

Optimal Strategy Synthesis in Stochastic Muller Games

Krishnendu Chatterjee



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/Eecs-2006-122

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/Eecs-2006-122.html>

October 4, 2006

Copyright © 2006, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Optimal Strategy Synthesis in Stochastic Müller Games ^{*}

Krishnendu Chatterjee

EECS, University of California, Berkeley, USA
c.krish@eecs.berkeley.edu

Abstract. The theory of graph games with ω -regular winning conditions is the foundation for modeling and synthesizing reactive processes. In the case of stochastic reactive processes, the corresponding stochastic graph games have three players, two of them (System and Environment) behaving adversarially, and the third (Uncertainty) behaving probabilistically. We consider two problems for stochastic graph games: the *qualitative* problem asks for the set of states from which a player can win with probability 1 (*almost-sure winning*); the *quantitative* problem asks for the maximal probability of winning (*optimal winning*) from each state. We consider ω -regular winning conditions formalized as Müller winning conditions. We present optimal memory bounds for *pure* (deterministic) almost-sure winning and optimal winning strategies in stochastic graph games with Müller winning conditions. We also present improved memory bounds for randomized almost-sure winning and optimal strategies. Our results are relevant in synthesis of stochastic reactive processes.

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

1 Introduction

A stochastic graph game [6] is played on a directed graph with three kinds of states: player-1, player-2, and probabilistic states. At player-1 states, player 1 chooses a successor state; at player-2 states, player 2 chooses a successor state; and at probabilistic states, a successor state is chosen according to a given probability distribution. The result of playing the game forever is an infinite path through the graph. If there are no probabilistic states, we refer to the game as a *2-player graph game*; otherwise, as a *2^{1/2}-player graph game*. There has been a long history of using 2-player graph games for modeling and synthesizing reactive processes [1, 17, 19]: a reactive system and its environment represent the two players, whose states and transitions are specified by the states and edges of a game graph. Consequently, 2^{1/2}-player graph games provide the theoretical foundation for modeling and synthesizing processes that are both reactive and stochastic [10, 18].

For the modeling and synthesis (or “control”) of reactive processes, one traditionally considers ω -regular winning conditions, which naturally express the temporal specifications and fairness assumptions of transition systems [13]. This paper focuses on 2^{1/2}-player graph games with respect to an important normal form of ω -regular winning conditions; namely *Müller* winning conditions [20].

In the case of 2-player graph games, where no randomization is involved, a fundamental determinacy result of Gurevich and Harrington [11] based on LAR (*latest appearance record*) construction ensures that, given an ω -regular winning condition, at each state, either player 1 has a strategy to ensure that the condition holds, or player 2 has a strategy to ensure that the condition does not hold. Thus, the problem of solving 2-player graph games consists in finding the set of *winning states*, from which player 1 can ensure that the condition holds. Along with the computation of the winning states, the characterization of complexity of winning strategies is a central question, since the winning strategies represent the implementation of the controller in the synthesis problem. The elegant algorithm of Zielonka [21] uses the LAR construction to compute winning sets in 2-player graph games with Müller conditions. In [7] the authors present an insightful analysis of Zielonka’s algorithm to present optimal memory bounds (matching upper and lower bound) for winning strategies in 2-player graph games with Müller conditions.

In the case of 2^{1/2}-player graph games, where randomization is present in the transition structure, the notion of winning needs to be clarified. Player 1 is said to *win surely* if she has a strategy that guarantees to achieve the winning condition against all player-2 strategies. While this is the classical notion of winning in the 2-player case, it is less meaningful in the presence of probabilistic states, because it makes all probabilistic choices adversarial (it treats them analogously to player-2 choices). To adequately treat probabilistic choice, we consider the *probability* with which player 1 can ensure that the winning condition is met. We thus define two solution problems for 2^{1/2}-player graph games: the *qualitative* problem asks for the set of states from which player 1 can ensure winning with probability 1; the *quantitative* problem asks for the maximal probability with which player 1 can ensure winning from each state (this probability is called the *value* of the game at a state). Correspondingly, we define *almost-sure winning strategies*, which enable player 1 to win with probability 1 whenever possible, and *optimal strategies*, which enable player 1 to win with maximal probability. The main result of this paper is an optimal memory bound for pure (deterministic) almost-sure and optimal strategies in 2^{1/2}-player graph games with Müller conditions. In fact we generalize the elegant analysis of [7] to present an upper bound for optimal strategies for 2^{1/2}-player graph games with Müller conditions that matches the lower bound for sure winning in 2-player games. As a consequence we generalize several results known for 2^{1/2}-player graph games: such as existence of pure memoryless optimal strategies for *parity* conditions [5, 22, 15] and Rabin conditions [4]. We present the result for almost-sure strategies in Section 3; and then generalize it to optimal strategies in Section 4. We also study the memory bounds for randomized strategies. In case of randomized strategies we improve the upper bound for almost-sure and optimal strategies as compared to pure strategies (Section 5). The problem of a matching upper and lower bound for almost-sure and optimal randomized strategies remains open.

2 Definitions

We consider several classes of turn-based games, namely, two-player turn-based probabilistic games ($2^{1/2}$ -player games), two-player turn-based deterministic games (2-player games), and Markov decision processes ($1^{1/2}$ -player games).

Notation. For a finite set A , a *probability distribution* on A is a function $\delta: A \rightarrow [0, 1]$ 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 $\text{Supp}(\delta) = \{x \in A \mid \delta(x) > 0\}$ the *support* of δ .

Game graphs. A *turn-based probabilistic game graph* ($2^{1/2}$ -player game graph) $G = ((S, E), (S_1, S_2, S_\circ), \delta)$ consists of a directed graph (S, E) , a partition (S_1, S_2, S_\circ) of the finite set S of states, and a probabilistic transition function $\delta: S_\circ \rightarrow \mathcal{D}(S)$, where $\mathcal{D}(S)$ denotes the set of probability distributions over the state space S . The states in S_1 are the *player-1* states, where player 1 decides the successor state; the states in S_2 are the *player-2* states, where player 2 decides the successor state; and the states in S_\circ are the *probabilistic* states, where the successor state is chosen according to the probabilistic transition function δ . We assume that for $s \in S_\circ$ and $t \in S$, we have $(s, t) \in E$ iff $\delta(s)(t) > 0$, and we often write $\delta(s, t)$ for $\delta(s)(t)$. For technical convenience we assume that every state in the graph (S, E) has at least one outgoing edge. For a state $s \in S$, we write $E(s)$ to denote the set $\{t \in S \mid (s, t) \in E\}$ of possible successors.

A set $U \subseteq S$ of states is called δ -closed if for every probabilistic state $u \in U \cap S_\circ$, if $(u, t) \in E$, then $t \in U$. The set U is called δ -live if for every nonprobabilistic state $s \in U \cap (S_1 \cup S_2)$, there is a state $t \in U$ such that $(s, t) \in E$. A δ -closed and δ -live subset U of S induces a *subgame graph* of G , indicated by $G \upharpoonright U$.

The *turn-based deterministic game graphs* (2-player game graphs) are the special case of the $2^{1/2}$ -player game graphs with $S_\circ = \emptyset$. The *Markov decision processes* ($1^{1/2}$ -player game graphs) are the special case of the $2^{1/2}$ -player game graphs with $S_1 = \emptyset$ or $S_2 = \emptyset$. We refer to the MDPs with $S_2 = \emptyset$ as *player-1 MDPs*, and to the MDPs with $S_1 = \emptyset$ as *player-2 MDPs*.

Plays and strategies. An infinite path, or *play*, of the game graph G is an infinite sequence $\omega = \langle s_0, s_1, s_2, \dots \rangle$ of states such that $(s_k, s_{k+1}) \in E$ for all $k \in \mathbb{N}$. We write Ω for the set of all plays, and for a state $s \in S$, we write $\Omega_s \subseteq \Omega$ for the set of plays that start from the state s .

A *strategy* for player 1 is a function $\sigma: S^* \cdot S_1 \rightarrow \mathcal{D}(S)$ that assigns a probability distribution to all finite sequences $\mathbf{w} \in S^* \cdot S_1$ of states ending in a player-1 state (the sequence represents a prefix of a play). Player 1 follows the strategy σ if in each player-1 move, given that the current history of the game is $\mathbf{w} \in S^* \cdot S_1$, she chooses the next state according to the probability distribution $\sigma(\mathbf{w})$. A strategy must prescribe only available moves, i.e., for all $\mathbf{w} \in S^*$, and $s \in S_1$ we have $\text{Supp}(\sigma(\mathbf{w} \cdot s)) \subseteq E(s)$. The strategies for player 2 are defined analogously. We denote by Σ and Π the set of all strategies for player 1 and player 2, respectively.

Once a starting state $s \in S$ and strategies $\sigma \in \Sigma$ and $\pi \in \Pi$ for the two players are fixed, the outcome of the game is a random walk $\omega_s^{\sigma, \pi}$ for which the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega$ is a measurable set of paths. Given strategies σ for player 1 and π for player 2, a play $\omega = \langle s_0, s_1, s_2, \dots \rangle$ is *feasible* if for every $k \in \mathbb{N}$ the following three conditions hold: (1) if $s_k \in S_\circ$, then $(s_k, s_{k+1}) \in E$; (2) if $s_k \in S_1$, then $\sigma(s_0, s_1, \dots, s_k)(s_{k+1}) > 0$; and (3) if $s_k \in S_2$ then $\pi(s_0, s_1, \dots, s_k)(s_{k+1}) > 0$. Given two strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a state $s \in S$, we denote by $\text{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ the set of feasible plays that start from s given strategies σ and π . For a state $s \in S$ and an event $\mathcal{A} \subseteq \Omega$, we write $\text{Pr}_s^{\sigma, \pi}(\mathcal{A})$ for the probability that a path belongs to \mathcal{A} if the game starts from the state s and the players follow the strategies σ and π , respectively. In the context of player-1 MDPs we often omit the argument π , because Π is a singleton set.

