reachability conditions, players must in general play with *randomized* (non-pure) strategies, which prescribe, at each round, a probability distribution over the moves to be played. Furthermore, optimal strategies may not exist: rather, for every real $\varepsilon > 0$, the players have ε-*optimal* strategies, which achieve the value of the game within ε. Even for relatively simple winning conditions, such as Büchi conditions, ε-optimal strategies need both randomization and infinite memory [8]. It is therefore not inconceivable that the complexity of concurrent ω-regular games might be considerably worse than NP $\cap$ coNP. The only known previous algorithm for computing the value of concurrent parity games is triple-exponential [8]: it was obtained via a reduction to the theory of the real closed field, by using decision procedures for the theory of reals with addition and multiplication [23, 1].

In this paper, we show that the problem of computing the value of a concurrent parity game is in NP $\cap$ coNP. More precisely, as the value of a concurrent game at a state can be an irrational number, we show that given an encoding of the game and of a rational $\varepsilon > 0$, the problem of approximating the value of the game within ε can be solved in NP $\cap$ coNP. This result generalizes the best known upper bound (NP $\cap$ coNP) for very restricted cases, such as turn-based deterministic parity games and turn-based stochastic reachability games, to the class of all concurrent parity games.

The basic idea behind the proof, which can no longer rely on the existence of pure memoryless optimal strategies, is as follows. We call a *value class* a maximal set of states where the game has the same value for player 1. By the results of [6] on *qualitative* winning (i.e., winning with probability 1), if the (player 1) value of the game is not constant 1 or 0, then there are two non-empty value classes W_1 and W_2 where the value is 1 and 0, respectively. We show that if the players play ε-optimal strategies, then $W_1 \cup W_2$ is reached with probability 1. Through a detailed analysis of the branching structure of the stochastic process of the game, we go on to show that we can construct a ε-optimal strategy by stitching together strategies, one per each value class. This gives us a polynomial witness for the resulting strategy and proves membership in NP; membership in NP $\cap$ coNP follows from the fact that the problem is symmetrical in players 1 and 2.

A detailed analysis of our proof gives us several new results about the structure of ε-optimal strategies in concurrent parity games. First, we show that concurrent games with coBüchi winning conditions admit memoryless ε-optimal strategies. This result completes the characterization of the memory requirements of the optimal strategies for concurrent ω-regular games: it was previously known that safety and reachability games admit memoryless ε-optimal strategies [11, 8], and that Büchi conditions may require infinite memory [8]. Second, we show that in concurrent parity games, the limit of the ε-optimal strategies for $\varepsilon \to 0$ is a memoryless strategy (which in general is not optimal). This result parallels the celebrated result of Mertens-Neyman [17] for concurrent games with limit-average objectives.

2 Definitions

Notation. For a countable set A, a *probability distribution* on A is a function $\delta\colon A \mapsto [0,1]$ 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 $G = \langle S, Moves, \Gamma_1, \Gamma_2, \delta\rangle$ consists of the following components:*

- *A finite state space S.*
- *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 move* a_2, *for all* $s, t \in S$ *and* $a_1 \in \Gamma_1(s)$, $a_2 \in \Gamma_2(s)$. ■

We distinguish the following special classes of concurrent game structures.

- A concurrent game structure G is *deterministic* if for all $s \in S$ and all $a_1 \in \Gamma_1(s)$, $a_2 \in \Gamma_2(s)$, there is a $t \in S$ such that $\delta(s, a_1, a_2)(t) = 1$.

- A concurrent game structure G is *turn-based* if at every state at most one player can choose among multiple moves; that is, if for every state $s \in S$ there exists at most one $i \in \{1, 2\}$ with $|\Gamma_i(s)| > 1$.

We define the size of the game structure G to be equal to the size of the transition function δ; specifically, $|G| = \sum_{s \in S} \sum_{a \in \Gamma_1(s)} \sum_{b \in \Gamma_2(s)} \sum_{t \in S} |\delta(s, a, b)(t)|$, where $|\delta(s, a, b)(t)|$ denotes the space to specify the probability distribution. We write n to denote the size of the state space, i.e., $n = |S|$. 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 choice of moves of the players the next state is always s. A state s is a *turn-based state* if there exists $i \in \{ 1, 2 \}$ such that $|\Gamma_i(s)| = 1$. Moreover, if $|\Gamma_2(s)| = 1$ then the state s is a player-1 turn-based state since the choice of moves for player 2 is trivial; and if $|\Gamma_1(s)| = 1$ then it is a player-2 turn-based state. We assume that the players act *non-cooperatively*, i.e., each player chooses her strategy independently and secretly from the other player, and is only interested in maximizing her own reward. 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 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.

2.1 Randomized strategies

A *selector* ξ for player $i \in \{1, 2\}$ is a function $\xi : S \mapsto \mathcal{D}(Moves)$ such that for all $s \in S$ and $a \in Moves$, 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 selector ξ is *pure* if for every $s \in S$ there is $a \in Moves$ such that $\xi(s)(a) = 1$; we denote by $\Lambda_i^P \subseteq \Lambda_i$ the set of pure selectors for player i. A *strategy* for player 1 is a function $\sigma : S^+ \to \Lambda_1$ 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 strategy σ for player i is *pure* if it yields only pure selectors, that is, is of type $S^+ \to \Lambda_i^P$. A strategy with memory can be described as a pair of functions: (a) memory update function $\sigma_u : S \times \mathtt{M} \to \mathtt{M}$, and (b) next move function $\sigma_m : S \times \mathtt{M} \to \Lambda_1$. A strategy with memory is finite memory if $\mathtt{M}$ is finite. 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. A strategy is pure memoryless if it is pure and memoryless. We denote by $\Sigma^P, \Sigma^F, \Sigma^{PM}$ the family of pure, finite-memory and pure memoryless strategies for player 1 respectively. Analogously we define the families of strategies for player 2. 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 follows the strategies σ and π. For $i \geq 0$, we also denote by $\Theta_i : \Omega_s \to S$ the random variable denoting the i-th state along a path.

2.2 Objectives

We specify objectives for the players by providing the set of *winning plays* $\Phi \subseteq \Omega$ for each player. In this paper we study only zero-sum games [20, 11], where the objectives of the two players are strictly competitive. In other words, it is implicit that if the objective of one player is Φ, then the objective of the other player is $\Omega \setminus \Phi$. Given a game graph G and an objective $\Phi \subseteq \Omega$, we write (G, Φ) for the game played on the graph G with the objective Φ for player 1.

A general class of objectives are the Borel objectives [12]. A *Borel objective* $\Phi \subseteq S^\omega$ is a Borel set in the Cantor topology on S^ω. In this paper we consider *ω-regular objectives* [26], which lie in the first $2\frac{1}{2}$ levels of the Borel hierarchy (i.e., in the intersection of Σ_3 and Π_3). The ω-regular objectives, and subclasses thereof, can be specified in the following forms. For a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega$, we define $\text{Inf}(\omega) = \{ s \in S \mid s_k = s \text{ for infinitely many } k \geq 0 \}$ to be the set of states that occur infinitely often in ω.

- *Reachability and safety objectives.* Given a set $T \subseteq S$ of "target" states, the reachability objective requires that some state of T be visited. The set of winning plays is thus $\text{Reach}(T) = \{ \omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0 \}$. Given a set $F \subseteq S$, the safety objective requires that only states of F be visited. Thus, the set of winning plays is $\text{Safe}(F) = \{ \omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in F \text{ for all } k \geq 0 \}$.

- *Büchi and coBüchi objectives.* Given a set $B \subseteq S$ of "Büchi" states, the Büchi objective requires that B is visited infinitely often. Formally, the set of winning plays is $\text{Büchi}(B) = \{ \omega \in \Omega \mid \text{Inf}(\omega) \cap B \neq \emptyset \}$. Given $C \subseteq S$, the coBüchi objective requires that all states visited infinitely often are in C. Formally, the set of winning plays is $\text{coBüchi}(C) = \{ \omega \in \Omega \mid \text{Inf}(\omega) \subseteq C \}$.

- *Parity objective.* For $c, d \in \mathbb{N}$, we let $[c..d] = \{ c, c+1, \ldots, d \}$. Let $p : S \mapsto [0..d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. The *Even parity objective* is defined as $\text{Parity}(p) = \{ \omega \in \Omega \mid \min\big(\text{Inf}(\omega)\big) \text{ is even } \}$, and the *Odd parity objective* as $\text{coParity}(p) = \{ \omega \in \Omega \mid \min\big(\text{Inf}(\omega)\big) \text{ is odd } \}$. Informally we say that a path ω satisfy the parity objective, $\text{Parity}(p)$, if $\omega \in \text{Parity}(p)$.

Note that for a priority function $p : V \to \{ 0, 1 \}$, an even parity objective $\text{Parity}(p)$ is equivalent to the Büchi objective $\text{Büchi}(p^{-1}(0))$, i.e., the Büchi set consists of the states with priority 0.

The ability to solve games with parity objectives suffices for solving games with arbitrary LTL (or ω-regular) objectives: in fact, it suffices to encode the ω-regular objective as a deterministic Rabin-chain automaton or parity automaton, solving then the game consisting of the synchronous product of the original game with the Rabin-chain automaton [18, 25].

