The Complexity of Quantitative Concurrent Parity Games *

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

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 decided 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 case of turn-based stochastic reachability games, for 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 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 games with parity conditions, by showing that memoryless strategies suffice for ε-optimality for coBüchi conditions.

*This research was supported in part by the AFOSR MURI grant F49620-00-1-0327 and the NSF ITR grant CCR-0225610.

†EECS, UC Berkeley, email: c_krish@eecs.berkeley.edu.

‡CE, UC Santa Cruz, email: luca@soe.ucsc.edu.

§EECS, UC Berkeley, and EPFL, Switzerland, email: tah@eecs.berkeley.edu.

1 Introduction

We consider *infinite recursive games* played between two players over a graph [23, 10, 17]. The games proceed in an infinite number of rounds. In 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 a *play*, consists of 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. Concurrent games provide a natural framework to model reactive systems with synchronous interactions [1].

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 [25]; the properties expressible by ω-regular languages include safety, reachability, and fairness. Games with ω-regular winning conditions have been applied to system synthesis [3, 22, 20] and verification [9, 1]. 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 [19, 26]. Hence concurrent games with parity conditions provide an adequate model for the synthesis of synchronous reactive systems.

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 parity conditions can be computed in NP $\cap$ coNP. This result was obtained for turn-based *deterministic* parity games, in which each move determines uniquely (instead of probabilistically) the successor state, in [9], and for turn-based stochastic *reachability* games in [6]; the case of turn-based stochastic parity games was shown in [4].

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 [4, 16]; 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 results 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, in 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 parity 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 parity 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 fields, and then using decision procedures for the theory of reals with addition and multiplication.[24, 2]. Even for the simpler Büchi winning conditions the previously known complexity was EXPTIME [8].

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 a rational r, for all rationals $\varepsilon > 0$, whether the value of the game is in the interval $[r-\varepsilon, r+\varepsilon]$ can be decided 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.[1]

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* the set of states where the game has the same value for player 1. By the results of [7] 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 an 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 symmetric 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 ε-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 [18] for concurrent games with limit-average objectives.

2 Definitions

Notation. For a countable set A, a *probability distribution* on A is a function $\delta : A \to [0,1]$ 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 2.1. (Concurrent Game Structures) *A (two-player)* concurrent game structure

[1] For turn-based deterministic parity games a bound of UP $\cap$ coUP is also known [12], but for turn-based stochastic reachability and turn-based stochastic parity games NP $\cap$ coNP is the best known bound.

$G = \langle S, M, \Gamma_1, \Gamma_2, \delta \rangle$ *consists of the following components:*

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

Plays. 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 that start from the state s.

Randomized strategies. A *selector* ξ for player $i \in \{1, 2\}$ is a function $\xi : S \to \mathcal{D}(M)$ such that for all $s \in S$ and $a \in M$, 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 exists $a \in M$ 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 $\sigma : S^+ \to \Lambda_1$ 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 strategy σ for player i is *pure* if it yields only pure selectors, that is, if it is of type $S^+ \to \Lambda_i^P$. 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 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. 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.

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 [21, 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 [13]. 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 $2^1/_2$ 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 $\mathrm{Inf}(\omega) = \{ s \in S \mid s_k = s \text{ for infinitely many } k \geq 0 \}$ 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 $\mathrm{Reach}(T) = \{ \omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0 \}$. Given a set $F \subseteq S$, the safety objective requires that only states of F be visited. Thus, the set of winning plays is $\mathrm{Safe}(F) = \{ \omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in F \text{ for all } k \geq 0 \}$.
- *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 objectives.* For $c, d \in \mathbb{N}$, we let $[c..d] = \{ c, c+1, \ldots, d \}$. Let $p : S \to [0..d]$ 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 $\text{Parity}(p) = \{ \omega \in \Omega \mid \min\big(p(\text{Inf}(\omega))\big) \text{ is even } \}$, and the *Odd parity objective* as $\text{coParity}(p) = \{ \omega \in \Omega \mid \min\big(p(\text{Inf}(\omega))\big) \text{ is odd } \}$. Informally we say that a path ω satisfies 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. Hence Büchi and coBüchi objectives are simpler and special cases of parity objectives.