We classify strategies according to their use of randomization and memory. The strategies that do not use randomization are called pure. A player-1 strategy σ is *pure* if for all $\mathbf{w} \in S^*$ and $s \in S_1$, there is a state $t \in S$ such that $\sigma(\mathbf{w} \cdot s)(t) = 1$. We denote by $\Sigma^P \subseteq \Sigma$ the set of pure strategies for player 1. A strategy that is not necessarily pure is called *randomized*. Let \mathbf{M} be a set called

memory, that is, \mathbf{M} is a set of memory elements. A player-1 strategy σ can be described as a pair of functions $\sigma = (\sigma_u, \sigma_m)$: a *memory-update* function $\sigma_u: S \times \mathbf{M} \rightarrow \mathbf{M}$ and a *next-move* function $\sigma_m: S_1 \times \mathbf{M} \rightarrow \mathcal{D}(S)$. We can think of strategies with memory as input/output automaton computing the strategies (see [7] for details). The strategy (σ_u, σ_m) is *finite-memory* if the memory \mathbf{M} is finite, and then we denote the size of the memory of the strategy σ by the size of its memory \mathbf{M} , i.e., $|\mathbf{M}|$. We denote by Σ^F the set of finite-memory strategies for player 1, and by Σ^{PF} the set of *pure finite-memory* strategies; that is, $\Sigma^{PF} = \Sigma^P \cap \Sigma^F$. The strategy (σ_u, σ_m) is *memoryless* if $|\mathbf{M}| = 1$; that is, the next move does not depend on the history of the play but only on the current state. A memoryless player-1 strategy can be represented as a function $\sigma: S_1 \rightarrow \mathcal{D}(S)$. A *pure memoryless strategy* is a pure strategy that is memoryless. A pure memoryless strategy for player 1 can be represented as a function $\sigma: S_1 \rightarrow S$. We denote by Σ^M the set of memoryless strategies for player 1, and by Σ^{PM} the set of pure memoryless strategies; that is, $\Sigma^{PM} = \Sigma^P \cap \Sigma^M$. Analogously we define the corresponding strategy families Π^P , Π^F , Π^{PF} , Π^M , and Π^{PM} for player 2.

Given a finite-memory strategy $\sigma \in \Sigma^F$, let G_σ be the game graph obtained from G under the constraint that player 1 follows the strategy σ . The corresponding definition G_π for a player-2 strategy $\pi \in \Pi^F$ is analogous, and we write $G_{\sigma,\pi}$ for the game graph obtained from G if both players follow the finite-memory strategies σ and π , respectively. Observe that given a $2^{1/2}$ -player game graph G and a finite-memory player-1 strategy σ , the result G_σ is a player-2 MDP. Similarly, for a player-1 MDP G and a finite-memory player-1 strategy σ , the result G_σ is a Markov chain. Hence, if G is a $2^{1/2}$ -player game graph and the two players follow finite-memory strategies σ and π , the result $G_{\sigma,\pi}$ is a Markov chain. These observations will be useful in the analysis of $2^{1/2}$ -player games.

Objectives. An *objective* for a player consists of an ω -regular set of *winning plays* $\Phi \subseteq \Omega$ [20]. In this paper we study zero-sum games [10, 18], where the objectives of the two players are complementary; that is, if the objective of one player is Φ , then the objective of the other player is $\bar{\Phi} = \Omega \setminus \Phi$. We consider ω -regular objectives specified as Müller objectives. For a play $\omega = \langle s_0, s_1, s_2, \dots \rangle$, let $\text{Inf}(\omega)$ be the set $\{s \in S \mid s = s_k \text{ for infinitely many } k \geq 0\}$ of states that appear infinitely often in ω . We use colors to define objectives as in [7]. A $2^{1/2}$ -player game $(G, C, \chi, \mathcal{F} \subseteq \mathcal{P}(C))$ consists of a $2^{1/2}$ -player game graph G , a finite set C of colors, a partial function $\chi: S \rightarrow C$ that assigns colors to some states, and a winning condition specified by a subset \mathcal{F} of the power set $\mathcal{P}(C)$ of colors. The winning condition defines subset $\Phi \subseteq \Omega$ of winning plays, defined as follows:

$$\text{Müller}(\mathcal{F}) = \{ \omega \in \Omega \mid \chi(\text{Inf}(\omega)) \in \mathcal{F} \}$$

that is the set of paths ω such that the colors appearing infinitely often in ω is in \mathcal{F} .

Remarks. A winning condition $\mathcal{F} \subseteq \mathcal{P}(C)$ has a *split* if there are sets $C_1, C_2 \in \mathcal{F}$ such that $C_1 \cup C_2 \notin \mathcal{F}$. A winning condition is a *Rabin winning* condition if it do not have splits, and it is a *Streett winning* condition if $\mathcal{P}(C) \setminus \mathcal{F}$ does not have a split. This notions coincide with the Rabin and Streett winning conditions usually defined in the literature (see [16, 7] for details). We now define the reachability, safety, Büchi and coBüchi objectives that will be useful in the proofs of our results.

- *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 $\text{Reach}(T) = \{ \omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0 \}$. Given a set $F \subseteq S$, the safety objective requires that only states of F be visited. Thus, the set of winning plays is $\text{Safe}(F) = \{ \omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega \mid s_k \in F \text{ for all } k \geq 0 \}$.
- *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 $\text{coBüchi}(C) = \{ \omega \in \Omega \mid \text{Inf}(\omega) \subseteq C \}$.

Sure, almost-sure, positive winning and optimality. Given a player-1 objective Φ , a strategy $\sigma \in \Sigma$ is *sure winning* for player 1 from a state $s \in S$ if for every strategy $\pi \in \Pi$ for player 2, we have $\text{Outcome}(s, \sigma, \pi) \subseteq \Phi$. A strategy σ is *almost-sure winning* for player 1 from the state s for the objective Φ if for every player-2 strategy π , we have $\Pr_s^{\sigma, \pi}(\Phi) = 1$. A strategy σ is *positive winning* for player 1 from the state s for the objective Φ if for every player-2 strategy π , we have $\Pr_s^{\sigma, \pi}(\Phi) > 0$. The sure, almost-sure and positive winning strategies for player 2 are defined analogously. Given an objective Φ , the *sure winning set* $\langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi)$ for player 1 is the set of states from which player 1 has a sure winning strategy. Similarly, the *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ and the *positive winning set* $\langle\langle 1 \rangle\rangle_{\text{pos}}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning and a positive winning strategy, respectively. The sure winning set $\langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi)$, the almost-sure winning set $\langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi)$ and the positive winning set $\langle\langle 2 \rangle\rangle_{\text{pos}}(\Omega \setminus \Phi)$ for player 2 are defined analogously. It follows from the definitions that for all $2^{1/2}$ -player game graphs and all objectives Φ , we have $\langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{\text{pos}}(\Phi)$. Computing sure, almost-sure and positive winning sets and strategies is referred to as the *qualitative analysis* of $2^{1/2}$ -player games [8].

Given ω -regular objectives $\Phi \subseteq \Omega$ for player 1 and $\Omega \setminus \Phi$ for player 2, we define the *value functions* $\langle\langle 1 \rangle\rangle_{\text{val}}$ and $\langle\langle 2 \rangle\rangle_{\text{val}}$ for the players 1 and 2, respectively, as the following functions from the state space S to the interval $[0, 1]$ of reals: for all states $s \in S$, let $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{\text{val}}(\Omega \setminus \Phi)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \Pr_s^{\sigma, \pi}(\Omega \setminus \Phi)$. In other words, the value $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s)$ gives the maximal probability with which player 1 can achieve her objective Φ from state s , and analogously for player 2. The strategies that achieve the value are called *optimal*: a strategy σ for player 1 is *optimal* from the state s for the objective Φ if $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi)$. The optimal strategies for player 2 are defined analogously. Computing values and optimal strategies is referred to as the *quantitative analysis* of $2^{1/2}$ -player games. The set of states with value 1 is called the *limit-sure winning set* [8]. For $2^{1/2}$ -player game graphs with ω -regular objectives the almost-sure and limit-sure winning sets coincide [4].

Let $\mathcal{C} \in \{P, M, F, PM, PF\}$ and consider the family $\Sigma^{\mathcal{C}} \subseteq \Sigma$ of special strategies for player 1. We say that the family $\Sigma^{\mathcal{C}}$ *suffices* with respect to a player-1 objective Φ on a class \mathcal{G} of game graphs for *sure winning* if for every game graph $G \in \mathcal{G}$ and state $s \in \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi)$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi \in \Pi$, we have $\text{Outcome}(s, \sigma, \pi) \subseteq \Phi$. Similarly, the family $\Sigma^{\mathcal{C}}$ *suffices* with respect to the objective Φ on the class \mathcal{G} of game graphs for (a) *almost-sure winning* if for every game graph $G \in \mathcal{G}$ and state $s \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi \in \Pi$, we have $\Pr_s^{\sigma, \pi}(\Phi) = 1$; (b) *positive winning* if for every game graph $G \in \mathcal{G}$ and state $s \in \langle\langle 1 \rangle\rangle_{\text{pos}}(\Phi)$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi \in \Pi$, we have $\Pr_s^{\sigma, \pi}(\Phi) > 0$; and (c) *optimality* if for every game graph $G \in \mathcal{G}$ and state $s \in S$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi)$. The notion of sufficiency for size of finite-memory strategies is obtained by referring to the size of the memory M of the strategies. The notions of sufficiency of strategies for player 2 is defined analogously.

Determinacy. For sure winning, the $1^{1/2}$ -player and $2^{1/2}$ -player games coincide with 2-player (deterministic) games where the random player (who chooses the successor at the probabilistic states) is interpreted as an adversary, i.e., as player 2. Theorem 1 and Theorem 2 state the classical determinacy results for 2-player and $2^{1/2}$ -player game graphs with Müller objectives. It follows from Theorem 2 that for all Müller objectives Φ , for all $\varepsilon > 0$, there exists an ε -optimal strategy σ_ε for player 1 such that for all π and all $s \in S$ we have $\Pr_s^{\sigma_\varepsilon, \pi}(\Phi) \geq \langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) - \varepsilon$.