Given any parity winning objective, we write Ω_e to denote $\text{Parity}(p)$; this set is measurable for any choice of strategies for the two players [27]. Similarly we write Ω_o to denote $\text{coParity}(p)$. Note that $\Omega_e \cap \Omega_o = \emptyset$ and

$\Omega_e \cup \Omega_o = \Omega$. Given a state s we write Ω_{es} to denote $\Omega_s \cap \Omega_e$ and similarly we write Ω_{os} to denote $\Omega_s \cap \Omega_o$. Hence, the probability that a path satisfies objective $\mathrm{Parity}(p)$ starting from state $s \in S$, given strategies σ, π for the players is $\mathrm{Pr}_s^{\sigma,\pi}(\Omega_{es})$. Given a state $s \in S$ and a parity winning objective, $\mathrm{Parity}(p)$, we are interested in finding the maximal probability with which player 1 can ensure that $\mathrm{Parity}(p)$ and player 2 can ensure that $\mathrm{coParity}(p)$ holds from s. We call such probability the *value of the game G at s* for player $i \in \{ 1, 2 \}$. The value for player 1 and player 2 are given by the function $\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e) : S \mapsto [0, 1]$ and $\langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o) : S \mapsto [0, 1]$, defined for all $s \in S$ by

$$\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Omega_{es})$$

$$\langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \mathrm{Pr}_s^{\sigma,\pi}(\Omega_{os}).$$

Note that the objectives of the player are complementary and hence we have a zero-sum game. Concurrent games satisfy a *quantitative* version of determinacy [15], stating that for all parity winning objectives, and all $s \in S$, we have

$$\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) + \langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o)(s) = 1.$$

A strategy σ for player 1 is *optimal* if for all $s \in S$ we have

$$\inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Omega_{es}) = \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s).$$

For $\varepsilon > 0$, a strategy σ for player 1 is *ε-optimal* if for all $s \in S$ we have

$$\inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Omega_{es}) \geq \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) - \varepsilon.$$

We define optimal and ε-optimal strategies for player 2 symmetrically. Note that the quantitative determinacy of concurrent games is equivalent to the existence of ε-optimal strategies for both players, for all $\varepsilon > 0$, at all states $s \in S$. We denote by $\langle\!\langle 1 \rangle\!\rangle_{limit} = \{ s \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) = 1 \}$ and $\langle\!\langle 2 \rangle\!\rangle_{limit} = \{ s \mid \langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o)(s) = 1 \}$, the set of states where player 1 and player 2 have values 1, respectively.

2.3 The branching structure of plays

Many of the arguments developed in this paper rely on a detailed analysis of the branching process resulting from the strategies chosen by the players, and from the probabilistic transition relation of the game. In order to make our arguments precise, we need some definitions. A play is *feasible* if each of its transitions could have arisen according to the transition relation of the game.

Definition 2 (Feasible plays and outcomes) *Given strategies σ for player 1 and π for player 2, a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ is* feasible *in a concurrent game graph G, if for every $k \in \mathbb{N}$ the following conditions hold: (1) $s_{k+1} \in \mathrm{Dest}(s_k, a_1, a_2)$; (2) $\sigma(s_0, s_1, \ldots, s_k)(a_1) > 0$ and (3) $\pi(s_0, s_1, \ldots, s_k)(a_2) > 0$. Given strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a state s, we denote by $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ the set of feasible plays that start from s, given strategies σ and π.* ∎

In order to make precise statements about the branching process arising from a game play, we define below trees labeled by game states.

Definition 3 (Infinite trees, S-labeled trees and trees for events) *An* infinite tree *is a set $\mathrm{Tr} \subseteq \mathbb{N}^*$ such that*

- *if $x \cdot i \in \mathrm{Tr}$ where $x \in \mathbb{N}^*$ and $i \in \mathbb{N}$ then $x \in \mathrm{Tr}$;*
- *for all $x \in \mathrm{Tr}$ there exists $i \in \mathbb{N}$ such that $x \cdot i \in \mathrm{Tr}$. We refer to $x \cdot i$ as a* successor *of x.*

We call the elements in Tr as nodes and the empty word ϵ is the root of the tree. An infinite path τ of Tr is a set $\tau \subseteq \mathrm{Tr}$ such that

- *$\epsilon \in \tau$;*
- *for every x in τ there is an unique $i \in \mathbb{N}$ such that $x \cdot i \in \tau$. Note that for every $i \in \mathbb{N}$, there is an unique element $x \in \tau$ such that $|x| = i$. We denote by τ_i the element $x \in \tau$ such that $|x| = i$.*

Given an infinite tree Tr and a node $x \in \mathrm{Tr}$, we denote by $\mathrm{Tr}(x)$ the sub-tree rooted at node x. Formally, $\mathrm{Tr}(x)$ denotes the set $\{ x' \in \mathrm{Tr} \mid x \text{ is a prefix of } x' \}$.

A S-labeled tree $\mathcal{T}$ is a pair $(\mathrm{Tr}, \langle\cdot\rangle)$, where Tr is a tree and $\langle\cdot\rangle : \mathrm{Tr} \to S$ maps each node of Tr to a state $s \in S$. Given a S-labelled tree $\mathcal{T}$, and a infinite path $\tau \subseteq \mathrm{Tr}$, we denote by $\langle\tau\rangle$ the play $\langle s_0, s_1, s_2, \ldots \rangle$, such that $s_0 = \langle\epsilon\rangle$ and for all $i > 0$ we have $s_i = \langle\tau_i\rangle$. A S-labeled tree $\mathcal{T}_s = (\mathrm{Tr}_s, \langle\cdot\rangle)$, where $\langle\epsilon\rangle = s$, represents a set of infinite paths, denoted as $\mathcal{L}(\mathcal{T}_s) \subseteq \Omega_s$, such that

$$\mathcal{L}(\mathcal{T}_s) = \{ \omega = \langle s_0 = s, s_1, s_2, \ldots \rangle \in \Omega_s \mid \exists \tau \subseteq \mathrm{Tr}_s.\ \langle\tau\rangle = \omega \}.$$

A S-labeled tree $\mathcal{T}_s$ represents an event $\mathcal{A} \subseteq \Omega_s$ if and only if $\mathcal{L}(\mathcal{T}_s) = \mathcal{A}$. We denote by $\mathcal{T}_{\mathcal{A},s}$ a S-labeled tree that represents an event $\mathcal{A} \subseteq \Omega_s$, and denote by $\mathrm{Tr}_{\mathcal{A},s}$ the tree of $\mathcal{T}_{\mathcal{A},s}$. ∎

Several of the following results will be phrased in terms of the S-labeled tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$, which represents the outcomes from $s \in S$ that result from player 1 using strategy σ and player 2 using strategy π, and that belong to a specified event $\mathcal{A}$.

Definition 4 (Trees for strategies) *Given a measurable event $\mathcal{A}$, strategies σ, π, a state s, such that $\Pr^{\sigma,\pi}_s(\mathcal{A}) > 0$, we denote by $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$ a S-labeled tree to represent $\mathcal{A} \cap \mathrm{Outcome}(s,\sigma,\pi)$, and we also denote by $\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ the tree of $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$. Given strategy σ,π, we denote by $\mathcal{T}^{\sigma,\pi}_s$ the S-labeled tree $\mathcal{T}^{\sigma,\pi}_{\mathrm{Outcome}(s,\sigma,\pi),s}$, and we also denote by $\mathrm{Tr}^{\sigma,\pi}_s$ the tree of $\mathcal{T}^{\sigma,\pi}_s$.* ∎

Notations. Let $\mathcal{T} = (\mathrm{Tr}, \langle\cdot\rangle)$ be a S-labeled tree and $x \in \mathrm{Tr}$ such that $|x| = n$. We denote by x_i the prefix of x of length i. We denote by $\mathrm{hist}(x) = (\langle\epsilon\rangle, \langle x_1\rangle, \ldots, \langle x_n\rangle)$, the history represented by the path from root to the node x. We denote by $\mathrm{Cone}(x) = \{ \omega = (s_0, s_1, s_2, \ldots,) \mid \langle x_i\rangle = s_i \text{ for all } 0 \le i \le n \}$ the set of paths with the prefix $\mathrm{hist}(x)$. Given a measurable event $\mathcal{A} \subseteq \Omega_s$, strategies σ and π such that $\Pr^{\sigma,\pi}_s(\mathcal{A}) > 0$, consider the S-labeled tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$ to represent $\mathcal{A} \cap \mathrm{Outcome}(s,\sigma,\pi)$. Consider the event $\mathcal{A}_{nil} = \{ \mathrm{Cone}(x) \mid x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}.\ \Pr^{\sigma,\pi}_s(\mathrm{Cone}(x) \cap \mathcal{A}) = 0 \}$. Since $\mathcal{A}_{nil}$ is the countable union of measurable sets each with measure 0 we have $\Pr^{\sigma,\pi}_s(\mathcal{A}_{nil} \cap \mathcal{A}) = 0$. Hence, in sequel without loss of generality given any event $\mathcal{A}$, we only consider the event $\mathcal{A} \setminus \mathcal{A}_{nil}$ and by a little abuse of notation use $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$ to represent the stochastic tree $\mathcal{T}^{\sigma,\pi}_{(\mathcal{A}\setminus\mathcal{A}_{nil}),s}$. Hence, without loss of generality we assume for any $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ we have $\Pr^{\sigma,\pi}_s(\mathrm{Cone}(x) \cap \mathcal{A}) > 0$. Henceforth, for any $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ we write $\Pr^{\sigma,\pi}_x(\mathcal{B} \mid \mathcal{A})$ to denote $\Pr^{\sigma,\pi}_s(\mathcal{B} \mid \mathrm{Cone}(x), \mathcal{A})$.

Definition 5 (Perennial ε-optimal strategies) *For all $\varepsilon > 0$, a strategy σ is a* perennial ε-optimal strategy *for player 1, from state s, if for all strategy π, for all node x in the stochastic tree $\mathrm{Tr}^{\sigma,\pi}_s$, we have $\Pr^{\sigma,\pi}_x(\Omega_{es}) \ge \langle\langle 1\rangle\rangle_{val}(\Omega_e)(\langle x\rangle) - \varepsilon$, i.e., in the stochastic sub-tree rooted at x player 1 is ensured the value of the game at $\langle x\rangle$ within ε-precision. Perennial ε-optimal strategies for player 2 are defined analogously. We denote by Σ_ε and Π_ε the set of perennial ε-optimal strategies for player 1 and player 2 respectively.* ∎

The ε-optimal strategies constructed for parity objectives in [8] are perennial ε-optimal strategies. This gives us the following Proposition.