Given any parity 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 $\text{Parity}(p)$ starting from state $s \in S$, given the strategies σ, π for the players is $\Pr_s^{\sigma,\pi}(\Omega_{es})$. Given a state $s \in S$ and a parity objective, $\text{Parity}(p)$, we are interested in finding the maximal probability with which player 1 can ensure that $\text{Parity}(p)$ and player 2 can ensure that $\text{coParity}(p)$ 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 $\langle\langle 1 \rangle\rangle_{val}(\Omega_e) : S \to [0,1]$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega_o) : S \to [0,1]$, defined for all $s \in S$ by $\langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \Pr_s^{\sigma,\pi}(\Omega_{es})$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega_o)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \Pr_s^{\sigma,\pi}(\Omega_{os})$. Note that the objectives of the players are complementary and hence we have a zero-sum game. Concurrent games satisfy a *quantitative* version of determinacy [15], stating that for all parity 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 $\inf_{\pi \in \Pi} \Pr_s^{\sigma,\pi}(\Omega_{es}) = \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s)$. For $\varepsilon > 0$, a strategy σ for player 1 is *ε-optimal* if for all $s \in S$ we have $\inf_{\pi \in \Pi} \Pr_s^{\sigma,\pi}(\Omega_{es}) \geq \langle\langle 1 \rangle\rangle_{val}(\Omega_e)(s) - \varepsilon$. 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 $s \in S$.

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 2.2. (FEASIBLE PLAYS AND OUTCOMES) *Given two 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 structure G if for every $k \in \mathbb{N}$ the following conditions hold for some $a_1 \in \Gamma_1(s_k)$ and $a_2 \in \Gamma_2(s_k)$: (1) $s_{k+1} \in \text{Dest}(s_k, a_1, a_2)$; (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 $\text{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ the set of feasible plays that start from s, given the strategies σ and π.* ∎

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

DEFINITION 2.3. (INFINITE TREES, S-LABELED TREES, AND TREES FOR EVENTS) *An* infinite tree *is a set $\text{Tr} \subseteq \mathbb{N}^*$ such that (a) if $x \cdot i \in \text{Tr}$, where $x \in \mathbb{N}^*$ and $i \in \mathbb{N}$, then $x \in \text{Tr}$; (b) for all $x \in \text{Tr}$ there exists $i \in \mathbb{N}$ such that $x \cdot i \in \text{Tr}$. 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 $\tau \subseteq \text{Tr}$ such that (a) $\epsilon \in \tau$; (b) 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 $x \in \text{Tr}$, we denote by $\text{Tr}(x)$ the sub-tree rooted at node x. Formally, $\text{Tr}(x)$ denotes the set $\{ x' \in \text{Tr} \mid x \text{ is a prefix of } x' \}$.*

A S-labeled tree $\mathcal{T}$ is a pair $(\text{Tr}, \langle\cdot\rangle)$, where Tr is a tree and $\langle\cdot\rangle : \text{Tr} \to S$ maps each node of Tr to a state $s \in S$. Given a S-labelled tree $\mathcal{T}$, and a infinite path $\tau \subseteq \text{Tr}$, 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 $\mathcal{T}_s = (\text{Tr}_s, \langle\cdot\rangle)$, where $\langle\epsilon\rangle = s$, represents a set of infinite paths, denoted as $\mathcal{L}(\mathcal{T}_s) \subseteq \Omega_s$, such that $\mathcal{L}(\mathcal{T}_s) = \{ \omega = \langle s_0 = s, s_1, s_2, \ldots \rangle \in \Omega_s \mid \exists \tau \subseteq \text{Tr}_s.\ \langle\tau\rangle = \omega \}$. A S-labeled tree $\mathcal{T}_s$ represents an event $\mathcal{A} \subseteq \Omega_s$ if and only if $\mathcal{L}(\mathcal{T}_s) = \mathcal{A}$. ∎