Theorem 1 (Qualitative determinacy [11]). *For all 2-player game graphs and Müller objectives Φ , we have $\langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \cap \langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi) = \emptyset$ and $\langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \cup \langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi) = S$. Moreover, on 2-player game graphs, the family of pure finite-memory strategies suffices for sure winning with respect to Müller objectives.*

Theorem 2 (Quantitative determinacy [14]). *For all $2^{1/2}$ -player game graphs, for all Müller winning conditions $\mathcal{F} \subseteq \mathcal{P}(C)$, and all states s , we have $\langle\langle 1 \rangle\rangle_{\text{val}}(\text{Müller}(\mathcal{F}))(s) + \langle\langle 2 \rangle\rangle_{\text{val}}(\Omega \setminus \text{Müller}(\mathcal{F}))(s) = 1$.*

3 Optimal Memory Bound for Pure Qualitative Winning Strategies

In this section we present optimal memory bounds for pure strategies with respect to qualitative (almost-sure and positive) winning for $2^{1/2}$ -player game graphs with Müller winning conditions. The result is obtained by a generalization of the result of [7] and depends on the novel constructions of Zielonka [21] for 2-player games. In [7] the authors use an insightful analysis of Zielonka's construction to present an upper bound (and also a matching lower bound) on memory of sure winning strategies in 2-player games with Müller objectives. In this section we generalize the result of [7] to show that the same upper bound holds for qualitative winning strategies in $2^{1/2}$ -player games with Müller objectives. We now introduce some notations and the Zielonka tree of a Müller condition.

Notation. Let $\mathcal{F} \subseteq \mathcal{P}(C)$ be a winning condition. For $D \subseteq C$ we define $(\mathcal{F} \upharpoonright D) \subseteq \mathcal{P}(D)$ as the set $\{D' \in \mathcal{F} \mid D' \subseteq D\}$. For a Müller condition $\mathcal{F} \subseteq \mathcal{P}(C)$ we denote by $\overline{\mathcal{F}}$ the complementary condition, i.e., $\overline{\mathcal{F}} = \mathcal{P}(C) \setminus \mathcal{F}$. Similarly for an objective Φ we denote by $\overline{\Phi}$ the complementary objective, i.e., $\overline{\Phi} = \Omega \setminus \Phi$.

Definition 1 (Zielonka tree of a winning condition [21]). *The Zielonka tree of a winning condition $\mathcal{F} \subseteq \mathcal{P}(C)$, denoted $\mathcal{Z}_{\mathcal{F},C}$, is defined inductively as follows:*

1. *If $C \notin \mathcal{F}$, then $\mathcal{Z}_{\mathcal{F},C} = \mathcal{Z}_{\overline{\mathcal{F}},C}$, where $\overline{\mathcal{F}} = \mathcal{P}(C) \setminus \mathcal{F}$.*
2. *If $C \in \mathcal{F}$, then the root of $\mathcal{Z}_{\mathcal{F},C}$ is labeled with C . Let C_0, C_1, \dots, C_{k-1} be all the maximal sets in $\{X \notin \mathcal{F} \mid X \subseteq C\}$. Then we attach to the root, as its subtrees, the Zielonka trees of $\mathcal{F} \upharpoonright C_i$, i.e., $\mathcal{Z}_{\mathcal{F} \upharpoonright C_i, C_i}$, for $i = 0, 1, \dots, k-1$.*

Hence the Zielonka tree is a tree with nodes labeled by sets of colors. A node of $\mathcal{Z}_{\mathcal{F},C}$ is a 0-level node if it is labeled with a set from \mathcal{F} , otherwise it is a 1-level node. In the sequel we write $\mathcal{Z}_{\mathcal{F}}$ to denote $\mathcal{Z}_{\mathcal{F},C}$ if C is clear from the context. ■

Definition 2 (The number $m_{\mathcal{F}}$ of Zielonka tree). *Let $\mathcal{F} \subseteq \mathcal{P}(C)$ be a winning condition and $\mathcal{Z}_{\mathcal{F}_0, C_0}, \mathcal{Z}_{\mathcal{F}_1, C_1}, \dots, \mathcal{Z}_{\mathcal{F}_{k-1}, C_{k-1}}$ be the subtrees attached to the root of the tree $\mathcal{Z}_{\mathcal{F},C}$, where $\mathcal{F}_i = \mathcal{F} \upharpoonright C_i \subseteq \mathcal{P}(C_i)$ for $i = 0, 1, \dots, k-1$. We define the number $m_{\mathcal{F}}$ inductively as follows*

$$m_{\mathcal{F}} = \begin{cases} 1 & \text{if } \mathcal{Z}_{\mathcal{F},C} \text{ does not have any subtrees,} \\ \max\{m_{\mathcal{F}_0}, m_{\mathcal{F}_1}, \dots, m_{\mathcal{F}_{k-1}}\} & \text{if } C \notin \mathcal{F}, (1\text{-level node}) \\ \sum_{i=1}^{k-1} m_{\mathcal{F}_i} & \text{if } C \in \mathcal{F}, (0\text{-level node}). \end{cases}$$

■

Our goal is to show that for winning conditions \mathcal{F} pure finite-memory qualitative winning strategies of size $m_{\mathcal{F}}$ exist in $2^{1/2}$ -player games. This proves the upper bound. The results of [7] already established the matching lower bound for 2-player games. This establishes the optimal bound of memory of qualitative winning strategies for $2^{1/2}$ -player games. We start with the key notion of *attractors* that will be crucial in our proofs.

Definition 3 (Attractors). *Given a $2^{1/2}$ -player game graph G and a set $U \subseteq S$ of states, such that $G \upharpoonright U$ is subgame, and $T \subseteq S$ we define $\text{Attr}_{1,\circ}(T, U)$ as follows:*

$$T_0 = T \cap U; \quad \text{and for } j \geq 0 \text{ we define } T_{j+1} \text{ from } T_j \text{ as}$$

$$T_{j+1} = T_j \cup \{s \in (S_1 \cup S_{\circ}) \cap U \mid E(s) \cap T_j \neq \emptyset\} \cup \{s \in S_2 \cap U \mid E(s) \cap U \subseteq T_j\}.$$

and $A = \text{Attr}_{1,\circ}(T,U) = \bigcup_{j \geq 0} T_j$. We obtain $\text{Attr}_{2,\circ}(T,U)$ by exchanging the roles of player 1 and player 2. A pure memoryless attractor strategy $\sigma^A : (A \setminus T) \cap S_1 \rightarrow S$ for player 1 on A to T is as follows: for $i > 0$ and a state $s \in (T_i \setminus T_{i-1}) \cap S_1$, the strategy $\sigma^A(s) \in T_{i-1}$ chooses a successor in T_{i-1} (which exists by definition). ■

Lemma 1 (Attractor properties). *Let G be a $2^{1/2}$ -player game graph and $U \subseteq S$ be a set of states such that $G \upharpoonright U$ is a subgame. For a set $T \subseteq S$ of states, let $Z = \text{Attr}_{1,\circ}(T,U)$. Then the following assertions hold.*

1. $G \upharpoonright (U \setminus Z)$ is a subgame.
2. Let σ^Z be a pure memoryless attractor strategy for player 1. For all strategies π for player 2 in the subgame $G \upharpoonright U$ and for all states $s \in U$ we have
 - (a) if $\Pr_s^{\sigma^Z, \pi}(\text{Reach}(Z)) > 0$, then $\Pr_s^{\sigma^Z, \pi}(\text{Reach}(T)) > 0$; and
 - (b) if $\Pr_s^{\sigma^Z, \pi}(\text{Büchi}(Z)) > 0$, then $\Pr_s^{\sigma^Z, \pi}(\text{Büchi}(T) \mid \text{Büchi}(Z)) = 1$.

Proof. We prove the following cases.

1. *Subgame property.* For a state $s \in U \setminus Z$, if $s \in S_1 \cup S_\circ$, then $E(s) \cap Z = \emptyset$, (otherwise s would have been in Z), i.e., $E(s) \cap U \subseteq U \setminus Z$. For a state $s \in S_2$, then $E(s) \cap (U \setminus Z) \neq \emptyset$ (otherwise s would have been in Z). It follows that $G \upharpoonright (U \setminus Z)$ is a subgame.
2. We now prove the two cases.
 - (a) *Positive probability reachability.* Let

$$\delta_{\min} = \min\{ \delta(s)(t) \mid s, t \in S_\circ, \delta(s)(t) > 0 \}.$$

Observe that $\delta_{\min} > 0$. Consider a strategy $\sigma_{1,\circ}^Z$ of both player 1 and the random player on Z as follows: player 1 follows an attractor strategy σ^Z on Z to T and for $s \in (T_i \setminus T_{i-1}) \cap S_\circ$, the random player chooses a successor $t \in T_{i-1}$. Such a successor exists by definition, and observe that such a choice is made in the game with probability at least δ_{\min} . The strategy $\sigma_{1,\circ}^Z$ ensures that for all states $s \in Z$ and for all strategies π for player 2 in $G \upharpoonright U$, the set $T \cap U$ is reached within $|Z|$ -steps. Given player 1 follows an attractor strategy σ^Z , the probability of the choice of $\sigma_{1,\circ}^Z$ is at least $\delta_{\min}^{|Z|}$. It follows that a pure memoryless attractor strategy σ^Z ensures that for all states $s \in Z$ and for all strategies π for player 2 in $G \upharpoonright U$ we have

$$\Pr_s^{\sigma^Z, \pi}(\text{Reach}(T)) \geq (\delta_{\min})^{|Z|} > 0.$$

The desired result follows.