Proposition 1 *For all $\varepsilon > 0$, we have $\Sigma_\varepsilon \neq \emptyset$ and $\Pi_\varepsilon \neq \emptyset$.*

3 Games with Reachability Objectives

In this section we show that the values of a concurrent parity game can be related to the ε-Nash equilibrium of a non-zero sum reachability game. This generalizes the well-known results in MDPs, stating that for all parity objectives the values of a MDP is equivalent to the value of reaching the set of states with value 1.

3.1 Non-zero sum reachability game

In sequel, we consider stochastic trees $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$ such that $\Pr^{\sigma,\pi}_s(\mathcal{A}) > 0$. Given a stochastic tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$, let κ be a subset of nodes, i.e., $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$. Analogous to the definition of reachability and safety we define the following notions of reachability and safety in the stochastic tree:

1. *Reachability in tree.* For a set $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, let

   $$\mathrm{ReachTree}(\kappa) = \{\langle\tau\rangle \mid \tau \text{ is an infinite path in } \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s} \text{ such that } \exists i \in \mathbb{N}.\ \tau_i \in \kappa\},$$

   denote the set of paths that reach the subset κ of nodes.

2. *Safety in tree.* For a set $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, let

   $$\mathrm{SafeTree}(\kappa) = \{\langle\tau\rangle \mid \tau \text{ is an infinite path in } \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s} \text{ such that } \forall i \in \mathbb{N}.\ \tau_i \in \kappa\},$$

   denote the set of paths that stay safe in the subset κ of nodes.

Given a positive integer k and a set $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, we define by $\mathrm{ReachTree}^k(\kappa) = \{\,\langle\tau\rangle \mid \exists\, x \in \tau.\ \exists\, i \leq k.\ x_i \in \kappa\,\}$, i.e., the set of paths that reaches κ within k steps.

Lemma 1 (Reachability Lemma) *Let $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$ be a stochastic tree.*

1. *For a set $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, if $\inf_{x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}} \Pr^{\sigma,\pi}_x(\mathit{ReachTree}(\kappa) \mid \mathcal{A}) > 0$, then $\Pr^{\sigma,\pi}_x(\mathit{ReachTree}(\kappa) \mid \mathcal{A}) = 1$, for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$.*

2. *For a set $U \subseteq S$, if $\inf_{x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}} \Pr^{\sigma,\pi}_x(\mathit{Reach}(U) \mid \mathcal{A}) > 0$, then $\Pr^{\sigma,\pi}_x(\mathit{Reach}(U) \mid \mathcal{A}) = 1$, for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$.*

Proof. We prove the first case and show that the second case is an immediate consequence.

1. Let $0 < c \leq \inf_{x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}} \mathrm{Pr}^{\sigma,\pi}_x(\mathrm{ReachTree}(\kappa) \mid \mathcal{A})$. Chose $0 < c' < c$. For every node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, there exists k_x such that $\mathrm{Pr}^{\sigma,\pi}_x(\mathrm{ReachTree}^{k_x}(\kappa) \mid \mathcal{A}) \geq c'$. Consider $k_1 = k_\epsilon$ (recall that ϵ is the root of the tree) and consider the frontier F_1 of $\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ at depth k_1. Given a frontier F at depth k, let $\overline{F}$ be the set of nodes x in F such that the path from the root to x has not visited a node in κ, i.e., none of $\epsilon, x_1, x_2, \ldots, x_{|x|}$ is in κ. For a frontier F_i, define $k_{i+1} = \max\{k_x \mid x \in \overline{F_i}\}$. Inductively, define the frontier F_{i+1} at depth $\sum_{j=1}^{i+1} k_j$. It follows that for $k = \sum_{i=1}^{n} k_i$ we have $\mathrm{Pr}^{\sigma,\pi}_s(\Omega \setminus \mathrm{ReachTree}^k(\kappa) \mid \mathcal{A}) \leq (1-c')^n$. Since $\lim_{n\to\infty}(1-c')^n = 0$, the desired result follows for the root of the tree. Since $\inf_{x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}} \mathrm{Pr}^{\sigma,\pi}_x(\mathrm{ReachTree}(\kappa) \mid \mathcal{A}) > 0$, it follows that for all node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ we have $\inf_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{ReachTree}(\kappa) \mid \mathcal{A}) > 0$. Arguing similarly for the subtree rooted at the node x the desired result follows.

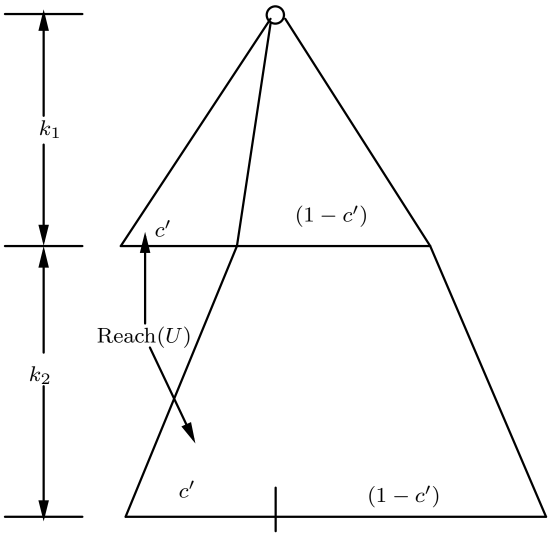

Figure 1: The Stochastic Tree for Reachability

2. Observe that with $\kappa = \{ x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s} \mid \langle x \rangle \in U \}$, we have $\mathrm{Reach}(U) = \mathrm{ReachTree}(\kappa)$. The result is immediate from part 1. ∎

Notations. Let $\mathcal{A} \subseteq \Omega_s$ be a measurable event such that $\mathrm{Pr}^{\sigma,\pi}_s(\mathcal{A}) > 0$. For a set $B \subseteq S$, let $\mathrm{InfSet}(B) = \{\omega \mid \mathrm{Inf}(\omega) \subseteq B\}$. For a set $B \subseteq S$, let $\mathrm{InfSetEq}(B) = \{\omega \mid \mathrm{Inf}(\omega) = B\}$. Given a node x in $\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, and $\varepsilon > 0$, we define $\mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x)$ as follows:

$$\mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x) = \{ B \subseteq S \mid \Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B) \mid \mathcal{A}) \geq 1 - \varepsilon \}.$$

Note that for $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$, such that $\varepsilon_1 \leq \varepsilon_2$, for any node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, if $B \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon_1}(x)$ then $B \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon_2}(x)$. We define by $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x) = \lim_{\varepsilon \to 0} \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x)$. The monotonicity property of $\mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}$ with respect to ε ensures that $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$ exists for all $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$.

Lemma 2 *For every node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, there is a unique minimal element of $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$ under the $\subset$ ordering.*

Proof. Consider a node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$. Let B_1 and B_2 be two distinct minimal elements in $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$. Consider any arbitrary $\varepsilon > 0$. It follows from the definition that we have $\Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B_i) \mid \mathcal{A}) \geq 1 - \frac{\varepsilon}{2}$, for $i \in \{ 1, 2 \}$. By definition we must have $\Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B_1 \cup B_2) \mid \mathcal{A}) \leq 1$. Hence we have the following equation:

$$\Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B_1) \mid \mathcal{A}) + \Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B_2) \mid \mathcal{A}) - \Pr^{\sigma,\pi}_x((\mathrm{InfSet}(B_1 \cap B_2)) \mid \mathcal{A}) \leq 1$$

Hence it follows that $\Pr^{\sigma,\pi}_x((\mathrm{InfSet}(B_1 \cap B_2)) \mid \mathcal{A}) \geq 1 - \varepsilon$. Hence for every $\varepsilon > 0$, we have $\Pr^{\sigma,\pi}_x(\mathrm{InfSet}(B_1 \cap B_2) \mid \mathcal{A}) \geq 1 - \varepsilon$. Hence, $B_1 \cap B_2 \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$. However, this is a contradiction to the assumption that B_1 and B_2 are distinct minimal elements of $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$. ∎

We define the function $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}} : \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s} \to 2^S$ that assigns to every node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ the minimum element of $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$. Formally, we have

$$\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x) = \bigcap_{B \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)} B = \lim_{\varepsilon \to 0} \bigcap_{B \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x)} B.$$

Proposition 2 *For every $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, for every successor x_1 of x we have $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x_1) \subseteq \mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)$.*

Proof. By definition for any nodes $x, x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, such that x_1 is a successor of x we have $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x_1) \subseteq \mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$. The result is an easy consequence of the above fact. ∎

Lemma 3 *Given a S-labeled tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$, for all node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, for all $\varepsilon > 0$, there is a set $B \subseteq S$, and $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$, such that*

$$\Pr^{\sigma,\pi}_{x_1}(\mathrm{InfSetEq}(B) \mid \mathcal{A}) \geq 1 - \varepsilon.$$

Proof. The proof is by induction on $|\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)|$.
Base Case. If $|\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)| = 1$, let $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x) = \{s\}$. Then for all nodes $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ we have $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSet}(\{s\}) \mid \mathcal{A}) \geq 1 - \varepsilon$, for all $\varepsilon > 0$. Thus for all nodes $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$, for all $\varepsilon > 0$, we have $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSetEq}(\{s\}) \mid \mathcal{A}) \geq 1-\varepsilon$.
Inductive Case. Suppose there exist a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x_1) \subsetneq \mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)$, then $|\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x_1)| < |\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)|$ and the result follows by inductive hypothesis at x_1. Otherwise for every node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ we have $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x_1) = \mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)$. Let the set $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x)$ be B. We have $\lim_{\varepsilon\to 0} \bigcap_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \left(\bigcap_{D \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x_1)} D\right) = B$.