Trees for outcomes and events. Let $\mathcal{T} = (\text{Tr}, \langle\cdot\rangle)$ be a S-labeled tree and consider $x \in \text{Tr}$ such that $|x| = n$. We denote by x_i the prefix of x of length i. We denote by $\text{hist}(x) = (\langle\epsilon\rangle, \langle x_1\rangle, \ldots, \langle x_n\rangle)$ the history represented by the path from the root to the node x. Given strategies σ and π, and a state s, a S-labelled tree $\mathcal{T}_s^{\sigma,\pi} = (\mathrm{Tr}_s^{\sigma,\pi}, \langle\cdot\rangle)$ to represent $\mathrm{Outcome}(s,\sigma,\pi)$ is defined as follows: (a) $\langle\epsilon\rangle = s$; (b) for $x \in \mathrm{Tr}_s^{\sigma,\pi}$, let $|x| = n$, and consider the set $U = \bigcup_{\{\ \sigma(\mathrm{hist}(x))(a_1)>0, \pi(\mathrm{hist}(x))(a_2)>0\ \}} \mathrm{Dest}(\langle x_n\rangle, a_1, a_2)$. The set of successors for x in the tree is $x \cdot j$ for $j \in \{\ 1,2,\ldots,|U|\ \}$, and the labeling function $\langle\cdot\rangle$ is a bijection from the successors of x to the set U of states. For an event $\mathcal{A}$, the stochastic tree, $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi} = (\mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}, \langle\cdot\rangle)$ is constructed from $\mathcal{T}_s^{\sigma,\pi}$ by retaining the set of paths $\mathcal{A} \cap \mathrm{Outcome}(s,\sigma,\pi)$.[2] We denote by $\mathrm{Cone}(x) = \{\ \omega = \langle s_0, s_1, s_2, \ldots\rangle \mid \langle x_i\rangle = s_i \text{ for all } 0 \le i \le n\ \}$ the set of paths with the prefix $\mathrm{hist}(x)$. Given a measurable event $\mathcal{A} \subseteq \Omega_s$ along with strategies σ and π such that $\mathrm{Pr}_s^{\sigma,\pi}(\mathcal{A}) > 0$, consider the S-labeled tree $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi}$ to represent $\mathcal{A} \cap \mathrm{Outcome}(s,\sigma,\pi)$. Consider the event $\mathcal{A}_{nil} = \{\mathrm{Cone}(x) \mid x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}.\ \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Cone}(x) \cap \mathcal{A}) = 0\}$. Since $\mathcal{A}_{nil}$ is the countable union of measurable sets each with measure 0 we have $\mathrm{Pr}_s^{\sigma,\pi}(\mathcal{A}_{nil} \cap \mathcal{A}) = 0$. Hence in the sequel, without loss of generality, given any event $\mathcal{A}$ we only consider the event $\mathcal{A} \setminus \mathcal{A}_{nil}$, and with a little abuse of notation we use $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi}$ to represent the stochastic tree $\mathcal{T}_{(\mathcal{A}\setminus\mathcal{A}_{nil}),s}^{\sigma,\pi}$. Furthermore, again without loss of generality, we assume that for any $x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}$ we have $\mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Cone}(x) \cap \mathcal{A}) > 0$. Henceforth, for any $x \in \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}$ we write $\mathrm{Pr}_x^{\sigma,\pi}(\mathcal{B} \mid \mathcal{A})$ to denote $\mathrm{Pr}_s^{\sigma,\pi}(\mathcal{B} \mid \mathrm{Cone}(x), \mathcal{A})$.

DEFINITION 2.4. (PERENNIAL ε-OPTIMAL STRATEGIES) *Given $\varepsilon > 0$, a strategy σ is a* perennial ε-optimal strategy *for player 1, from state s, if for all strategies π and for all nodes x in the stochastic tree* $\mathrm{Tr}_s^{\sigma,\pi}$, *we have* $\mathrm{Pr}_x^{\sigma,\pi}(\Omega_{es}) \ge \langle\!\langle 1\rangle\!\rangle_{val}(\Omega_e)(\langle x\rangle) - \varepsilon$, *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. The perennial ε-optimal strategies for player 2 are defined analogously. We denote by Σ_ε and Π_ε the sets 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 leads to the following result.