- (b) *Almost-sure Büchi property.* Given a pure memoryless attractor strategy σ^Z , if the set Z is visited ℓ -times, then by the previous part we have that T is reached at least once with probability $1 - (1 - \delta_{\min}^{|Z|})^\ell$, which goes to 1 as $\ell \rightarrow \infty$. Hence for all states s and strategies π in $G \upharpoonright U$, given $\Pr_s^{\sigma^Z, \pi}(\text{Büchi}(Z)) > 0$, we have $\Pr_s^{\sigma^Z, \pi}(\text{Reach}(T) \mid \text{Büchi}(Z)) = 1$. Since given the event that Z is visited infinitely often (i.e., $\text{Büchi}(Z)$) the set T is reached with probability 1 from all states, it follows that the set T is visited infinitely often with probability 1. Formally, for all states s and strategies π in $G \upharpoonright U$, given $\Pr_s^{\sigma^Z, \pi}(\text{Büchi}(Z)) > 0$, we have $\Pr_s^{\sigma^Z, \pi}(\text{Büchi}(T) \mid \text{Büchi}(Z)) = 1$.

The result of the lemma follows. ■

Lemma 1 shows that the complement of an attractor is a subgame; and a pure memoryless attractor strategy ensures that if the attractor of a set T is reached with positive probability, then T is reached with positive probability, and given that the attractor of T is visited infinitely often, then T is visited infinitely often with probability 1. We now present the main result of this section (upper bound on memory for qualitative winning strategies). A matching lower bound follows from the results of [7] for 2-player games (see Theorem 4).




Fig. 1. The sets of the construction.

Theorem 3 (Qualitative forgetful determinacy). *Let $(G, C, \chi, \mathcal{F})$ be a $2^{1/2}$ -player game with Müller winning condition \mathcal{F} for player 1. Let $\Phi = \text{Müller}(\mathcal{F})$, and consider the following sets*

$$\begin{aligned} W_1^{>0} &= \langle\langle 1 \rangle\rangle_{\text{pos}}(\Phi); & W_1 &= \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi); \\ W_2^{>0} &= \langle\langle 2 \rangle\rangle_{\text{pos}}(\bar{\Phi}); & W_2 &= \langle\langle 2 \rangle\rangle_{\text{almost}}(\bar{\Phi}). \end{aligned}$$

The following assertions hold.

1. We have (a) $W_1^{>0} \cup W_2 = S$ and $W_1^{>0} \cap W_2 = \emptyset$; and (b) $W_2^{>0} \cup W_1 = S$ and $W_2^{>0} \cap W_1 = \emptyset$.
2. (a) Player 1 has a pure strategy σ with memory of size $m_{\mathcal{F}}$ such that for all states $s \in W_1^{>0}$ and for all strategies π for player 2 we have $\Pr_s^{\sigma, \pi}(\Phi) > 0$; and (b) player 2 has a pure strategy π with memory of size $m_{\bar{\mathcal{F}}}$ such that for all states $s \in W_2$ and for all strategies σ for player 1 we have $\Pr_s^{\sigma, \pi}(\bar{\Phi}) = 1$.
3. (a) Player 1 has a pure strategy σ with memory of size $m_{\mathcal{F}}$ such that for all states $s \in W_1$ and for all strategies π for player 2 we have $\Pr_s^{\sigma, \pi}(\Phi) = 1$; and (b) player 2 has a pure strategy π with memory of size $m_{\bar{\mathcal{F}}}$ such that for all states $s \in W_2^{>0}$ and for all strategies σ for player 1 we have $\Pr_s^{\sigma, \pi}(\bar{\Phi}) > 0$.

Proof. The first part of the result is a consequence of Theorem 2. We will concentrate on the proof for the result for part 2. The last part (part 3) follows from a symmetric argument.

The proof goes by induction on the structure of the Zielonka tree $Z_{\mathcal{F}, C}$ of the winning condition \mathcal{F} . We assume that $C \notin \mathcal{F}$. The case when $C \in \mathcal{F}$ can be proved by a similar argument: if $C \in \mathcal{F}$, then we consider $\hat{c} \notin C$ and consider the condition $\hat{\mathcal{F}} = \mathcal{F} \subseteq \mathcal{P}(C \cup \{\hat{c}\})$ and $C \cup \{\hat{c}\} \notin \hat{\mathcal{F}}$. Hence we consider without loss of generality $C \notin \mathcal{F}$ and let C_0, C_1, \dots, C_{k-1} be the label of the subtrees attached to the root C , i.e., C_0, C_1, \dots, C_{k-1} are maximal subset of colors that appear in \mathcal{F} . We will define by induction a non-decreasing sequence of sets $(U_j)_{j \geq 0}$ as follows. Let $U_0 = \emptyset$ and for $j > 0$ we define U_j below:

1. $A_j = \text{Attr}_{1, \circ}(U_{j-1}, S)$ and $X_j = S \setminus A_j$;
2. $D_j = C \setminus C_{j \bmod k}$ and $Y_j = X_j \setminus \text{Attr}_{2, \circ}(\chi^{-1}(D_j), X_j)$;
3. let Z_j be the set of positive winning states for player 1 in $(G \upharpoonright Y_j, C_{j \bmod k}, \chi, \mathcal{F} \upharpoonright C_{j \bmod k})$, (i.e., $Z_j = \langle\langle 1 \rangle\rangle_{\text{pos}}(\text{Müller}(\mathcal{F} \upharpoonright C_{j \bmod k}))$ in $G \upharpoonright Y_j$); hence $(Y_j \setminus Z_j)$ is almost-sure winning for player 2 in the subgame; and
4. $U_j = A_j \cup Z_j$.

Fig 1 describes all these sets. The property of attractors and almost-sure winning states ensure certain edges are forbidden between the sets. This is shown in Fig 2. We start with a few observations of the construction.




Fig. 2. The sets of the construction with forbidden edges.

1. *Observation 1.* For all $s \in S_2 \cap Z_j$, we have $E(s) \subseteq Z_j \cup A_j$. This follows from the following case analysis.
 - Since Y_j is a complement of an attractor set $\text{Attr}_{2, \circ}$, it follows that for all states $s \in S_2 \cap Y_j$ we have $E(s) \cap X_j \subseteq Y_j$. It follows that $E(s) \subseteq Y_j \cup A_j$.
 - Since player 2 can win almost-surely from the set $Y_j \setminus Z_j$, if a state $s \in Y_j \cap S_2$ has an edge to $Y_j \setminus Z_j$, then $s \in Y_j \setminus Z_j$.
2. *Observation 2.* For all $s \in X_j \cap (S_1 \cup S_\circ)$ we have (a) $E(s) \cap A_j = \emptyset$; else s would have been in A_j ; and (b) if $s \in Y_j \setminus Z_j$, then $E(s) \cap Z_j = \emptyset$ (else $s \in Z_j$).
3. *Observation 3.* For all $s \in Y_j \cap S_\circ$ we have $E(s) \subseteq Y_j$.

We will denote by \mathcal{F}_i the winning condition $\mathcal{F} \upharpoonright C_i$, for $i = 0, 1, \dots, k-1$, and $\overline{\mathcal{F}}_i = \mathcal{P}(C_i) \setminus \mathcal{F}_i$. By induction hypothesis on $\mathcal{F}_i = \mathcal{F} \upharpoonright C_{j \bmod k}$, player 1 has a pure positive winning strategy of size $m_{\mathcal{F}_i}$ from Z_j and player 2 has a pure almost-sure winning strategy of size $m_{\overline{\mathcal{F}}_i}$ from $Y_j \setminus Z_j$. Let $W = \bigcup_{j \geq 0} U_j$. We will show in Lemma 2 that player 1 has a pure positive winning strategy of size $m_{\mathcal{F}}$ from W ; and then in Lemma 3 we will show that player 2 has a pure almost-sure winning strategy of size $m_{\overline{\mathcal{F}}}$ from $S \setminus W$. This completes the proof. We now prove the Lemmas 2 and 3. ■

Lemma 2. *Player 1 has a pure positive winning strategy of size $m_{\mathcal{F}}$ from the set W .*

Proof. By induction hypothesis player 1 has a pure positive winning strategy σ_{j-1}^U of size $m_{\mathcal{F}}$ from U_{j-1} . From the set $A_j = \text{Attr}_{1, \circ}(U_{j-1}, S)$, player 1 has a pure memoryless attractor strategy σ_j^A to bring the game to U_{j-1} with positive probability (Lemma 1(part 2.(a))), and then use σ_{j-1}^U and ensure winning with positive probability from the set A_j . Let σ_j^Z be the pure positive winning strategy for player 1 of size $m_{\mathcal{F}_i}$, where $i = j \bmod k$. We now show the combination of strategies σ_{j-1}^U , σ_j^A and σ_j^Z ensure positive probability winning for player 1 from U_j . If the play starts at a state $s \in Z_j$, then player 1 follows σ_j^Z . If the play stays in Y_j for ever, then the strategy σ_j^Z ensures that player 1 wins with positive probability. By observation 1 of Theorem 3, for all states $s \in Y_j \cap S_2$, we have $E(s) \subseteq Y_j \cup A_j$. Hence if the play leaves Y_j , then player 2 must chose an edge to A_j . In A_j player 1 can use the attractor strategy σ_j^A followed by σ_{j-1}^U to ensure positive probability win. Hence if the play is in Y_j for ever with probability 1, then σ_j^Z ensures positive probability win, and if the play reaches A_j with positive probability, then σ_j^A followed by σ_{j-1}^U ensures positive probability win.