- Suppose we have $\inf_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}(\{s\}) \mid \mathcal{A}) > 0$, for all states $s \in B$. Then it follows from Lemma 1 that for all nodes $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ we have $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}(\{ s \}) \mid \mathcal{A}) = 1$. Hence for all nodes $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ we have $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSetEq}(B) \mid \mathcal{A}) = 1$.

- Otherwise, consider a state $s \in B$ such that $\inf_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}(\{ s \}) \mid \mathcal{A}) = 0$. Hence it follows, for every $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSet}(B \setminus \{ s \}) \mid \mathcal{A}) \geq 1 - \varepsilon$. Formally, we have $\lim_{\varepsilon\to 0} \bigcap_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \left(\bigcap_{D \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x_1)} D\right) \subseteq B \setminus \{ s \}$. This is a contradiction to the fact that for all nodes $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ we have $\mathcal{M}^{\sigma,\pi}_{\mathcal{A}}(x_1) = B$ (i.e., $\lim_{\varepsilon\to 0} \bigcap_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \left(\bigcap_{D \in \mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x_1)} D\right) = B$).

The desired result follows. ∎

Lemma 4 *For every stochastic tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$, for every node $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ one of the following conditions hold:*

1. *for all $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\Omega_{es} \mid \mathcal{A}) \geq 1 - \varepsilon$;*

2. *for all $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\Omega_{os} \mid \mathcal{A}) \geq 1 - \varepsilon$.*

Proof. It follows from Lemma 3 that for all $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$, and a set B such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSetEq}(B) \mid \mathcal{A}) \geq 1-\varepsilon$. If $\min(p(B))$ is even then condition 1 is satisfied; otherwise condition 2 is satisfied. ∎

We now show that solving the zero-sum parity game is equivalent to computing the states where the value of the players are 1 and then solving some special ε-Nash equilibrium of a non-zero sum reachability game.

Consider a game graph G with winning objectives, Ω_e for player 1 and Ω_o for player 2. In sequel we denote by $W_1 = \langle\!\langle 1 \rangle\!\rangle_{limit}$ and $W_2 = \langle\!\langle 2 \rangle\!\rangle_{limit}$. We will prove that if both the player play one of their perennial ε-optimal strategies, with $\varepsilon \to 0$, then the probability of Ω_e being satisfied is equal to the probability of reaching W_1 and the probability of Ω_o being satisfied is equal to the probability of reaching W_2. For a set $T \subseteq S$ we denote by $\overline{T}$ the set $S \setminus T$. Given a state s and a set T of vertices we write $\mathrm{Safe}_s(T)$ to denote $\mathrm{Safe}(T) \cap \Omega_s$ and $\mathrm{Reach}_s(T)$ to denote $\mathrm{Reach}(T) \cap \Omega_s$.

Lemma 5 (Reachability with ε-optimal strategies) *Given a game G, consider a strategy pair $(\sigma, \pi) \in \Sigma_\varepsilon \times \Pi_\varepsilon$, with $\varepsilon \to 0$. For all states s, for all node $x \in \mathrm{Tr}_s^{\sigma,\pi}$ we have $\Pr_x^{\sigma,\pi}(\mathit{Safe}_s(\overline{W_1 \cup W_2})) = 0$.*

Proof. Let $0 < 2 \cdot \eta < \alpha = \min\{\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s), \langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o)(s) \mid s \in \overline{W_1 \cup W_2}\}$, i.e., α is the least positive value for player 1 or player 2. Consider a strategy pair $(\sigma, \pi) \in \Sigma_\eta \times \Pi_\eta$, i.e., the strategies are perennial η-optimal strategies. Let $U_s^{\sigma,\pi} = \{x \in \mathrm{Tr}_s^{\sigma,\pi} \mid s \in \overline{W_1 \cup W_2}$ and $\Pr_x^{\sigma,\pi}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) > 0\}$. If $U_s^{\sigma,\pi}$ is empty the desired result follows. Assume for the sake of contradiction that $U_s^{\sigma,\pi}$ is non-empty. Let x be a node in $U_s^{\sigma,\pi}$ and consider the S-labeled subtree $\mathcal{T}_s^{\sigma,\pi}(x)$ rooted at x. Since $\Pr_x^{\sigma,\pi}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) > 0$, we must have $\inf_{x_1 \in \mathrm{Tr}_s^{\sigma,\pi}(x)} \Pr_{x_1}^{\sigma,\pi}(\mathrm{Reach}_s(W_1 \cup W_2)) = 0$. Otherwise, it follows from Lemma 1 that if $\inf_{x_1 \in \mathrm{Tr}_s^{\sigma,\pi}(x)} \Pr_{x_1}^{\sigma,\pi}(\mathrm{Reach}_s(W_1 \cup W_2)) > 0$, then $\Pr_x^{\sigma,\pi}(\mathrm{Reach}_s(W_1 \cup W_2)) = 1$. Since $\inf_{x_1 \in \mathrm{Tr}_s^{\sigma,\pi}(x)} \Pr_{x_1}^{\sigma,\pi}(\mathrm{Reach}_s(W_1 \cup W_2)) = 0$ we have $\sup_{x_1 \in \mathrm{Tr}_s^{\sigma,\pi}(x)} \Pr_{x_1}^{\sigma,\pi}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) = 1$. Consider a node $x_1 \in \mathrm{Tr}_s^{\sigma,\pi}(x)$ such that $\Pr_{x_1}^{\sigma,\pi}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) \geq 1 - \eta$. Let $\mathcal{A}$ be the event $\mathrm{Safe}_s(\overline{W_1 \cup W_2})$. Since σ and π are perennial η-optimal strategy, and $\Pr_{x_1}^{\sigma,\pi}(\mathcal{A}) \geq 1 - \eta$, it follows that for every node $x_2 \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}(x_1)$ we have $\Pr_{x_2}^{\sigma,\pi}(\Omega_{es} \mid \mathcal{A}) \geq c_1 \geq (\alpha - 2\eta) > 0$ and $\Pr_{x_2}^{\sigma,\pi}(\Omega_{os} \mid \mathcal{A}) \geq c_2 \geq (\alpha - 2\eta) > 0$. This implies that for all node $x_2 \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}(x_1)$ we have $\Pr_{x_2}^{\sigma,\pi}(\Omega_{es} \mid \mathcal{A}) \leq 1 - c_2$ and $\Pr_{x_2}^{\sigma,\pi}(\Omega_{os} \mid \mathcal{A}) \leq 1 - c_1$. It follows from Lemma 4 that for every $\varepsilon > 0$, there is a node $x_2 \in \mathrm{Tr}_s^{\sigma,\pi}(x_1)$ such that either $\Pr_{x_2}^{\sigma,\pi}(\Omega_{es} \mid \mathcal{A}) \geq 1 - \varepsilon$ or $\Pr_{x_2}^{\sigma,\pi}(\Omega_{os} \mid \mathcal{A}) \geq 1 - \varepsilon$. Since c_1 and c_2 are constants greater than 0, we have a contradiction. Hence $U_s^{\sigma,\pi} = \emptyset$ and the Lemma follows. ∎

Lemma 6 *Given a game G, let the winning objectives of player 1 and player 2 be Ω_e and Ω_o, respectively. Then*

$$\lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\mathit{Reach}_s(W_1)) = \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s)$$

$$\lim_{\varepsilon \to 0} \sup_{\pi \in \Pi_\varepsilon} \inf_{\sigma \in \Sigma_\varepsilon} \Pr_s^{\sigma,\pi}(\mathit{Reach}_s(W_2)) = \langle\!\langle 2 \rangle\!\rangle_{val}(\Omega_o)(s)$$

Proof. Given any strategy σ and π we have the following equality:

$$\begin{array}{rcl} \Pr_s^{\sigma,\pi}(\Omega_{es}) & = & \Pr_s^{\sigma,\pi}(\Omega_{es} \ \cap \ \mathrm{Safe}_s(\overline{W_1 \cup W_2})) \\ & + & \Pr_s^{\sigma,\pi}(\Omega_{es} \ \cap \ \mathrm{Reach}_s(W_1 \cup W_2)) \end{array}$$

It follows from the definition of ε-optimal strategies and determinacy of parity games [15, 8] that for all state s we have $\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) = \lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es})$. For any state s have the following containment relation: $\Omega_{es} \cap \mathrm{Safe}_s(\overline{W_1 \cup W_2}) \subseteq \mathrm{Safe}_s(\overline{W_1 \cup W_2})$. It follows from Lemma 5 that $\lim_{\varepsilon \to 0} \sup_{(\sigma,\pi) \in \Sigma_\varepsilon \times \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \cap \mathrm{Safe}_s(\overline{W_1 \cup W_2})) = 0$. Hence we have

$$\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) = \lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \ \cap \ \mathrm{Reach}_s(W_1 \cup W_2))$$

Since σ and π are ε-optimal strategies we have the following two facts:

$$\lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \ \cap \ \mathrm{Reach}_s(W_2)) = 0$$

$$\lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \mid \mathrm{Reach}_s(W_1)) = 1.$$

This gives us the following equality:

$$\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) = \lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \ \cap \ \mathrm{Reach}_s(W_1))$$

The right hand side of the equality can be expressed as

$$\lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\Omega_{es} \mid \mathrm{Reach}_s(W_1)) \Pr_s^{\sigma,\pi}(\mathrm{Reach}_s(W_1))$$

$$= \lim_{\varepsilon \to 0} \sup_{\sigma \in \Sigma_\varepsilon} \inf_{\pi \in \Pi_\varepsilon} \Pr_s^{\sigma,\pi}(\mathrm{Reach}_s(W_1)).$$

This gives us the desired result. ∎

Consider the following variants of the game G, a game G_A and G_R as follows, with the same state space as G and the states in W_1 and W_2 changed to absorbing states. G_A is a zero-sum parity game and the priority for each state in W_1 is 0 and for each state in W_2 is 1, and for all the other states is same as the priority of the game G. Note that for every state s the value for player 1 and player 2 for the game G and G_A are the same. The game G_R is a non-zero sum reachability game and the winning objectives of both the players are reachability objectives: the objective for player 1 is $\mathrm{Reach}(W_1)$ and the objective for player 2 is $\mathrm{Reach}(W_2)$. Note that the game G_R is not zero-sum in the following sense: there are infinite paths ω such that $\omega \notin \mathrm{Reach}(W_1)$ and $\omega \notin \mathrm{Reach}(W_2)$ and each player gets a payoff 0 for the path ω. We define ε-Nash equilibrium of the game G_R and relate some special ε-Nash equilibrium of G_R with the values of G.