PROPOSITION 2.1. *For all $\varepsilon > 0$, we have $\Sigma_\varepsilon \neq \emptyset$ and $\Pi_\varepsilon \neq \emptyset$.*

3 Results

In this section we construct polynomial witnesses for perennial ε-optimal strategies and describe a polynomial procedure to verify the witnesses. As an immediate consequence, 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 reveal several key characteristics of perennial ε-optimal strategies. In general, perennial ε-optimal strategies require infinite memory [7, 8]. We show that even though the perennial ε-optimal strategies require infinite memory in general, there exist perennial ε-optimal strategies that in the limit for $\varepsilon \to 0$ converge to memoryless strategies. This result parallels with the celebrated result of Mertens-Neyman [18] for concurrent games with limit-average objectives, which states that there exist ε-optimal strategies that in the limit coincide with memoryless strategies (the memoryless strategy correspond to the memoryless optimal strategies in the discounted game with discount factor very close to 0). However, the memoryless strategies to which the ε-optimal strategies converge are not necessarily ε-optimal themselves.

In concurrent games with safety objectives, optimal memoryless strategies always 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 [7]. In contrast, we show that for all $\varepsilon > 0$, memoryless ε-optimal strategies exist 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 for concurrent parity games with with 3 or more priorities, ε-optimal strategies require in general infinite memory. Our results thus complete the characterization of the memory requirements of ε-optimal strategies in concurrent parity games.

Reachability properties. Several key properties of perennial ε-optimal strategies will follow by analyzing the behavior of the strategies with respect to some reachability and safety objectives. In the sequel, we consider stochastic trees $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi}$ such that $\mathrm{Pr}_s^{\sigma,\pi}(\mathcal{A}) > 0$. Given a stochastic tree $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi}$, let κ be a subset of nodes, i.e., $\kappa \subseteq \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}$. 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 $\kappa \subseteq \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi}$, let $\mathrm{ReachTree}(\kappa) = \{\ \langle\tau\rangle \mid \tau \text{ is an infinite path in } \mathrm{Tr}_{\mathcal{A},s}^{\sigma,\pi} \text{ such that exists } i \in \mathbb{N}.\ \tau_i \in \kappa\ \}$, denote the set of paths that reach the subset κ of nodes.

[2] Note that the stochastic tree $\mathcal{T}_{\mathcal{A},s}^{\sigma,\pi}$ is not constructed by extending every finite prefix of paths.

2. *Safety in tree.* For a set $\kappa \subseteq \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, let $\mathrm{SafeTree}(\kappa) = \{ \langle\tau\rangle \mid \tau$ is an infinite path in $\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$ such that for all $i \in \mathbb{N}.\ \tau_i \in \kappa \}$, 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 3.1. (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 $\inf_{x\in\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}} \mathrm{Pr}^{\sigma,\pi}_x(\mathit{ReachTree}(\kappa) \mid \mathcal{A}) > 0$, then $\mathrm{Pr}^{\sigma,\pi}_x(\mathit{ReachTree}(\kappa) \mid \mathcal{A}) = 1$, for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$.*

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

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. ∎

Notation. 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 $\mathrm{InfSet}(B) = \{\omega \mid \mathrm{Inf}(\omega) \subseteq B\}$ and $\mathrm{InfSetEq}(B) = \{\omega \mid \mathrm{Inf}(\omega) = B\}$. 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 $\mathcal{C}^{\sigma,\pi}_{\mathcal{A},\varepsilon}(x) = \{ B \subseteq S \mid \mathrm{Pr}^{\sigma,\pi}_x(\mathrm{InfSet}(B) \mid \mathcal{A}) \geq 1-\varepsilon \}$. Note that for $\varepsilon_1 > 0$ and $\varepsilon_2 > 0$ such that $\varepsilon_1 \leq \varepsilon_2$, for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, 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 every $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$.

LEMMA 3.2. *For all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, there is a unique minimal element of $\mathcal{C}^{\sigma,\pi}_{\mathcal{A}}(x)$ under $\subset$ ordering.*

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

LEMMA 3.3. *Given a S-labeled tree $\mathcal{T}^{\sigma,\pi}_{\mathcal{A},s}$, for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}$, 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 $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSetEq}(B) \mid \mathcal{A}) \geq 1-\varepsilon$.*

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 exists 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 3.1 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 $\inf_{x_1\in\mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}(\{s\}) \mid \mathcal{A}) = 0$. For every $\varepsilon > 0$, there must be a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{InfSet}(B \setminus \{ s \}) \mid \mathcal{A}) \geq 1-\varepsilon$. Thus, 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 3.4. *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: (a)* *for all $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\Omega_{es} \mid \mathcal{A}) \geq 1-\varepsilon$; or (b) for all $\varepsilon > 0$, there is a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\Omega_{os} \mid \mathcal{A}) \geq 1-\varepsilon$.*