We now formally present σ_j^U . Let $\sigma_j^Z = (\sigma_{j,u}^Z, \sigma_{j,m}^Z)$ be the strategy obtained from inductive hypothesis; defined on Z_j of size $m_{\mathcal{F}_i}$, where $i = j \bmod k$, $\sigma_{j,u}^Z$ is the memory-update function and $\sigma_{j,m}^Z$ is the next-move function of σ_j^Z . We assume the memory $M_{\mathcal{F}_i}$ of σ_j^Z to be the set

$\{1, 2, \dots, m_{\mathcal{F}_i}\}$. The strategy $\sigma_j^A : (A_j \setminus U_{j-1}) \cap S_1 \rightarrow A_j$ is a pure memoryless attractor strategy on A_j to U_{j-1} . The strategy σ_j^U is as follows: the memory-update function is as follows

$$\sigma_{j,u}^U(s, m) = \begin{cases} \sigma_{j-1,u}^U(s, m) & s \in U_{j-1} \\ \sigma_{j-1,u}^Z(s, m) & s \in Z_j, m \in M_{\mathcal{F}_i} \\ 1 & \text{otherwise} \end{cases}$$

the next-move function is as follows

$$\sigma_{j,m}^U(s, m) = \begin{cases} \sigma_{j-1,m}^U(s, m) & s \in U_{j-1} \cap S_1 \\ \sigma_{j-1,m}^Z(s, m) & s \in Z_j \cap S_1, m \in M_{\mathcal{F}_i} \\ \sigma_{j-1,m}^Z(s, 1) & s \in Z_j \cap S_1, m \notin M_{\mathcal{F}_i} \\ \sigma_j^A(s) & s \in (A \setminus U_{j-1}) \cap S_1; \\ t & s \in (S \setminus U_j) \cap S_1; t \in E(s). \end{cases}$$

The strategy σ_j^U formally defines the strategy we described and proves the result. ■

Lemma 3. *Player 2 has a pure almost-sure winning strategy of size $m_{\overline{\mathcal{F}}}$ from the set $S \setminus W$.*

Proof. Let $\ell \in \mathbb{N}$ be such that $\ell \bmod k = 0$ and $W = U_{\ell-1} = U_\ell = U_{\ell+1} = \dots = U_{\ell+k-1}$. From the equality $W = U_{\ell-1} = U_\ell$ we have $\text{Attr}_{1,\circ}(W, S) = W$. Let us denote by $\overline{W} = S \setminus W$. Hence $G \upharpoonright \overline{W}$ is a subgame (by Lemma 1), and also for all $s \in \overline{W} \cap (S_1 \cup S_\circ)$ we have $E(s) \subseteq \overline{W}$. The equality $U_{\ell+i-1} = U_{\ell+i}$ implies that $Z_{\ell+i} = \emptyset$. Hence for all $i = 0, 1, \dots, k-1$, we have $Z_{\ell+i} = \emptyset$. By inductive hypothesis for all $i = 0, 1, \dots, k-1$, player 2 has a pure almost-sure winning strategy π^i of size $m_{\overline{\mathcal{F}}_i}$ in the game $(G \upharpoonright Y_{\ell+i}, C_i, \chi, \mathcal{F} \upharpoonright C_i)$.

We now describe the construction of a pure almost-sure winning strategy π^* for player 2 in \overline{W} . For $D_i = C \setminus C_i$ we denote by $\widehat{D}_i = \chi^{-1}(D_i)$ the set of states with colors D_i . If the play starts in a state in $Y_{\ell+i}$, for $i = 0, 1, \dots, k-1$, then player 2 uses the almost-sure winning strategy π^i . If the play leaves $Y_{\ell+i}$, then the play must reach $\overline{W} \setminus Y_{\ell+i} = \text{Attr}_{2,\circ}(\widehat{D}_i, \overline{W})$, since player 1 and random states do not have edges to W . In $\text{Attr}_{2,\circ}(\widehat{D}_i, \overline{W})$, player 2 plays a pure memoryless attractor strategy to reach the set \widehat{D}_i with positive probability. If the set \widehat{D}_i is reached, then a state in $Y_{(\ell+i+1) \bmod k}$ or in $\text{Attr}_{2,\circ}(\widehat{D}_{(i+1) \bmod k}, \overline{W})$ is reached. If $Y_{(\ell+i+1) \bmod k}$ is reached $\pi^{(i+1) \bmod k}$ is followed, and otherwise the pure memoryless attractor strategy to reach the set $\widehat{D}_{(i+1) \bmod k}$ with positive probability is followed. Of course, the play may leave $Y_{(\ell+i+1) \bmod k}$, and reach $Y_{(\ell+i+2) \bmod k}$, and then we would repeat the reasoning, and so on. Let us analyze various cases to prove that π^* is almost-sure winning for player 2.

1. If the play finally settles in some $Y_{\ell+i}$, for $i = 0, 1, \dots, k-1$, then from this moment player 2 follows π^i and ensures that the objective $\overline{\Phi}$ is satisfied with probability 1. Formally, for all states $s \in \overline{W}$, for all strategies σ for player 1 we have $\Pr_s^{\sigma, \pi^*}(\overline{\Phi} \mid \text{coBüchi}(Y_{\ell+i})) = 1$. This holds for all $i = 0, 1, \dots, k-1$ and hence for all states $s \in \overline{W}$, for all strategies σ for player 1 we have $\Pr_s^{\sigma, \pi^*}(\overline{\Phi} \mid \bigcup_{0 \leq i \leq k-1} \text{coBüchi}(Y_{\ell+i})) = 1$.
2. Otherwise, for all $i = 0, 1, \dots, k-1$, the set $\overline{W} \setminus Y_{\ell+i} = \text{Attr}_{2,\circ}(\widehat{D}_i, \overline{W})$ is visited infinitely often. By Lemma 1, given $\text{Attr}_{2,\circ}(\widehat{D}_i, \overline{W})$ is visited infinitely often, then the attractor strategy ensures that the set \widehat{D}_i is visited infinitely often with probability 1. Formally, for all states $s \in \overline{W}$, for all strategies σ for player 1, for all $i = 0, 1, \dots, k-1$, we have $\Pr_s^{\sigma, \pi^*}(\text{Büchi}(\widehat{D}_i) \mid \text{Büchi}(\overline{W} \setminus Y_{\ell+i})) = 1$; and also $\Pr_s^{\sigma, \pi^*}(\text{Büchi}(\widehat{D}_i) \mid \bigcap_{0 \leq i \leq k-1} \text{Büchi}(\overline{W} \setminus Y_{\ell+i})) = 1$. It follows that for all states $s \in \overline{W}$, for all strategies σ for player 1 we have $\Pr_s^{\sigma, \pi^*}(\bigcap_{0 \leq i \leq k-1} \text{Büchi}(\widehat{D}_i) \mid \bigcap_{0 \leq i \leq k-1} \text{Büchi}(\overline{W} \setminus Y_{\ell+i})) = 1$. Hence the play visits states with colors not in C_i with probability 1. Hence the set of colors visited infinitely often is not contained in any C_i . Since C_0, C_1, \dots, C_{k-1} are all the maximal subsets of \mathcal{F} , we have the set of colors visited infinitely often is not in \mathcal{F} with probability 1, and hence player 2 wins almost-surely.

Hence it follows that for all strategies σ and for all states $s \in (S \setminus W)$ we have $\Pr_s^{\sigma, \pi^*}(\overline{\Phi}) = 1$. To complete the proof we present precise description of the strategy π^* with memory of size $m_{\mathcal{F}}$. Let $\pi^i = (\pi_u^i, \pi_m^i)$ be an almost-sure winning strategy for player 2 for the subgame on $Y_{\ell+i}$ with memory $M_{\mathcal{F}_i}$. By definition we have $m_{\mathcal{F}} = \sum_{i=0}^{k-1} m_{\mathcal{F}_i}$. Let $M_{\mathcal{F}} = \bigcup_{i=0}^{k-1} (M_{\mathcal{F}_i} \times \{i\})$. This set is not exactly the set $\{1, 2, \dots, m_{\mathcal{F}}\}$, but has the same cardinality (which suffices for our purpose). We define the strategy π^* as follows:

$$\pi_u^*(s, (m, i)) = \begin{cases} \pi_u^i(s, (m, i)) & s \in Y_{\ell+i} \\ (1, i+1 \bmod k) & \text{otherwise.} \end{cases}$$

$$\pi_m^*(s, (m, i)) = \begin{cases} \pi_m^i(s, (m, i)) & s \in Y_{\ell+i} \\ \pi^{L_i}(s) & s \in L_i \setminus \widehat{D}_i \\ s_i & s \in \widehat{D}_i, s_i \in E(s) \cap \overline{W}. \end{cases}$$

where $L_i = \text{Attr}_{2, \circ}(\widehat{D}_i, \overline{W})$; π^{L_i} is a pure memoryless attractor strategy on L_i to \widehat{D}_i , and s_i is a successor state of s in \overline{W} (such a state exists since \overline{W} induces a subgame). This formally represents π^* and the size of π^* satisfies the required bound. Observe that the disjoint sum of all $M_{\mathcal{F}_i}$ was required since $Y_{\ell}, Y_{\ell+1}, \dots, Y_{\ell+k-1}$ may not be disjoint and the strategy π^* need to know which Y_j the play is in. ■

A careful analysis of the strategy synthesis (the proof of Theorem 3) also implicitly presents an algorithm to compute the almost-sure winning and positive winning states.

Corollary 1 (Strategy synthesis algorithm). *There is an algorithm that given a game $(G, C, \chi, \mathcal{F})$ computes an almost-sure winning strategy and the almost-sure winning set in $O((|S| + |E|) \cdot d^h)$ time; where d is the maximum degree of a node and h is the height of the Zielonka tree $Z_{\mathcal{F}}$.*

Lower bound. In [7] the authors show a matching lower bound for sure winning strategies in 2-player games. It may be noted that in 2-player games any pure almost-sure winning or any pure positive winning strategy is also a sure winning strategy. This observation along with the result of [7] gives us the following result.

Theorem 4 (Lower bound [7]). *For all Müller winning conditions $\mathcal{F} \subseteq \mathcal{P}(C)$, there is a 2-player game $(G, C, \chi, \mathcal{F})$ (with a 2-player game graph G) such that every pure almost-sure and positive winning strategy for player 1 requires memory of size at least $m_{\mathcal{F}}$; and every pure almost-sure and positive winning strategy for player 2 requires memory of size at least $m_{\mathcal{F}}$.*

4 Optimal Memory Bound for Pure Optimal Strategies

In this section we extend the sufficiency results for families of strategies from almost-sure winning to optimality with respect to all Müller objectives. In the following, we fix a $2^{1/2}$ -player game graph G . We first present a useful proposition and then some definitions. Since Müller objectives are infinitary objectives the following proposition is immediate.