Definition 6 (ε-Nash equilibrium in G_R) *A strategy profile $(\sigma^*, \pi^*) \in \Sigma \times \Pi$ is an ε-Nash equilibrium at state s if the following two conditions hold:*

$$\forall \sigma \in \Sigma.\ \Pr_s^{\sigma^*,\pi^*}(Reach_s(W_1)) \geq \Pr_s^{\sigma,\pi^*}(Reach_s(W_1)) - \varepsilon$$

$$\forall \pi \in \Pi.\ \Pr_s^{\sigma^*,\pi^*}(Reach_s(W_2)) \geq \Pr_s^{\sigma^*,\pi}(Reach_s(W_2)) - \varepsilon \ \blacksquare$$

Theorem 1 (Nash equilibrium of reachability game G_R associated with the parity game G) *The following assertion hold for the game G_R.*

1. *For all $\varepsilon > 0$, there is an ε-Nash equilibrium $(\sigma_\varepsilon^*, \pi_\varepsilon^*)$ such that for all states s we have*

$$\lim_{\varepsilon \to 0} \Pr_s^{\sigma_\varepsilon^*,\pi_\varepsilon^*}(Reach_s(W_1)) = \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s)$$

$$\lim_{\varepsilon \to 0} \Pr_s^{\sigma_\varepsilon^*,\pi_\varepsilon^*}(Reach_s(W_2)) = \langle\langle 2 \rangle\rangle_{val}(\Omega_o)(s).$$

Proof. It follows from Lemma 6 and Proposition 1. $\blacksquare$

4 Strategy Characterization and Computational Complexity

In this section we construct polynomial witnesses for perennial ε-optimal strategies and describe polynomial procedure to verify the witnesses. An immediate consequence is the fact that the values of concurrent parity games can be decided within ε-precision in NP $\cap$ coNP. Since the values can be irrational, one can only hope to ε-approximate the values. Our proof techniques reveals several key characteristics of the perennial ε-optimal strategies. In general perennial ε-optimal strategies require infinite memory in general [6, 8]. We show that though the perennial ε-optimal strategies require infinite memory in general, there exist perennial ε-optimal strategies that in limit coincide with some memoryless strategies. This result parallels with the celebrated result of Mertens-Neyman [17] for concurrent games with limit-average objectives, that states there exists ε-optimal strategies that in limit coincide with some memoryless strategies (the memoryless strategy correspond to the memoryless optimal strategies in the discounted game with discount factor tends to 0). It may be noted that the memoryless strategies that the perennial ε-optimal strategies coincide, is itself not necessarily ε-optimal.

In concurrent games with safety objective optimal memoryless strategies exist, and the optimal strategies in general require randomization [11]. In case of concurrent games with reachability objectives optimal strategies need not exist, but memoryless ε-optimal strategies exist for all $\varepsilon > 0$ [11], and the ε-optimal strategies require randomization. In case of concurrent games with Büchi objectives, ε-optimal strategies require infinite memory in general [6]. In contrast we show that for all $\varepsilon > 0$, memoryless ε-optimal strategies exit for all concurrent games with coBüchi objectives. It follows from the simpler case of reachability objectives that optimal strategies need not exist and ε-optimal strategies require randomization. It follows from the results on Büchi objectives that in concurrent games with parity objectives with 3 or more priorities ε-optimal strategies require infinite memory in general. Our result thus completes the precise memory requirements of ε-optimal strategies in concurrent parity games.

4.1 Reduction to Qualitative Witness

The notion of *local optimality* will play an important role in our construction of polynomial witnesses. Informally, a selector function ξ is *locally optimal* if it is optimal in the one-step matrix game where each state is assigned a reward value $\langle\langle 1\rangle\rangle_{val}(\Omega_e)(s)$. A *locally optimal strategy* is a strategy that consists of locally optimal selectors. A *locally ε-optimal strategy* is a strategy that has a total deviation from locally-optimal selectors of at most ε. Locally optimal selectors and strategies play a role in the construction of polynomial witnesses, since local optimality is a notion that can be checked in polynomial time.

We note that *local ε-optimality* and *ε-optimality* are very different notions. *Local ε-optimality* consists in the approximation of a local selector; a locally ε-optimal strategy provides no guarantee of yielding a probability of winning the game close to the optimal one. On the other hand, a ε-optimal strategy is a strategy that guarantees a probability of winning close to the optimal one; there are no constraints on its local structure. The construction of polynomial witnesses will depend on constructing a relation between the notion of local ε-optimality (which is polynomially checkable) and global ε-optimality (which yields a value close to the value of the game).

Definition 7 (Locally ε-optimal selectors and strategies) *A selector ξ is* locally optimal *if for all $s \in S$ and $a_2 \in \Gamma_s(s)$ we have*

$$\mathrm{E}[\langle\langle 1\rangle\rangle_{val}(\Omega_e)(\Theta_1) \mid s, \xi(s), a_2] \geq \langle\langle 1\rangle\rangle_{val}(\Omega_e)(s).$$

We denote by Λ^ℓ the set locally-optimal selectors.

A strategy σ is locally optimal *if for every history $\langle s_0, s_1, \ldots, s_k\rangle$ we have $\sigma(s_0, s_1, \ldots, s_k) \in \Lambda^\ell$, i.e., player 1 plays a locally optimal selector at every stage of the play. We denote by Σ^ℓ the set of locally optimal strategies.*

A strategy σ_ε is locally ε-optimal *if for every strategy $\pi \in \Pi$ and for every $\omega = \langle s_0, s_1, s_2, \ldots, \rangle \in \mathrm{Outcome}(s, \sigma_\varepsilon, \pi)$ we have*

$$\sum_{k=0}^{\infty} \big(\max\{(\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s_k) - \mathrm{E}[\langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(\Theta_{k+1}) \mid s_k, \sigma_\varepsilon(\omega_k), \pi(\omega_k)]), 0\} \big) \leq \varepsilon,$$

where $\omega_k = \langle s_0, s_1, \ldots, s_k\rangle$. We denote by Σ^ℓ_ε the set of locally ε-optimal strategies. ■

A *value class* of the game is the set of all states where the game has a given value.

Definition 8 (Value class) *A value class $\mathrm{VC}(r)$ is the set of states s such that the value for player 1 is r. Formally, $\mathrm{VC}(r) = \{s \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) = r\}$. Note that for any game there are at most $|S|$ many value classes. By $\mathrm{VC}^{<r}$ we denote the set $\{s \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) < r\}$ and similarly we use $\mathrm{VC}^{>r}$ to denote the set $\{s \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Omega_e)(s) > r\}$.* ■

Intuitively, we can picture the game as a "quilt" of value classes. Two of the value classes correspond to values 1 (player 1 wins with probability arbitrarily close to 1) and 0 (player 2 wins with probability arbitrarily close to 1); the other value classes correspond to intermediate values. We construct a polynomial witness in a piece-meal fashion. We first show that we can construct, for each intermediate value class, a strategy that with probability arbitrarily close to 1 guarantees either leaving the class, or winning without leaving the class. Such a strategy can be constructed using results from [6], and has a concise (polynomial) witness. Second, we show that the above strategy can be constructed so that when the class is left, it is left via a locally ε-optimal selector. By stitching together the strategies constructed in this fashion for the various value classes, we will obtain a single polynomial witness for the complete game. The construction of a strategy in a value class relies on the following reduction.

Reduction. For a state s we define the set of allowable actions as follows

$$\begin{array}{rcl} \mathrm{AllowActs}(s) & = & \{\gamma \subseteq \Gamma_1(s) : \text{ such that there is an optimal selector} \\ & & \xi_1^\ell \in \Lambda^\ell \text{ and } \mathrm{Supp}(\xi_1^\ell) = \gamma\} \end{array}$$

Let $G = (S, Moves, \Gamma_1, \Gamma_2, \delta)$ be a concurrent game with parity objectives Parity(p) and coParity(p) for player 1 and player 2 respectively, and let the priority function be p. Consider a value class VC(r) with $0 < r < 1$. We construct a concurrent game $\widetilde{G}_r = (\widetilde{S}_r, \widetilde{Moves}, \widetilde{\Gamma_1}, \widetilde{\Gamma_2}, \widetilde{\delta})$ with a priority function $\widetilde{p}$ as follows:

1. **State space.** Given a state s let $\mathrm{AllowActs}(s) = \{\gamma_1, \gamma_2, \ldots, \gamma_k\}$. Then we have

$$\begin{array}{rcl} \widetilde{S}_r & = & \{\, \widetilde{s} \mid s \in \mathrm{VC}(r) \,\} \cup \{\, w_1, w_2 \,\} \\ & & \cup \{\, (\widetilde{s}, i) \mid s \in \mathrm{VC}(r), i \in \{\, 1, 2, \ldots, k \,\} \text{ and } \mathrm{AllowActs}(s) = \{\gamma_1, \gamma_2, \ldots, \gamma_k\} \,\} \end{array}$$

2. **Priority function.**
   (a) $\widetilde{p}(\widetilde{s}) = p(s)$ for all $s \in \mathrm{VC}(r)$.
   (b) $\widetilde{p}((\widetilde{s}, i)) = p(s)$ for all $(\widetilde{s}, i) \in \widetilde{S}_r$.
   (c) $\widetilde{p}(w_1) = 0$ and $\widetilde{p}(w_2) = 1$.