Lemma 3.4 is an easy consequence of Lemma 3.3. In the sequel, we denote by $W_1 = \{ s \mid \langle\!\langle 1\rangle\!\rangle_{val}(\Omega_e)(s) = 1 \}$ and $W_2 = \{ s \mid \langle\!\langle 2\rangle\!\rangle_{val}(\Omega_o)(s) = 1 \}$ the set of states where player 1 and player 2 can achieve value 1, respectively. We will prove that if both players play one of their perennial ε-optimal strategies, with $\varepsilon \to 0$, then the play reaches $W_1 \cup W_2$ with probability 1. 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 3.5. (REACHABILITY WITH ε-OPTIMAL STRATEGIES) *Given a game structure G, consider a strategy pair $(\sigma,\pi) \in \Sigma_\varepsilon \times \Pi_\varepsilon$, for sufficiently small ε. For all states s and for all nodes $x \in \mathrm{Tr}^{\sigma,\pi}_s$ we have $\mathrm{Pr}^{\sigma,\pi}_x(Safe_s(\overline{W_1 \cup W_2})) = 0$.*

Proof. Fix $\eta > 0$, such that $0 < 2\cdot\eta < \alpha = \min\{\langle\!\langle 1\rangle\!\rangle_{val}(\Omega_e)(s), \langle\!\langle 2\rangle\!\rangle_{val}(\Omega_o)(s) \mid s \in \overline{W_1 \cup W_2}\}$, 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 $U^{\sigma,\pi}_s = \{x \in \mathrm{Tr}^{\sigma,\pi}_s \mid s \in \overline{W_1 \cup W_2}$ and $\mathrm{Pr}^{\sigma,\pi}_x(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) > 0\}$. If $U^{\sigma,\pi}_s$ is empty the desired result follows.

Assume for the sake of contradiction that $U^{\sigma,\pi}_s$ is non-empty. Let x be a node in $U^{\sigma,\pi}_s$ and consider the S-labeled subtree $\mathcal{T}^{\sigma,\pi}_s(x)$ rooted at x. Since $\mathrm{Pr}^{\sigma,\pi}_x(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) > 0$, we must have $\inf_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_s(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}_s(W_1 \cup W_2)) = 0$, or $\sup_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_s(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) = 1$. In fact, from Lemma 3.1 we have that $\inf_{x_1 \in \mathrm{Tr}^{\sigma,\pi}_s(x)} \mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Reach}_s(W_1 \cup W_2)) > 0$ implies $\mathrm{Pr}^{\sigma,\pi}_x(\mathrm{Reach}_s(W_1 \cup W_2)) = 1$.

Consider a node $x_1 \in \mathrm{Tr}^{\sigma,\pi}_s(x)$ such that $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathrm{Safe}_s(\overline{W_1 \cup W_2})) \geq 1-\eta$. Let $\mathcal{A}$ be the event $\mathrm{Safe}_s(\overline{W_1 \cup W_2})$. Since σ and π are perennial η-optimal strategies, and $\mathrm{Pr}^{\sigma,\pi}_{x_1}(\mathcal{A}) \geq 1-\eta$, it follows that for every node $x_2 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x_1)$ we have $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{es} \mid \mathcal{A}) \geq c_1 \geq (\alpha - 2\eta) > 0$ and $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{os} \mid \mathcal{A}) \geq c_2 \geq (\alpha - 2\eta) > 0$. This implies that for all nodes $x_2 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x_1)$ we have $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{es} \mid \mathcal{A}) \leq 1-c_2$ and $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{os} \mid \mathcal{A}) \leq 1-c_1$. It follows from Lemma 3.4 that for every $\varepsilon > 0$, there is a node $x_2 \in \mathrm{Tr}^{\sigma,\pi}_{\mathcal{A},s}(x_1)$ such that either $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{es} \mid \mathcal{A}) \geq 1-\varepsilon$ or $\mathrm{Pr}^{\sigma,\pi}_{x_2}(\Omega_{os} \mid \mathcal{A}) \geq 1-\varepsilon$. Since c_1 and c_2 are constants greater than 0, we have a contradiction. Hence $U^{\sigma,\pi}_s = \emptyset$ and the result follows. ∎

Reduction to qualitative witness. The notion of *local optimality* plays 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 of 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, an 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. Our polynomial witnesses will consist in strategies that are locally ε-optimal (which can be checked in polynomial time), and that have a particular structure that ensures their global ε-optimality.