Proposition 1 (Optimality conditions). *For all Müller objectives Φ , for every $s \in S$ the following conditions hold.*

1. *If $s \in S_1$, then for all $t \in E(s)$ we have $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) \geq \langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(t)$, and for some $t \in E(s)$ we have $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(t)$.*
2. *If $s \in S_2$, then for all $t \in E(s)$ we have $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) \leq \langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(t)$, and for some $t \in E(s)$ we have $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(t)$.*

3. If $s \in S_{\circ}$, then $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = (\sum_{t \in E(s)} \langle\langle 1 \rangle\rangle_{val}(\Phi)(t) \cdot \delta(s)(t))$.

Similar conditions hold for the value function $\langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)$ of player 2.

Definition 4 (Value classes). Given a Müller objective Φ , for every real $r \in [0, 1]$ the value class with value r is $VC(\Phi, r) = \{s \in S \mid \langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = r\}$ is the set of states with value r for player 1. For $r \in [0, 1]$ we denote by $VC(\Phi, > r) = \bigcup_{q > r} VC(\Phi, q)$ the value classes greater than r and by $VC(\Phi, < r) = \bigcup_{q < r} VC(\Phi, q)$ value classes smaller than r . ■

Definition 5 (Boundary probabilistic states). Given a value class $VC(\Phi, r)$ a state $s \in VC(\Phi, r) \cap S_{\circ}$ is a boundary probabilistic state if $E(s) \cap (S \setminus VC(\Phi, r)) \neq \emptyset$, i.e., the probabilistic state has an edge out of the value class. For a value class $VC(\Phi, r)$ we denote by $Bnd(\Phi, r)$ the set of boundary probabilistic states of value class r . ■

Observation. It follows from Proposition 1 that for a state $s \in Bnd(\Phi, r)$ we have $E(s) \cap VC(\Phi, > r) \neq \emptyset$ and $E(s) \cap VC(\Phi, < r) \neq \emptyset$, i.e., the boundary probabilistic states have edges to higher and lower value classes. It follows that for all Müller objectives Φ we have $Bnd(\Phi, 1) = \emptyset$ and $Bnd(\Phi, 0) = \emptyset$.

Reduction of a value class. Given a value class $VC(\Phi, r)$, let $Bnd(\Phi, r)$ be the set of boundary probabilistic states in $VC(\Phi, r)$. We denote by $G_{Bnd(\Phi, r)}$ the subgame where every boundary probabilistic state in $Bnd(\Phi, r)$ is converted to an absorbing state (state with a self-loop). We denote by $G_{\Phi, r} = G_{Bnd(\Phi, r)} \upharpoonright VC(\Phi, r)$: this is a subgame since every value class is δ -live by Proposition 1, and δ -closed as all states in $Bnd(\Phi, r)$ is converted to absorbing states.

Lemma 4 (Almost-sure reduction). Let G be a $2^{1/2}$ -player game graph and $\mathcal{F} \subseteq \mathcal{P}(C)$ be a Müller winning condition. Let $\Phi = \text{Müller}(\mathcal{F})$. For $0 < r < 1$, the following assertions hold.

1. Player 1 wins almost-surely for objective $\Phi \cup \text{Reach}(Bnd(\Phi, r))$ from all states in $G_{\Phi, r}$, i.e., $\langle\langle 1 \rangle\rangle_{almost}(\Phi \cup \text{Reach}(Bnd(\Phi, r))) = VC(\Phi, r)$ in the subgame $G_{\Phi, r}$.
2. Player 2 wins almost-surely for objective $\bar{\Phi} \cup \text{Reach}(Bnd(\Phi, r))$ from all states in $G_{\Phi, r}$, i.e., $\langle\langle 2 \rangle\rangle_{almost}(\bar{\Phi} \cup \text{Reach}(Bnd(\Phi, r))) = VC(\Phi, r)$ in the subgame $G_{\Phi, r}$.

Proof. We prove the first part and the second part follows from symmetric arguments. The result is obtained through an argument by contradiction. Let $0 < r < 1$, and let

$$q = \max\{\langle\langle 1 \rangle\rangle_{val}(\Phi)(t) \mid t \in E(s) \setminus VC(\Phi, r), s \in VC(\Phi, r) \cap S_1\},$$

that is, q is the maximum value a successor state t of a player 1 state $s \in VC(\Phi, r)$ such that the successor state t is not in $VC(\Phi, r)$. By Proposition 1 we must have $q < r$. Hence by escaping from the value class $VC(\Phi, r)$ player 1 gets to see a state with value at most $q < r$. We consider the subgame $G_{\Phi, r}$. Let $U = VC(\Phi, r)$ and $Z = Bnd(\Phi, r)$. Assume towards contradiction, there exists a state $s \in U$ such that $s \notin \langle\langle 1 \rangle\rangle_{almost}(\Phi \cup \text{Reach}(Z))$. Then we have $s \in (U \setminus Z)$ and $\langle\langle 2 \rangle\rangle_{val}(\bar{\Phi} \cap \text{Safe}(U \setminus Z))(s) > 0$. It follows from the results of [2] that for all Müller objectives Ψ , if $\langle\langle 2 \rangle\rangle_{val}(\Psi)(s) > 0$, then for some state s_1 we have $\langle\langle 2 \rangle\rangle_{val}(\Psi)(s_1) = 1$. Observe that in $G_{\Phi, r}$ we have all states in Z are absorbing states, and hence the objective $\bar{\Phi} \cap \text{Safe}(U \setminus Z)$ is equivalent to the objective $\bar{\Phi} \cap \text{coBüchi}(U \setminus Z)$, which is a Müller objective. It follows that there exists a state $s_1 \in (U \setminus Z)$ such that $\langle\langle 2 \rangle\rangle_{val}(\bar{\Phi} \cap \text{Safe}(U \setminus Z)) = 1$. Hence there exists a strategy $\hat{\pi}$ for player 2 in $G_{\Phi, r}$ such that for all strategies $\hat{\sigma}$ for player 1 in $G_{\Phi, r}$ we have $\Pr_{s_1}^{\hat{\sigma}, \hat{\pi}}(\bar{\Phi} \cap \text{Safe}(U \setminus Z)) = 1$. We will now construct a strategy π^* for player 2 as a combination of the strategy $\hat{\pi}$ and a strategy in the original game G . By Martin's determinacy result (Theorem 2), for all $\varepsilon > 0$, there exists an ε -optimal strategy π_ε for player 2 in G such that for all $s \in S$ and for all strategies σ for player 1 we have

$$\Pr_s^{\sigma, \pi_\varepsilon}(\bar{\Phi}) \geq \langle\langle 2 \rangle\rangle_{val}(\bar{\Phi})(s) - \varepsilon.$$

Let $r - q = \alpha > 0$, and let $\varepsilon = \frac{\alpha}{2}$ and consider an ε -optimal strategy for player 2 in G . The strategy π^* in G is constructed as follows: for a history w that remains in U , player 2 follows $\hat{\pi}$; and if the history reaches $(S \setminus U)$, then player 2 follows the strategy π_ε . Formally, for a history $w = \langle s_1, s_2, \dots, s_k \rangle$ we have

$$\pi^*(w) = \begin{cases} \hat{\pi}(w) & \text{if for all } 1 \leq j \leq k. s_j \in U; \\ \pi_\varepsilon(s_j, s_{j+1}, \dots, s_k) & \text{where } j = \min\{i \mid s_i \notin U\} \end{cases}$$

We consider the case when the play starts at s_1 . The strategy π^* ensures the following: if the game stays in U , then the strategy $\hat{\pi}$ is followed, and given the play stays in U , the strategy $\hat{\pi}$ ensures with probability 1 that $\bar{\Phi}$ is satisfied and $\text{Bnd}(\Phi, r)$ is not reached. Hence if the game escapes U , then it reaches a state with value at most q for player 1. We consider an arbitrary strategy σ for player 1 and consider the following cases.

1. If $\Pr_{s_1}^{\sigma, \pi^*}(\text{Safe}(U)) = 1$, then we have $\Pr_{s_1}^{\sigma, \pi^*}(\bar{\Phi} \cap \text{Safe}(U)) = \Pr_{s_1}^{\sigma, \hat{\pi}}(\bar{\Phi} \cap \text{Safe}(U)) = 1$. Hence we also have $\Pr_{s_1}^{\sigma, \hat{\pi}}(\bar{\Phi}) = 1$, i.e., we have $\Pr_{s_1}^{\sigma, \pi^*}(\Phi) = 0$.
2. If $\Pr_{s_1}^{\sigma, \pi^*}(\text{Reach}(S \setminus U)) = 1$, then the play reaches a state with value for player 1 at most q and the strategy π_ε ensures that $\Pr_{s_1}^{\sigma, \pi^*}(\Phi) \leq q + \varepsilon$.
3. If $\Pr_{s_1}^{\sigma, \pi^*}(\text{Safe}(U)) > 0$ and $\Pr_{s_1}^{\sigma, \pi^*}(\text{Reach}(S \setminus U)) > 0$, then we condition on both these events and have the following:

$$\begin{aligned} \Pr_{s_1}^{\sigma, \pi^*}(\Phi) &= \Pr_{s_1}^{\sigma, \pi^*}(\Phi \mid \text{Safe}(U)) \cdot \Pr_{s_1}^{\sigma, \pi^*}(\text{Safe}(U)) \\ &\quad + \Pr_{s_1}^{\sigma, \pi^*}(\Phi \mid \text{Reach}(S \setminus U)) \cdot \Pr_{s_1}^{\sigma, \pi^*}(\text{Reach}(S \setminus U)) \\ &\leq 0 + (q + \varepsilon) \cdot \Pr_{s_1}^{\sigma, \pi^*}(\text{Reach}(S \setminus U)) \\ &\leq q + \varepsilon. \end{aligned}$$

The above inequalities are obtained as follows: given the event $\text{Safe}(U)$, the strategy π^* follows $\hat{\pi}$ and ensures that $\bar{\Phi}$ is satisfied with probability 1 (i.e., Φ is satisfied with probability 0); else the game reaches states where the value for player 1 is at most q , and then the analysis is similar to the previous case.