3. **Moves assignment.**
   (a) $\widetilde{\Gamma_1}(\widetilde{s}) = \{\, 1, 2, \ldots, k \,\}$ such that $\mathrm{AllowActs}(s) = \{\gamma_1, \gamma_2, \ldots, \gamma_k\}$ and $\widetilde{\Gamma_2}(\widetilde{s}) = \{a_2\}$. Note that every $\widetilde{s} \in \widetilde{S}_r$ is a player-1 turn-based state.
   (b) $\widetilde{\Gamma_1}((\widetilde{s}, i)) = \{i\} \cup (\Gamma_1(s) \setminus \gamma_i)$ and $\widetilde{\Gamma_2}((\widetilde{s}, i)) = \Gamma_2(s)$. At state $(\widetilde{s}, i)$ all the moves in γ_i are collapsed to one move i and the moves not in γ_i exist in the set of available moves.

4. **Transition function.**
   (a) The states w_1 and w_2 are absorbing states. Observe that player 1 have value 1 and 0 at state w_1 and w_2 respectively.
   (b) For any state $\widetilde{s}$ we have $\widetilde{\delta}(\widetilde{s}, i, a_2)((\widetilde{s}, i)) = 1$. Hence at state $\widetilde{s}$ player 1 can decide which element in AllowActs(s) to play and if player 1 chooses move i the game proceed to state $(\widetilde{s}, i)$.
   (c) *Transition function at state* $(\widetilde{s}, i)$.
      i. For any move $a_2 \in \Gamma_2(s)$, if there is a move $a_1 \in \gamma_i$ such that $\sum_{s' \notin \mathrm{VC}(r)} \delta(s, a_1, a_2)(s') > 0$, then $\widetilde{\delta}((\widetilde{s}, i), i, a_2)(w_1) = 1$. The above transition specifies that if for a move a_2 for player 2 and a move $a_1 \in \gamma_i$ for player 1, if the game G proceeds to a different value class with positive probability then

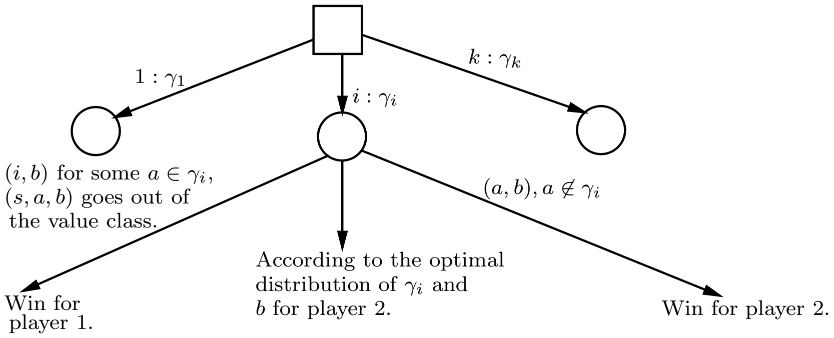

Figure 2: Reduction to limit-sure

in $\widetilde{G}_r$ the game proceeds to the state w_1, which has value 1 for player 1, with probability 1. Note, that since $a_1 \in \gamma_i$ and $\gamma_i \in \text{AllowActs}(s)$, if in G the game proceeds to a different value class with positive probability it also proceeds to $\text{VC}^{>r}$ with positive probability.

ii. For any move $a_2 \in \Gamma_2(s)$, if for every move $a_1 \in \gamma_i$ we have $\sum_{s' \in \text{VC}(r)} \delta(s, a_1, a_2)(s') = 1$ then

$$\widetilde{\delta}((\widetilde{s}, i), i, a_2)(\widetilde{s'}) = \sum_{a_1 \in \gamma_i} \xi_1^\ell(a_1) \cdot \delta(s, a_1, a_2)(s')$$

where ξ_1^ℓ is an locally optimal selector with $\text{Supp}(\xi_1^\ell) = \gamma_i$.

iii. For any move $a_1 \in (\Gamma_1(s) \setminus \gamma_i)$, for any move $a_2 \in \Gamma_2(s)$ we have:

$$\widetilde{\delta}((\widetilde{s}, i), a_1, a_2)(\widetilde{s'}) = \delta(s, a_1, a_2)(s') \text{ for } s' \in \text{VC}(r);$$

$$\widetilde{\delta}((\widetilde{s}, i), a_1, a_2)(w_2) = \sum_{s' \notin \text{VC}(r)} \delta(s, a_1, a_2)(s').$$

The reduction is illustrated in Fig. 2.

Fact 1. If player 1 follows a strategy σ_ε such that at any state $(\widetilde{s}, i)$ it plays action i with probability 1 then for every strategy π for player 2 we have $\text{Pr}_{\widetilde{s}}^{\sigma_\varepsilon,\pi}(\text{Reach}(w_2)) = 0$.

Lemma 7 *For every $r > 0$, for any state $s \in \text{VC}(r)$, the state $\widetilde{s}$ is limit-sure winning in the game $\widetilde{G}_r$ for player 1, i.e., from state $\widetilde{s}$ player 1 can win with probability arbitrarily close to 1.*

Proof. Let σ_ε be a locally ε-optimal and perennial ε-optimal strategy in G, i.e., $\sigma_\varepsilon \in \Sigma_\varepsilon^\ell \cap \Sigma_\varepsilon$ (the fact that $\Sigma_\varepsilon^\ell \cap \Sigma_\varepsilon \neq \emptyset$ follows from the results of [8]). Assume for the sake of contradiction that $U \subseteq \widetilde{S}_r \cap \{ \widetilde{s} \mid s \in \mathrm{VC}(r) \}$ is a non-empty set such that player 2 wins with bounded-positive probability. Let $\widetilde{\pi}$ be a perennial bounded-positive optimal strategy for player 2 from the set U. We construct a *projected* strategy $\widetilde{\sigma_\varepsilon}$ for player 1 in $\widetilde{G}_r$ and a *extended* strategy π_ε for player 2 in G as follows:

1. Strategy $\widetilde{\sigma_\varepsilon}$ in the game $\widetilde{G}_r$:
   - $\widetilde{\sigma_\varepsilon}(\widetilde{s_0}, (\widetilde{s_0}, i_0), \widetilde{s_1}, (\widetilde{s_1}, i_1), \ldots, \widetilde{s_k})(j) = 1$ if and only if $\gamma_j = \arg\max_{\gamma \in \mathrm{AllowActs}(s_k)} \sum_{a \in \gamma} \sigma_\varepsilon(s_0, s_1, \ldots, s_k)(a)$.
   - $\widetilde{\sigma_\varepsilon}(\widetilde{s_0}, (\widetilde{s_0}, i_0), \widetilde{s_1}, (\widetilde{s_1}, i_1), \ldots, \widetilde{s_k}, (\widetilde{s_k}, j))(j) = \sum_{a \in \gamma_j} \sigma_\varepsilon(s_0, s_1, \ldots, s_k)(a)$ and for all $a' \notin \gamma_j$ we have $\widetilde{\sigma_\varepsilon}(\widetilde{s_0}, (\widetilde{s_0}, i_0), \widetilde{s_1}, (\widetilde{s_1}, i_1), \ldots, \widetilde{s_k}, (\widetilde{s_k}, j))(a') = \sigma_\varepsilon(s_0, s_1, \ldots, s_k)(a')$.

2. Strategy π_ε in the game G:
   - $\pi_\varepsilon(s_0, s_1, \ldots, s_k) = \widetilde{\pi}(\widetilde{s_0}, (\widetilde{s_0}, i_0), \widetilde{s_1}, (\widetilde{s_1}, i_1), \ldots, \widetilde{s_k}))$ such that for all $0 \le l \le k$, we have $\widetilde{\sigma_\varepsilon}(\widetilde{s_0}, (\widetilde{s_0}, i_0), \widetilde{s_1}, (\widetilde{s_1}, i_1), \ldots, \widetilde{s_l}))(i_l) = 1$.