DEFINITION 3.1. (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 of 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.* ∎

Observe that a strategy that at each round i chooses a locally optimal selector with probability at least $(1-\varepsilon_i)$, with $\sum_{i=0}^{\infty} \varepsilon_i \leq \varepsilon$, is a locally ε-optimal strategy. A *value class* of the game is the set of all states where the game has a given value. A value class $\mathrm{VC}(r)$ is the set of states s such that the value for player 1 is r. Formally, $\mathrm{VC}(r) = \{s \mid \langle\!\langle 1\rangle\!\rangle_{val}(\Omega_e)(s) = r\}$. 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 [7], and has a 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. Let $G = (S, M, \Gamma_1, \Gamma_2, \delta)$ be a concurrent game with parity objectives $\mathrm{Parity}(p)$ and $\mathrm{coParity}(p)$ for player 1 and player 2 respectively, and let the priority function be p. For a state $s \in S$, we define the set of allowable supports $\mathrm{OptSupps}(s) = \{ \gamma \subseteq \Gamma_1(s) \mid \exists \xi_1^\ell \in \Lambda^\ell . \mathrm{Supp}(\xi_1^\ell) = \gamma \}$ to be the set of supports of locally optimal selectors. For every $s \in S$, we assume that we have a fixed way to enumerate $\mathrm{OptSupps}(s) = \{ \gamma_1, \gamma_2, \ldots, \gamma_n \}$. Consider a value class $\mathrm{VC}(r)$ with $0 < r < 1$. We construct a concurrent game $\widetilde{G}_r = (\widetilde{S}_r, \widetilde{M}, \widetilde{\Gamma_1}, \widetilde{\Gamma_2}, \widetilde{\delta})$ with a priority function $\widetilde{p}$ as follows:

1. *State space.* $\widetilde{S}_r = \{ \widetilde{s} \mid s \in \mathrm{VC}(r) \} \cup \{ w_1, w_2 \} \cup \{ (\widetilde{s}, i) \mid s \in \mathrm{VC}(r), i \in \{ 1, 2, \ldots, |\mathrm{OptSupps}(s)| \} \}$.
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 $(\widetilde{s}, i) \in \widetilde{S}_r$; and (c) $\widetilde{p}(w_1) = 0$ and $\widetilde{p}(w_2) = 1$.
3. *Moves assignment.*
   (a) $\widetilde{\Gamma_1}(\widetilde{s}) = \{1, 2, \ldots, |\mathrm{OptSupps}(s)|\}$ 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) $\widetilde{\Gamma_1}((\widetilde{s}, i)) = \{ i \} \cup (\Gamma_1(s) \setminus \gamma_i)$, where $\mathrm{OptSupps}(s) = \{ \gamma_1, \gamma_2, \ldots, \gamma_n \}$, 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 all the moves not in γ_i are still available.
4. *Transition function.*
   (a) The states w_1 and w_2 are absorbing states. Observe that player 1 has value 1 at state w_1, and value 0 at state w_2.
   (b) For any state $\widetilde{s}$ we have $\widetilde{\delta}(\widetilde{s}, i, a_2)((\widetilde{s}, i)) = 1$: at state $\widetilde{s}$, player 1 can decide which element of $\mathrm{OptSupps}(s)$ to play, and if player 1 chooses move i the game proceed to state $(\widetilde{s}, i)$.
   (c) *Transition function at state* $(\widetilde{s}, i)$. Let $\mathrm{OptSupps}(s) = \{ \gamma_1, \gamma_2, \ldots, \gamma_n \}$.
      i. For all moves $a_2 \in \Gamma_2(s)$, if there is $a_1 \in \gamma_i$ 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, when a pair of moves a_1, a_2 with $a_1 \in \gamma_i$ is played, if the game G proceeds with positive probability to a different value class, then the game $\widetilde{G}_r$ proceeds to the state w_1, which has value 1 for player 1. Note that since $a_1 \in \gamma_i$ and $\gamma_i \in \mathrm{OptSupps}(s)$, if the game G proceeds to a different value class with positive probability, it proceeds to $\bigcup_{k>r} \mathrm{VC}(k)$ with positive probability.
      ii. For all moves $a_2 \in \Gamma_2(s)$, if for all $a_1 \in \gamma_i$ we have $\sum_{s' \in \mathrm{VC}(r)} \delta(s, a_1, a_2)(s') = 1$, 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 a locally optimal selector with $\mathrm{Supp}(\xi_1^\ell) = \gamma_i$.
      iii. For all moves $a_1 \in (\Gamma_1(s) \setminus \gamma_i)$ and $a_2 \in \Gamma_2(s)$ we let $\widetilde{\delta}((\widetilde{s}, i), a_1, a_2)(\widetilde{s'}) = \delta(s, a_1, a_2)(s')$ for $s' \in \mathrm{VC}(r)$; furthermore, we let $\widetilde{\delta}((\widetilde{s}, i), a_1, a_2)(w_2) = \sum_{s' \notin \mathrm{VC}(r)} \delta(s, a_1, a_2)(s')$.