Hence for all strategies σ we have

$$\Pr_{s_1}^{\sigma, \pi^*}(\Phi) \leq q + \varepsilon = q + \frac{\alpha}{2} = r - \frac{\alpha}{2}.$$

Hence we must have $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s_1) \leq r - \frac{\alpha}{2}$. Since $\alpha > 0$ and $s_1 \in \text{VC}(\Phi, r)$ (i.e., $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s_1) = r$), we have a contradiction. The desired result follows. ■

Lemma 5 (Almost-sure to optimality [4]). *Let G be a $2^{1/2}$ -player game graph and $\mathcal{F} \subseteq \mathcal{P}(C)$ be a Müller winning condition. Let $\Phi = \text{Müller}(\mathcal{F})$. Let σ be a strategy such that*

- σ is an almost-sure winning strategy from the almost-sure winning states ($\langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ in G);
- and
- σ is an almost-sure winning strategy for objective $\Phi \cup \text{Reach}(\text{Bnd}(\Phi, r))$ in the game $G_{\Phi, r}$, for all $0 < r < 1$.

Then σ is an almost-sure winning strategy.

Proof. We prove the result for the case when σ is memoryless (randomized memoryless). The case when σ is finite-memory with memory \mathbb{M} , the arguments can be repeated on the game $G \times \mathbb{M}$ (the usual synchronous product of G and the memory \mathbb{M}).

Consider the player-2 MDP G_σ with the objective Müller($\overline{\mathcal{F}}$) for player 2. In MDPs with Müller objectives randomized memoryless optimal strategies exist [3]. We fix a randomized memoryless optimal strategy π for player 2 in G_σ . Let $W_1 = \langle\langle 1 \rangle\rangle_{almost}(\overline{\Phi})$ and $W_2 = \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$. We consider the Markov chain $G_{\sigma,\pi}$ and analyze the recurrent states of the Markov chain.

Recurrent states in $G_{\sigma,\pi}$. Let U be a closed, connected recurrent set in $G_{\sigma,\pi}$ (i.e., U is a bottom strongly connected component in the $G_{\sigma,\pi}$). Let $q = \max\{r \mid \text{VC}(\Phi, r) \cap U \neq \emptyset\}$, i.e., for all $q' > q$ we have $\text{VC}(\Phi, q') \cap U = \emptyset$ or in other words $\text{VC}(\Phi, > q) \cap U = \emptyset$. For a state $s \in U \cap \text{VC}(\Phi, q)$ we have the following cases.

1. If $s \in S_1$, then $\text{Supp}(\sigma(s)) \subseteq \text{VC}(\Phi, q)$. This is because in the game $G_{\Phi,q}$ the edges of player 1 consists of edges in the value class $\text{VC}(\Phi, q)$
2. If $s \in S_\circ$ and $s \in \text{Bnd}(\Phi, q)$, then it means that $U \cap \text{VC}(\Phi, q') \neq \emptyset$, for some $q' > q$: this is because $E(s) \cap \text{VC}(\Phi, > q) \neq \emptyset$ for $s \in \text{Bnd}(\Phi, q)$ and U is closed. This is not possible since by assumption on U we have $U \cap \text{VC}(\Phi, > q) = \emptyset$. Hence we have $s \in S_\circ \cap (U \setminus \text{Bnd}(\Phi, q))$, and $E(s) \subseteq \text{VC}(\Phi, q)$.
3. If $s \in S_2$, then since $U \cap \text{VC}(\Phi, > q) = \emptyset$, it follows by Proposition 1 that $\text{Supp}(\pi(s)) \subseteq \text{VC}(\Phi, q)$.

Hence for all $s \in U \cap \text{VC}(\Phi, q)$ we have all successors of U in $G_{\sigma,\pi}$ are in $\text{VC}(\Phi, q)$, and moreover $U \cap \text{Bnd}(\Phi, q) = \emptyset$, i.e., U is contained in a value class and does not intersect with the boundary probabilistic states. By the property of strategy σ , if $U \cap (S \setminus W_2) \neq \emptyset$, then for all $s \in U$ we have $\text{Pr}_s^{\sigma,\pi}(\Phi) = 1$: this is because for all $r > 0$, the strategy σ is almost-sure winning for objective $\Phi \cup \text{Reach}(\text{Bnd}(\Phi, r))$ in $G_{\Phi,r}$. Since σ is a fixed strategy and π is optimal against σ , it follows that if $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) < 1$, then $\text{Pr}_s^{\sigma,\pi}(\Phi) < 1$. Hence it follows that $U \cap (S \setminus (W_1 \cup W_2)) = \emptyset$. Hence the recurrent states of $G_{\sigma,\pi}$ are contained in $W_1 \cup W_2$, i.e., we have $\text{Pr}_s^{\sigma,\pi}(\text{Reach}(W_1 \cup W_2)) = 1$. Since σ is an almost-sure winning strategy in W_1 , we have $\text{Pr}_s^{\sigma,\pi}(\overline{\Phi}) = \text{Pr}_s^{\sigma,\pi}(\text{Reach}(W_2))$. Hence the strategy π maximizes the probability to reach W_2 in the MDP G_σ .

Analyzing reachability in G_σ . Since in G_σ player 2 maximizes the probability to reachability to W_2 , we analyze the player-2 MDP G_σ with objective $\text{Reach}(W_2)$ for player 2. For every state s consider a real-valued variable $x_s = 1 - \langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \langle\langle 2 \rangle\rangle_{val}(\overline{\Phi})(s)$. The following constraints are satisfied

$$\begin{aligned} x_s &= \sum_{t \in \text{Supp}(\sigma(s))} x_t \cdot \sigma(s)(t) & s \in S_1; \\ x_s &= \sum_{t \in E(s)} x_t \cdot \delta(s)(t) & s \in S_\circ; \\ x_s &\geq x_t & s \in S_2; \\ x_s &= 1 & s \in W_2; \end{aligned}$$

The first equality follows as for all $r \in [0, 1]$ and for all $s \in S \cap \text{VC}(\Phi, r)$ we have $\text{Supp}(\sigma(s)) \subseteq \text{VC}(\Phi, r)$. The next equality and the first inequality follows from Proposition 1. Since the values for MDPs with reachability objective is characterized as the least value vector satisfying the above constraints [10], it follows that for all $s \in S$ and for all strategies $\pi \in \Pi$ we have

$$\text{Pr}_s^{\sigma,\pi}(\text{Reach}(W_2)) \leq x_s = \langle\langle 2 \rangle\rangle_{val}(\overline{\Phi})(s).$$

Hence we have $\text{Pr}_s^{\sigma,\pi}(\overline{\Phi}) \leq \langle\langle 2 \rangle\rangle_{val}(\overline{\Phi})(s)$, i.e., $\text{Pr}_s^{\sigma,\pi}(\Phi) \geq 1 - \langle\langle 2 \rangle\rangle_{val}(\overline{\Phi})(s) = \langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$. Thus we obtain that σ is an optimal strategy. ■

Müller reduction for $G_{\Phi,r}$. Given a Müller winning condition \mathcal{F} and the objective $\Phi = \text{Müller}(\mathcal{F})$, we consider the game $G_{\Phi,r}$ with the objective $\Phi \cup \text{Reach}(\text{Bnd}(\Phi, r))$ for player 1. We present a simple reduction to a game with objective Φ . The reduction is achieved as follows: without loss of generality we assume $\mathcal{F} \neq \emptyset$, and let $F \in \mathcal{F}$ and $F = \{c_1^F, c_2^F, \dots, c_f^F\}$. We construct a game graph $\tilde{G}_{\Phi,r}$ with objective Φ for player 1 as follows: convert every state $s_j \in \text{Bnd}(\Phi, r)$ to a cycle $U_j = \{s_1^j, s_2^j, \dots, s_f^j\}$ with $\chi(s_i^j) = c_i^F$, i.e., once s_j is reached the cycle U_j is repeated with $\chi(U_j) \in \mathcal{F}$. An almost-sure winning strategy in $G_{\Phi,r}$ with objective $\Phi \cup \text{Reach}(\text{Bnd}(\Phi, r))$, is an almost-sure winning strategy in $\tilde{G}_{\Phi,r}$ with objective Φ ; and vice-versa. The present reduction

along with Lemma 4 and Lemma 5 gives us Lemma 6. Lemma 6 along with Theorem 3 gives us Theorem 5.

Lemma 6. *For all Müller winning conditions \mathcal{F} , the following assertions hold.*

1. *If the family of pure finite-memory strategies of size $\ell_{\mathcal{F}}^P$ suffices for almost-sure winning on $2^{1/2}$ -player game graphs, then the family of pure finite-memory strategies of size $\ell_{\mathcal{F}}^P$ suffices for optimality on $2^{1/2}$ -player game graphs.*
2. *If the family of randomized finite-memory strategies of size $\ell_{\mathcal{F}}^R$ suffices for almost-sure winning on $2^{1/2}$ -player game graphs, then the family of randomized finite-memory strategies of size $\ell_{\mathcal{F}}^R$ suffices for optimality on $2^{1/2}$ -player game graphs.*

Theorem 5. *For all Müller winning conditions \mathcal{F} , the family of pure finite-memory strategies of size $m_{\mathcal{F}}$ suffices for optimality on $2^{1/2}$ -player game graphs.*

5 An Improved Bound for Randomized Strategies

We now show that if a player plays randomized strategies, then the upper bound on memory for optimal strategies can be improved. We first present the notions of an upward closed restriction of a Zielonka tree. The number $m_{\mathcal{F}}^U$ of such restrictions of the Zielonka tree will be in general lower than the number $m_{\mathcal{F}}$ of Zielonka trees, and we show that randomized strategies with memory of size $m_{\mathcal{F}}^U$ suffices for optimality.

Upward closed sets. A set $\mathcal{F} \subseteq \mathcal{P}(C)$ is *upward closed* if for all $F \in \mathcal{F}$ and all $F \subseteq F_1$ we have $F_1 \in \mathcal{F}$, i.e., if a set F is in \mathcal{F} , then all supersets F_1 of F are in \mathcal{F} as well.