Given a set of states $\widetilde{C} \subseteq \widetilde{S}_r \setminus \{ w_1, w_2 \}$ we denote by $C_G = \{ s \mid \widetilde{s} \in \widetilde{C}$ or, for some i. $(\widetilde{s}, i) \in \widetilde{C} \}$. Suppose, for some state $\widetilde{s}$ we have $\Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\mathrm{Safe}_{\widetilde{s}}(\widetilde{C})) > 0$, for some set $\widetilde{C} \subseteq \widetilde{S}_r \setminus \{ w_1, w_2 \}$. Then by construction of π_ε we have $\Pr_s^{\sigma_\varepsilon, \pi_\varepsilon}(\mathrm{Safe}_s(C_G)) > 0$. It follows from argument similar to Lemma 5 that there is a node $x \in \mathrm{Tr}_s^{\sigma_\varepsilon, \pi_\varepsilon}$ such that $\Pr_x^{\sigma, \pi}(\mathrm{Safe}_s(C_G)) \ge 1 - \varepsilon'$, with $\varepsilon' \to 0$. Let us denote by $\mathcal{A}$ the event $\mathrm{Safe}_s(C_G)$. Note that for event $\mathcal{A}$, the strategy pair $(\sigma_\varepsilon, \pi_\varepsilon)$ is well-defined. Since σ_ε is a perennial ε-optimal strategy, for all nodes $x_1 \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma_\varepsilon, \pi_\varepsilon}(x)$ we have $\Pr_{x_1}^{\sigma_\varepsilon, \pi_\varepsilon}(\Omega_{os} \mid \mathcal{A}) \le c_2$, for $c_2 < 1$. Since $\widetilde{\pi}$ is a perennial bounded positive optimal strategy in $\widetilde{G}_r$ for all nodes $x_1 \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma_\varepsilon, \pi_\varepsilon}(x)$ we have $\Pr_{x_1}^{\sigma_\varepsilon, \pi_\varepsilon}(\Omega_{es} \mid \mathcal{A}) \le c_1$ and for $c_1 < 1$. However, this is a contradiction to Lemma 4. Hence for every state $\widetilde{s} \in \widetilde{S}_r \setminus \{ w_1, w_2 \}$ we have $\Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\mathrm{Safe}_{\widetilde{s}}(\widetilde{S}_r \setminus \{ w_1, w_2 \})) = 0$, and hence $\Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\mathrm{Reach}_{\widetilde{s}}(\{ w_1, w_2 \})) = 1$. Since, $\sigma_\varepsilon \in \Sigma_\varepsilon^\ell$ it follows from the construction of the game $\widetilde{G}_r$, Fact 1. and the property of locally ε-optimal strategies that $\Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\mathrm{Reach}_{\widetilde{s}}(\{ w_2 \})) \le \varepsilon$. Thus $\Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\Omega_e \cap \Omega_{\widetilde{s}}) \ge \Pr_{\widetilde{s}}^{\widetilde{\sigma_\varepsilon}, \widetilde{\pi}}(\mathrm{Reach}_{\widetilde{s}}(\{ w_1 \})) \ge 1 - \varepsilon$. This is a contradiction to the assumption that $\widetilde{\pi}$ is bounded positive optimal. ∎

Limit-sure witness [6]. The witness strategy σ for a limit-sure game constructed in [6] consists of the following parts: a *ranking function* of the states, and a *ranking function* of the actions at a state. The ranking functions were described by a μ-calculus formula. The witness strategy σ at round k of a play, at a state s, plays the actions of the least rank at s with positive-bounded probabilities and other actions with vanishingly small probabilities (as function of ε), in appropriate proportion as described by the ranking function. Hence, the strategy σ can be described as

$$\sigma = (1 - \varepsilon_k)\sigma_\ell + \varepsilon_k \cdot \sigma_d(\varepsilon_k),$$

where σ_ℓ is any selector with ξ such that $\mathrm{Supp}(\xi)$ is the set actions with least rank, and $\sigma_d(\varepsilon_k)$ denotes a selector with $\mathrm{Supp}(\sigma_d(\varepsilon_k)) = \Gamma_1 \setminus \mathrm{Supp}(\sigma_\ell)$. Hence the strategy σ plays the moves in $\mathrm{Supp}(\sigma_d(\varepsilon_k))$ with vanishingly small probability as $\varepsilon_k \to 0$. We denote by *limit-sure witness move set* the set of actions with the least rank, i.e., $\mathrm{Supp}(\sigma_\ell)$. It follows from the above construction that as $\varepsilon \to 0$, the limit-sure winning strategy σ converges to the memoryless selector σ_ℓ, i.e., the limit of the limit-sure witness strategy is a memoryless strategy.

Lemma 8 *At any state $(\widetilde{s}, i)$, if the limit-sure witness move set for player 1 is γ, if $(\gamma \setminus \{ i \}) \neq \emptyset$, then $(\gamma \setminus \{ i \}) \in \mathrm{AllowActs}(s)$.*

Proof. Consider a move $a \in \gamma \setminus \{ i \}$. If there is a move $b \in \Gamma_2(s)$ such that $\widetilde{\delta}((\widetilde{s}, i), a, b)(w_2) > 0$, we would obtain a contraction to the hypothesis that player 1, at $(\widetilde{s}, i)$, wins with probability arbitrarily close to 1. Hence, we have $\mathrm{Dest}(s, a, b) \subseteq \mathrm{VC}(r)$ for every move $b \in \Gamma_2(s)$, leading to the result. ∎

Lemma 9 (Union-closure of $\mathrm{AllowActs}(s)$) *For all state s, if $\gamma_1 \in \mathrm{AllowActs}(s)$ and $\gamma_2 \in \mathrm{AllowActs}(s)$, then $\gamma_1 \cup \gamma_2 \in \mathrm{AllowActs}(s)$.*

Proof. It follows from the properties of "one-step" matrix games that if ξ_1 and ξ_2 are optimal strategies for a player, then any convex combination of ξ_1 and ξ_2 is also an optimal strategy. Thus it follows that if $\xi_1 \in \Lambda^\ell$ and $\xi_2 \in \Lambda^\ell$, then there exist $\xi \in \Lambda^\ell$ such that $\mathrm{Supp}(\xi) = \mathrm{Supp}(\xi_1) \cup \mathrm{Supp}(\xi_2)$. The lemma follows. ∎

Lemma 10 *At any state $(\widetilde{s}, i)$, if the limit-sure witness move set for player 1 is γ, then $\gamma_j = \big((\gamma \setminus \{ i \}) \cup \gamma_i\big) \in \mathrm{AllowActs}(s)$.*

Proof. The Lemma follows from Lemma 8 and Lemma 9. ∎

Lemma 11 *For every state $\widetilde{s}$ there is a pure memoryless move j for player 1 and limit-sure winning strategy σ such that* $\mathrm{Supp}(\sigma)(\widetilde{s}) = \{ j \}$ *and the limit-sure witness move set at* $(\widetilde{s}, j) = \{ j \}$.

Proof. The existence of pure memoryless move is a consequence of the fact that every state $\widetilde{s}$ is a player-1 turn-based state and the witness construction in [6]. The rest follows from Lemma 10. ∎

Definition 9 (Value-class qualitative optimal strategy) *A strategy σ_ε is a value-class qualitative optimal strategy for a value-class* $\mathrm{VC}(r)$, *with* $0 < r < 1$, *if*

1. *σ_ε is locally ε-optimal.*
2. *Let π be an arbitrary strategy for player 2. For a state $s \in \overline{W_1 \cup W_2}$, for all node x in $\mathrm{Tr}_s^{\sigma_\varepsilon,\pi}$ such that $\langle x \rangle \in \mathrm{VC}(r)$, $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathit{Safe}(\mathrm{VC}(r))) \geq 1 - \varepsilon$.*

A strategy σ_ε is value-class qualitative optimal if it is value-class qualitative optimal for all value class $0 < r < 1$. ∎

The existence of value-class qualitative optimal strategies follows from Lemma 7 and Lemma 11.

Lemma 12 *The set of value-class qualitative optimal strategies is non-empty.*

Lemma 13 *Let σ_ε be a locally ε-optimal strategy. For all strategy π for player 2, for all node $x \in \mathrm{Tr}_s^{\sigma_\varepsilon,\pi}$, if $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathit{Reach}(W_1 \cup W_2)) = 1$, then $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathit{Reach}(W_1)) \geq \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\langle x \rangle) - \varepsilon$.*

Proof. The results then follows from the fact that the sequence $(\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\Theta_i))_i$ is a sub-martingale under σ_ε and π. ∎

The following Lemma shows that the value-class qualitative optimal strategies for different value classes can be "stitched" or composed together to produce a perennial ε-optimal strategy. This will allow us to produce witness for individual value classes and compose them to obtain a witness for perennial ε-optimal strategy.

Lemma 14 (Stitching Lemma) *Let σ_ε be a value-class qualitative optimal strategy and perennial ε-optimal for all state in W_1. Then σ_ε is a perennial ε-optimal strategy.*

Proof. Consider any strategy π for player 2 and consider the stochastic tree $\mathcal{T}_s^{\sigma_\varepsilon,\pi}$ for any state s. For a node x we define the set $\mathrm{SafeVal}(x) = \{\ \omega = \langle s_0, s_1, \ldots, \rangle \in \mathrm{Cone}(x) \mid \forall\ k \geq |x|.s_k \in \mathrm{VC}(r),\ \text{where } \langle x \rangle \in \mathrm{VC}(r)\ \}$ as the set of paths that stays safe in the value class $\mathrm{VC}(r)$ of $\langle x \rangle$ from x. Note that $\mathrm{Cone}(x) \setminus \mathrm{SafeVal}(x)$ denotes the set of paths that leaves the value class $\mathrm{VC}(r)$ from x. Let $\alpha = \max\{\ \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) \mid s \in (S \setminus W_1\ \}$, i.e., α is the maximum value for player 1 that is less than 1. Consider the following set of nodes

$$\kappa_1 = \{\ x \in \mathrm{Tr}_s^{\sigma_\varepsilon,\pi} \mid \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{SafeVal}(x))) \geq \alpha\ \}$$
$$\kappa_2 = \{\ x \in \mathrm{Tr}_s^{\sigma_\varepsilon,\pi} \mid \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{Cone}(x) \setminus \mathrm{SafeVal}(x)) > 1 - \alpha\ \}$$

Note that $\kappa_1 = \mathrm{Tr}_s^{\sigma_\varepsilon,\pi} \setminus \kappa_2$ and hence for any node $x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}$ we have $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{ReachTree}(\kappa_1)) + \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{SafeTree}(\kappa_2)) = 1$. Consider the event $\mathcal{A} = \mathrm{SafeTree}(\kappa_2)$. Since σ_ε is a locally ε-optimal strategy it follows that if a play leaves a value class $\mathrm{VC}(r)$ with probability at least $(1 - \alpha) > 0$, then it reaches $\mathrm{VC}^{>r}$ with positive bounded probability. It follows that $\kappa_2 \subseteq \{\ x \mid \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{Reach}(W_1 \cup W_2)) \geq c > 0\ \}$. Hence, it follows that for all node $x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma_\varepsilon,\pi}$ we have $\inf_{x_1 \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma_\varepsilon,\pi}} \mathrm{Pr}_{x_1}^{\sigma_\varepsilon,\pi}(\mathrm{Reach}(W_1 \cup W_2) \mid \mathcal{A}) > 0$. It follows from Lemma 1 that for all node $x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma_\varepsilon,\pi}$ we have $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{Reach}(W_1 \cup W_2) \mid \mathcal{A}) = 1$. Since σ_ε is locally ε-optimal, it follows from Lemma 13 that

$$\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathrm{Reach}(W_1 \cup W_2)) \geq \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{Reach}(W_1)) \geq \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\langle x \rangle) - \varepsilon.$$