LEMMA 3.6. *For all $0 < r < 1$ and all $s \in \mathrm{VC}(r)$, the state $\widetilde{s}$ is limit-sure winning for player 1 in the game $\widetilde{G}_r$, i.e., from state $\widetilde{s}$ player 1 can win with probability arbitrarily close to 1.*

Limit-sure witness. The witness strategy for a limit-sure game constructed in [7] consists of two parts: a *ranking function* of the states, and a *ranking function* of the actions at a state. These ranking functions were described by a μ-calculus formula. At the round k of a play, the witness strategy σ plays at a state s the actions with least rank with positive-bounded probabilities, and the other actions with vanishingly small probabilities as $\varepsilon \to 0$. 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 a memoryless strategy such that, at each state s, $\mathrm{Supp}(\sigma_\ell(s))$ is the set of actions with least rank at s. We denote by *limit-sure witness move set* the set of actions with the least rank, i.e., at each s the set of moves $\mathrm{Supp}(\sigma_\ell(s))$. It follows from the above construction that as $\varepsilon \to 0$, the limit-sure winning strategy σ converges to the memoryless selector σ_ℓ.

LEMMA 3.7. *In the game $\widetilde{G}_r$, there is a limit-sure winning strategy with support $i \in \{ 1, 2, \ldots, |\mathrm{OptSupps}(s)| \}$ at $\widetilde{s}$, and with limit-sure witness move set γ_i at $(\widetilde{s}, i)$.*

DEFINITION 3.2. (VALUE-CLASS QUALITATIVE ε-OPTIMAL STRATEGIES) *For $\varepsilon > 0$, a strategy σ_ε is a value-class qualitative ε-optimal strategy for a value-class $\mathrm{VC}(r)$, with $0 < r < 1$, if (a) σ_ε is locally ε-optimal, and (b) for all nodes x in*

$\mathrm{Tr}_s^{\sigma_\varepsilon,\pi}$ *with* $\langle x\rangle \in \mathrm{VC}(r)$ *and all* $\pi \in \Pi$ *we have* $\mathrm{Pr}_x^{\sigma_\varepsilon,\pi}(\Omega_{es} \mid \mathit{Safe}(\mathrm{VC}(r))) \geq 1-\varepsilon$. *A strategy* σ_ε *is value-class qualitative* ε*-optimal if it is value-class qualitative* ε*-optimal for all value classes* $\mathrm{VC}(r)$*, for all* $0<r<1$. ∎

Lemma 3.8 states that the value-class qualitative ε-optimal strategies for different value classes can be "stitched" or composed together to produce a perennial ε-optimal strategy. This allows us to produce witness for individual value classes and compose them to obtain a witness for perennial ε-optimal strategy. The key argument is as follows: given a value-class qualitative ε-optimal strategy for any strategy π for player 2 if the game stays in a value class then player 1 wins with probability at least $1-\varepsilon$; otherwise, the game leaves the value class according to the locally ε-optimal strategy, and reaches W_1 with probability at least the value of the game, within ε-precision.