Upward closed restriction of Zielonka tree. The upward closed restriction of a Zielonka tree for a Müller winning condition $\mathcal{F} \subseteq \mathcal{P}(C)$, denoted as $\mathcal{Z}_{\mathcal{F},C}^U$, is obtained by making upward closed conditions as leaves. Formally, we define $\mathcal{Z}_{\mathcal{F},C}^U$ inductively as follows:

1. if \mathcal{F} is upward closed, then $\mathcal{Z}_{\mathcal{F},C}^U$ is leaf labeled \mathcal{F} (i.e., it has no subtrees);
2. otherwise
 - (a) if $C \notin \mathcal{F}$, then $\mathcal{Z}_{\mathcal{F},C}^U = \mathcal{Z}_{\overline{\mathcal{F}},C}^U$, where $\overline{\mathcal{F}} = \mathcal{P}(C) \setminus \mathcal{F}$.
 - (b) if $C \in \mathcal{F}$, then the root of $\mathcal{Z}_{\mathcal{F},C}^U$ is labeled with C ; and let C_0, C_1, \dots, C_{k-1} be all the maximal sets in $\{X \notin \mathcal{F} \mid X \subseteq C\}$; then we attach to the root, as its subtrees, the Zielonka upward closed restricted trees $\mathcal{Z}_{\mathcal{F},C}^U$ of $\mathcal{F} \upharpoonright C_i$, i.e., $\mathcal{Z}_{\mathcal{F} \upharpoonright C_i, C_i}^U$, for $i = 0, 1, \dots, k-1$.

The number $m_{\mathcal{F}}^U$ for $\mathcal{Z}_{\mathcal{F},C}^U$ is the number defined as the number $m_{\mathcal{F}}$ was defined for the tree $\mathcal{Z}_{\mathcal{F},C}$.

We will prove randomized strategies of size $m_{\mathcal{F}}^U$ suffices for optimality. To prove this result, we first prove that randomized strategies of size $m_{\mathcal{F}}^U$ suffices for almost-sure winning. The result then follows from Lemma 6. To prove the result for almost-sure winning we take a closer look at the proof of Theorem 3. The inductive proof characterizes that if existence of randomized memoryless strategies can be proved for $2^{1/2}$ -player games with Müller winning conditions that appear in the leaves of the Zielonka tree, then the inductive proof generalizes to give a bound as in Theorem 3. Hence to prove upper bound $m_{\mathcal{F}}^U$ for almost-sure winning, it suffices to show that randomized memoryless strategies suffices for upward closed Müller winning conditions. In [3] it was shown that for all $2^{1/2}$ -player games randomized memoryless strategies suffices for almost-sure winning for upward closed objectives (see Appendix for a proof). This gives us Theorem 6.

Theorem 6. *For all Müller winning conditions \mathcal{F} , the family of randomized finite-memory strategies of size $m_{\mathcal{F}}^U$ suffices for optimality on $2^{1/2}$ -player game graphs.*

Remark. In general we have $m_{\mathcal{F}}^U < m_{\mathcal{F}}$. Consider for example $\mathcal{F} \subseteq \mathcal{P}(C)$, where $C = \{c_1, c_2, \dots, c_k\}$. For the Müller winning condition $\mathcal{F} = \{C\}$. We have $m_{\mathcal{F}}^U = 1$, and $m_{\mathcal{F}} = |C|$.

6 Conclusion

In this work we present optimal memory bounds for pure almost-sure, positive and optimal strategies for $2^{1/2}$ -player games with Müller winning conditions. We also present improved memory bounds for randomized strategies. Unlike the results of [7] our results do not extend to infinite state games: for example, the results of [9] showed that even for $2^{1/2}$ -player pushdown games optimal strategies need not exist, and for $\varepsilon > 0$ even ε -optimal strategies may require infinite memory. For lower bound of randomized strategies the constructions of [7] do not work: in fact for the family of games used for lower bounds in [7] randomized memoryless almost-sure winning strategies exist. However, it is known that there exist Müller winning conditions $\mathcal{F} \subseteq \mathcal{P}(C)$, such that randomized almost-sure winning strategies may require memory $|C|!$ [12]. However, whether a matching lower bound of size $m_{\mathcal{F}}^U$ can be proved in general, or whether the upper bound of $m_{\mathcal{F}}^U$ can be improved and a matching lower bound can be proved for randomized strategies with memory remains open.

References

1. J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the AMS*, 138:295–311, 1969.
2. K. Chatterjee. Concurrent games with tail objectives. In *CSL'06*, Springer.
3. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. Trading memory for randomness. In *QEST'04* IEEE Computer Society Press, 2004.
4. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. The complexity of stochastic Rabin and Streett games. In *ICALP'05* vol. 3580 of *LNCS*, pages 878–890. Springer.
5. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA'04*, pages 114–123. SIAM, 2004.
6. A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
7. S. Dziembowski, M. Jurdziński, and I. Walukiewicz. How much memory is needed to win infinite games? In *LICS'97*, pages 99–110. IEEE Computer Society Press, 1997.
8. L. de Alfaro and T.A. Henzinger. Concurrent ω -regular games. In *LICS'00*, pages 141–154. IEEE Computer Society, 2000.
9. K. Etessami and M. Yannakakis Recursive Markov decision processes and recursive stochastic games. In *ICALP'05* vol. 3580 of *LNCS*, pages 891–903. Springer.
10. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
11. Y. Gurevich and L. Harrington. Trees, automata, and games. In *STOC'82*, pages 60–65. ACM, 1982.
12. R. Majumdar. *Symbolic algorithms for verification and control*. PhD Thesis, UC Berkeley, 2003.
13. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1992.
14. D.A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63:1565–1581, 1998.
15. A.K. McIver and C.C. Morgan. Games, probability, and the quantitative μ -calculus $qm\mu$. In *LPAR'02*, volume 2514 of *LNAI*, pages 292–310. Springer, 2002.
16. D. Niwiński Fixed-point characterization of infinite behavior of finite-state systems. In *Theoretical Computer Science*, 189(1-2): 1–69, 1997.
17. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM, 1989.
18. 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.
19. P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25:206–230, 1987.
20. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3 (*Beyond Words*), pages 389–455. Springer, 1997.
21. W. Zielonka Infinite games on finitely coloured graphs with applications to automata on infinite trees. In *Theoretical Computer Science*, 200(1-2): 135–183, 1998.
22. W. Zielonka. Perfect-information stochastic parity games. In *FoSSaCS'04*, volume 2987 of *LNCS*, pages 499–513. Springer, 2004.

Appendix

Theorem 7 ([3]). *The family of randomized memoryless strategies suffices for almost-sure winning with respect to upward closed objectives on $2^{1/2}$ -player game graphs.*

Proof. Consider a $2^{1/2}$ -player game graph G and the game $(G, C, \chi, \mathcal{F})$ with an upward closed objective $\Phi = \text{Müller}(\mathcal{F})$ for player 1, i.e., \mathcal{F} is upward closed. Let $W_1 = \langle\langle 1 \rangle\rangle_{almost}(\Phi)$ be the set of almost-sure winning states for player 1 in G . We have $S \setminus W_1 = \langle\langle 2 \rangle\rangle_{pos}(\bar{\Phi})$ and hence any almost-sure winning strategy for player 1 ensures that from W_1 the set $S \setminus W_1$ is not reached with positive probability. Hence we only require to consider strategies σ for player 1 such that for all $w \in W_1^*$ and $s \in W_1$ we have $\text{Supp}(\sigma(w \cdot s)) \subseteq W_1$. Consider a randomized memoryless strategy σ for player 1 such that for a state $s \in W_1$ it chooses uniformly at random all successors in W_1 . Observe that for a state $s \in (S_2 \cup S_\circ) \cap W_1$ we have $E(s) \subseteq W_1$; otherwise s would not have been in W_1 . Consider the MDP $G_\sigma \upharpoonright W_1$. Since it is a player-2 MDP with the Müller objective $\bar{\Phi}$ and randomized memoryless optimal strategies exist in MDPs [3], we fix a memoryless counter-optimal strategy π for player 2 in $G_\sigma \upharpoonright W_1$. Now consider the player-1 MDP $G_\pi \upharpoonright W_1$. Consider a memoryless strategy σ' in $G_\pi \upharpoonright W_1$. We first present an observation: since the strategy σ chooses all successors in W_1 uniformly at random and for all $s \in W_1 \cap S_1$ we have $\text{Supp}(\sigma'(s)) \subseteq \text{Supp}(\sigma(s))$, it follows that for every closed recurrent set U' in the Markov chain $G_{\sigma', \pi} \upharpoonright W_1$ there is a closed recurrent set U in the Markov chain $G_{\sigma, \pi} \upharpoonright W_1$ with $U' \subseteq U$. We now prove that σ is an almost-sure winning strategy by showing that all recurrent set of states U in $G_{\sigma, \pi} \upharpoonright W_1$ is winning for player 1, i.e., $\chi(U) \in \mathcal{F}$. Assume towards contradiction, there is a closed recurrent set U in $G_{\sigma, \pi} \upharpoonright W_1$ with $\chi(U) \notin \mathcal{F}$. Consider the player-1 MDP $G_\pi \upharpoonright W_1$. Since randomized memoryless optimal strategies exist in MDPs [3], we fix a memoryless counter-optimal strategy σ' for player 1. By observation for any closed recurrent set U' in $G_{\sigma', \pi}$ such that $U' \cap U \neq \emptyset$ we have $U' \subseteq U$; and moreover, $\chi(U') \subseteq \chi(U)$ and $\chi(U') \notin \mathcal{F}$, since \mathcal{F} is upward closed and $\chi(U) \notin \mathcal{F}$. It then follows that player 2 wins with probability 1 in from a non-empty set U' (a closed recurrent set $U' \subseteq U$) of states in the Markov chain $G_{\sigma', \pi}$. Since π is a fixed strategy for player 2 and the strategy σ' is counter-optimal for player 1, this contradicts that $U' \subseteq U \subseteq \langle\langle 1 \rangle\rangle_{almost}(\Phi)$. It follows that every closed recurrent set U in $G_{\sigma, \pi} \upharpoonright W_1$ is winning for player 1 and the result follows. ■