Since σ_ε is a value-class qualitative optimal strategy we have $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathrm{Safe}(\mathrm{VC}(r)) \geq (1-\varepsilon)$. Therefore, for all node x in κ_1 we have $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es}) \geq \alpha \cdot (1 - \varepsilon) > \alpha - \varepsilon$, since $\alpha < 1$. Thus for all node x we have, $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathrm{ReachTree}(\kappa_1)) > \alpha - \varepsilon$. For all node x we have

$$\begin{array}{rcl}
\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es}) & \geq & \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathrm{SafeTree}(\kappa_2)) \cdot \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{SafeTree}(\kappa_2)) \\
& + & \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathrm{ReachTree}(\kappa_1)) \cdot \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{ReachTree}(\kappa_1)) \\
& \geq & (\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\langle x \rangle) - \varepsilon) \cdot \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{SafeTree}(\kappa_2)) \\
& + & (\alpha - \varepsilon) \cdot \mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\mathrm{ReachTree}(\kappa_1))
\end{array}$$

Since $\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\langle x \rangle) \leq \alpha$ we have $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es}) \geq \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(\langle x \rangle) - \varepsilon$. Hence σ_ε is a perennial ε-optimal strategy. ∎

The following Theorem follows from existence of memoryless limit-sure winning strategies for concurrent games with coBüchi objectives [6] and the existence of perennial ε-optimal strategies obtained by composing value-class qualitative optimal strategies across value classes (Lemma 14).

Theorem 2 (Memoryless ε-optimal strategies for coBüchi objectives) *For every $\varepsilon > 0$, memoryless ε-optimal strategies exist for all coBüchi objectives on all concurrent games.*

The following Theorem states that there exist perennial ε-optimal strategies that in limit coincide with locally optimal selector, i.e., a memoryless strategy with locally optimal selectors. This parallels the results of Mertens-Neyman [17] for concurrent games with limit-average objectives.

Theorem 3 (Limit of ε-optimal strategies) *For every $\varepsilon > 0$, there exist perennial ε-optimal strategy $\sigma_\varepsilon \in \Sigma_\varepsilon$, such that the sequence of the strategies σ_ε converge to a locally optimal selector $\overline{\sigma}$ as $\varepsilon \to 0$, i.e., $\lim_{\varepsilon \to 0} \sigma_\varepsilon = \overline{\sigma}$, where $\overline{\sigma} \in \Sigma^\ell$ and $\overline{\sigma}$ is memoryless.*

Proof. For arbitrary $\varepsilon > 0$, consider the perennial ε-optimal strategy σ_ε constructed as a value-class qualitative optimal strategy. The fact that the value-class qualitative optimal strategy is a perennial ε-optimal strategy follows from Lemma 14. The result then follows from Lemma 11 and the fact that the limit-sure winning strategies coincide in limit with a memoryless selector σ_ℓ such that $\mathrm{Supp}(\sigma_\ell)$ is the set of least-rank actions of the limit-sure witness. ∎

Witness for perennial ε-optimal strategies. The witness for a perennial ε-optimal strategy σ_ε is presented as a value-class qualitative optimal strategy (recall Lemma 14). The existence of a value-class qualitative optimal strategy is guaranteed by Lemma 12. The witness consists of the limit-sure winning strategy witness in the game $\widetilde{G}_r$, for all $0 < r < 1$, and of a locally ε-optimal strategy. The witness can be described as follows:

- *Limit-sure witness.* The limit-sure witness in the game $\widetilde{G}_r$, for $r > 0$, is constructed as the the witness described in [6]. Observe that the game $\widetilde{G}_r$ can be exponential in the size of the game G, since the set $\mathrm{AllowActs}(s)$ can be exponential. To obtain efficient polynomial witness we make the following key observation: at every state $\widetilde{s}$ there is a pure memoryless move i for player 1 (Lemma 11) in the limit-sure witness strategy. Hence player 1 constructs a game $\widetilde{G}'_r$ such that every state $\widetilde{s}$ there is only a single successor $(\widetilde{s}, i)$, where i is a pure memoryless move in the limit-sure witness in $\widetilde{G}_r$. The graph $\widetilde{G}'_r$ is linear in the size of the game G. The witness in state $(\widetilde{s}, i)$ is the witness as described in [6]: the witness consists of a *ranking function* of the actions and a *ranking function* of the state space. The witness is polynomial and can be verified in polynomial time in size of the game graph.

- *Locally ε-optimal witness.* The locally ε-optimal witness consists of the following:

1. The values of the game at every state s, within ε precision.
2. The locally optimal selector $\overline{\sigma} \in \Sigma^{\ell}$. Note that the selector $\overline{\sigma}$ may specify probabilities that are irrational. The locally optimal selector $\overline{\sigma}$ is ε-approximated by a k-uniform selector $\overline{\sigma}_k$, where a k-uniform selector is a selector such that the associated probabilities of the distribution are multiple of $\frac{1}{k}$. It follows from [4, 14], that k is polynomial in the size of the game graph and $\frac{1}{\varepsilon}$. The strategy $\overline{\sigma}_k$ must satisfy the constraint that $\mathrm{Supp}(\overline{\sigma}_k)$ is exactly the set of actions with the least rank as described by the limit-sure witness. The verification of the witness can be achieved in polynomial time, since checking local optimality involves verifying that $\overline{\sigma}_k$ is optimal for the "one-step" game with respect to the values at every state.

It follows from above that there are polynomial witness for perennial ε-optimal strategies and the witness can be verified in polynomial time. This shows that the values of concurrent parity games can be decided with in ε-precision in NP. Since concurrent parity games are closed under complementation the decision procedure is also in coNP. This gives us the following Theorem.

Theorem 4 (Computational complexity of concurrent parity games) *For all constant $\varepsilon > 0$,*

1. *for all rational r, whether $\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) \geq r - \varepsilon$ can be decided in NP $\cap$ coNP.*
2. *the value functions $\langle\langle 1 \rangle\rangle_{val}(\Omega_e)$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega_o)$ can be approximated within ε-precision in time exponential in G and polynomial in $\frac{1}{\varepsilon}$.*

The previous best known algorithm to approximate values is triple exponential in the size of the game graph and logarithmic in $\frac{1}{\varepsilon}$ [8].

References

[1] S. Basu. New results on quantifier elimination over real closed fields and applications to constraint d atabases. *Journal of the ACM*, 46(4):537–555, 1999.

[2] J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the AMS*, 138:295–311, 1969.

[3] K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA 04: ACM-SIAM Symposium on Discrete Algorithms*, pages 114–123, 2004. Technical Report: UCB/CSD-3-1280 (October 2003).

[4] K. Chatterjee, R.Majumdar, and M. Jurdziński. On nash equilibria in stochastic games. In *CSL 04: Computer Science Logic*, pages 26–40, 2004.

[5] A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

[6] L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *Proceedings of the 15th Annual Symposium on Logic in Computer Science*, pages 141–154. IEEE Computer Society Press, 2000.

[7] L. de Alfaro, T.A. Henzinger, and F.Y.C. Mang. Detecting errors before reaching them. In *CAV 00*, volume 1855 of *LNCS*, pages 186–201. Springer, 2000.

[8] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *STOC 01: Symposium on Theory of Computing*, pages 675–683. ACM Press, 2001. Full version to appear in JCSS.

[9] E.A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *Proceedings of the 32th Annual Symposium on Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.

[10] H. Everett. Recursive games. In *Contributions to the Theory of Games III*, volume 39 of *Annals of Mathematical Studies*, pages 47–78, 1957.

[11] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer-Verlag, 1997.

[12] A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.

[13] O. Kupferman and M.Y. Vardi. Module checking. In *Proc. 8th Conference on Computer Aided Verification*, volume 1102 of *LNCS*, pages 75–86. Springer, 1996.

[14] R.J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *EC 03: Electronic Commerce*, pages 36–41. ACM Press, 2003.

[15] D.A. Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998.

[16] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.

[17] J.F. Mertens and A. Neyman. Stochastic games. *International Journal of Game Theory*, 10:53–56, 1981.

[18] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In *5th Symposium on Computation Theory*, LNCS 208, pages 157–168. Springer-Verlag, 1984.

[19] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings of the 16th Annual Symposium on Principles of Programming Languages*, pages 179–190. ACM Press, 1989.

[20] T.E.S. Raghavan and J.A. Filar. Algorithms for stochastic games — a survey. *ZOR — Methods and Models of Operations Research*, 35:437–472, 1991.

[21] P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25(1):206–230, 1987.

[22] L.S. Shapley. Stochastic games. *Proc. Nat. Acad. Sci. USA*, 39:1095–1100, 1953.

[23] A. Tarski. *A Decision Method for Elementary Algebra and Geometry.* University of California Press, Berkeley and Los Angeles, 1951.

[24] W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–191. Elsevier, 1990.

[25] W. Thomas. On the synthesis of strategies in infinite games. In *STACS 95: Theoretical Aspects of Computer Science*, volume 900 of *Lecture Notes in Computer Science*, pages 1–13. Springer-Verlag, 1995.

[26] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.

[27] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *Proceedings of the 26th Annual Symposium on Foundations of Computer Science*, pages 327–338. IEEE Computer Society Press, 1985.