LEMMA 3.8. (STITCHING LEMMA) *Let* σ_ε *be a value-class qualitative* ε*-optimal strategy that is also perennial* ε*-optimal for all states in* W_1. *Then* σ_ε *is a perennial* ε*-optimal strategy.*

Theorem 3.1 follows from existence of memoryless limit-sure winning strategies for concurrent games with coBüchi objectives [7] and the existence of perennial ε-optimal strategies obtained by composing value-class qualitative ε-optimal strategies across value classes (Lemma 3.8).

THEOREM 3.1. (MEMORYLESS ε-OPTIMAL STRATEGIES FOR COBÜCHI OBJECTIVES) *For every real* $\varepsilon > 0$*, memoryless* ε*-optimal strategies exist for all coBüchi objectives on all concurrent game structures.*

Theorem 3.2 states that there exist perennial ε-optimal strategies that in the limit coincide with a locally optimal selector, i.e., a memoryless strategy with locally optimal selectors. The result follows from Lemma 3.8, which proves the existence of perennial ε-optimal strategies as value-class qualitative ε-optimal strategies, and from the properties of limit-sure winning strategies.

THEOREM 3.2. (LIMIT OF ε-OPTIMAL STRATEGIES) *For all concurrent game structures with parity objectives, for every real* $\varepsilon > 0$*, there exists a perennial* ε*-optimal strategy* $\sigma_\varepsilon \in \Sigma_\varepsilon$ *such that the sequence of the strategies* σ_ε *converges to a locally optimal selector* $\overline{\sigma}$ *as* $\varepsilon \to 0$*, i.e.,* $\lim_{\varepsilon\to 0}\sigma_\varepsilon = \overline{\sigma}$*, for* $\overline{\sigma} \in \Sigma^\ell$.

Witness for perennial ε-optimal strategies. The witness for a perennial ε-optimal strategy σ_ε is presented as a value-class qualitative ε-optimal strategy (from Lemma 3.8). The existence of a value-class qualitative ε-optimal strategy follows from Lemma 3.6 and Lemma 3.7. 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 witness described in [7]. Observe that the game $\widetilde{G}_r$ can be exponential in the size of the game G, since the set $\mathrm{OptSupps}(s)$ 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 3.7) 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 [7]: 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 values of the game at all state s, within ε-precision and the locally optimal selector $\overline{\sigma} \in \Sigma^\ell$. 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 $\frac{1}{j}, j \leq k$. It follows from [5, 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. 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].

THEOREM 3.3. (COMPLEXITY OF CONCURRENT PARITY GAMES) *For all concurrent game structures G, for all parity objectives Ω_e and Ω_o, and for all rationals $\varepsilon > 0$,*

1. *for all rationals r, whether $\langle\langle 1\rangle\rangle_{val}(\Omega_e)(s) \in [r-\varepsilon, r+\varepsilon]$ 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 with precision ε-precision in time exponential in $|G|$ and polynomial in $\frac{1}{\varepsilon}$.*

References

[1] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.

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

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

[4] K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA 04*, pages 114–123. ACM-SIAM, 2004.

[5] K. Chatterjee, R.Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In *CSL 04*, volume LNCS 3210, pages 26–40. Springer, 2004.

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

[7] L. de Alfaro and T.A. Henzinger. Concurrent omega-regular games. In *LICS 00*, pages 141–154. IEEE Computer Society Press, 2000.

[8] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of Computer and System Sciences*, 68:374–397, 2004.

[9] E.A. Emerson and C. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS 91*, 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, 1997.

[12] M. Jurdzinski. Deciding the winner in parity games is in UP $\cap$ co-UP. *Information Processing Letters*, 68:119–124, 1998.

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

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

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

[16] A. McIver and C. Morgan. Games, probability, and the quantitative μ-calculus $q\mu$. In *LPAR 02*, volume LNCS 2514, pages 292–310. Springer, 2002.

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

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

[19] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In *Symposium on Computation Theory 84*, volume LNCS 208, pages 157–168. Springer, 1984.

[20] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL 89*, pages 179–190. ACM Press, 1989.

[21] 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.

[22] 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.

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

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

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

[26] W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, pages 389–455. Springer, 1997.

[27] M.Y. Vardi. Automatic verification of probabilistic concurrent finite-state systems. In *FOCS 85*, pages 327–338. IEEE Computer Society Press, 1985.