The Complexity of Stochastic Rabin and Streett Games [⋆,⋆⋆]

Krishnendu Chatterjee[1], Luca de Alfaro[2], and Thomas A. Henzinger[1,3]

[1] EECS, University of California, Berkeley, USA
[2] CE, University of California, Santa Cruz, USA
[3] EPFL, Switzerland
{c_krish,tah}@eecs.berkeley.edu, luca@soe.ucsc.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 show that for Rabin winning conditions, both problems are in NP. As these problems were known to be NP-hard, it follows that they are NP-complete for Rabin conditions, and dually, coNP-complete for Streett conditions. The proof proceeds by showing that pure memoryless strategies suffice for qualitatively and quantitatively winning stochastic graph games with Rabin conditions. This insight is of interest in its own right, as it implies that controllers for Rabin objectives have simple implementations. We also prove that for every ω-regular condition, optimal winning strategies are no more complex than almost-sure winning strategies.

1 Introduction

A stochastic graph game [5] 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, 14, 16]: a reactive system and its environment represent the two players, whose states and transitions are specified by the states and edges

⋆ This research was supported in part by the ONR grant N00014-02-1-0671, the AFOSR MURI grant F49620-00-1-0327, and the NSF grant CCR-0225610.
⋆⋆ Full proofs are available in [2].

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 [9, 15].

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 [11]. This paper focuses on the complexity of solving $2^1\!/_2$-player graph games with respect to two important normal forms of ω-regular winning conditions: *Rabin conditions* and *Streett conditions* [17]. Rabin and Streett conditions are dual (i.e., complementary), and their practical relevance stems from the fact that their form corresponds to the form of fairness conditions for transition systems.

In the case of 2-player graph games, where no randomization is involved, a fundamental determinacy result 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 [10]. 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. This problem is known to be in NP $\cap$ coNP for parity conditions, to be NP-complete for Rabin conditions [8], and consequently, to be coNP-complete for Streett conditions. The proofs of inclusion in NP rely on the existence of pure (i.e., deterministic) memoryless winning strategies, which act as polynomial witnesses. The existence of pure memoryless winning strategies is also of independent interest, as such strategies can be simply and effectively implemented by a controller. Note that for Streett conditions, winning strategies in general require memory.

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) [7]. 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 that, in $2^1\!/_2$-player graph games, both the qualitative and the quantitative solution problems are NP-complete in the case of Rabin conditions, and coNP-complete in the case of Streett conditions. The NP-hardness for Rabin conditions follows from the NP-hardness of 2-player games with Rabin conditions [8]; we establish the membership in NP. Both questions are known to be in NP $\cap$ coNP for the more restrictive, self-dual case of parity conditions [4, 13, 18], whose exact complexity is an important open problem.

Our proof of membership in NP for stochastic Rabin games relies on establishing the existence of pure memoryless almost-sure winning and optimal strategies. The corresponding result for stochastic parity games has been proved only recently [4, 13, 18], and these proofs rely on the self-duality of parity conditions. For Rabin conditions, a new proof approach is required. First, we show the existence of pure memoryless almost-sure winning strategies in stochastic Rabin games by a reduction from $2\frac{1}{2}$-player games to 2-player. The reduction preserves the ability of player 1 to win with probability 1, but it does not preserve the maximal probability of winning. The proof technique is combinatorial and uses graph-theoretic arguments to account for the fact that Rabin conditions are not closed under complementation. Second, to show the existence of pure memoryless optimal strategies in stochastic Rabin games, we partition the game graph into value classes, each consisting of states where the value of the game is identical. We prove that if the players play according to optimal strategies, then the game leaves every intermediate value class (in which the value is neither 0 nor 1) with probability 1. We then use the qualitative result on almost-sure winning to establish the existence of pure memoryless optimal strategies.

We emphasize that, as mentioned earlier, the existence of pure memoryless strategies is relevant in its own right, as such strategies consist in mappings that associate with each player-1 state a unique successor, without need for randomization or memory; such mappings are easily implemented in controllers. Furthermore, our techniques lead us to a more general result, which states a strong connection between certain qualitative and quantitative games: we show that for every ω-regular winning condition in a $2\frac{1}{2}$-player game graph, if a restricted family of strategies suffices for almost-sure winning, then it suffices also for optimality. Hence future research on $2\frac{1}{2}$-player games with ω-regular conditions can focus on qualitatively (i.e., almost-sure) winning strategies, and our result generalizes these strategies to quantitatively winning (i.e., optimal) strategies.

2 Definitions

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

Game graphs. A *turn-based probabilistic game graph* ($2\frac{1}{2}$-*player game graph*) $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$ consists of a directed graph (S, E), a partition (S_1, S_2, $S_\bigcirc$) of the finite set S of states, and a probabilistic transition function δ: $S_\bigcirc \to \mathcal{D}(S)$, 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 $S_\bigcirc$ are the *probabilistic* states, where the successor state is chosen according to the probabilistic transition function δ. We assume that for $s \in S_{\bigcirc}$ 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 $u \in U \cap S_{\bigcirc}$, 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 $S_{\bigcirc} = \emptyset$. 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 $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ 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 σ: $S^* \cdot S_1 \to \mathcal{D}(S)$ that assigns a probability distribution to all finite sequences $\boldsymbol{w} \in S^* \cdot S_1$ 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 $\boldsymbol{w} \in S^* \cdot S_1$, she chooses the next state according to the probability distribution $\sigma(\boldsymbol{w})$. A strategy must prescribe only available moves, i.e., for all $\boldsymbol{w} \in S^*$, $s \in S_1$, and $t \in S$, if $\sigma(\boldsymbol{w} \cdot s)(t) > 0$, then $(s,t) \in E$. 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 $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ is *feasible* if for every $k \in \mathbb{N}$ the following three conditions hold: (1) if $s_k \in S_{\bigcirc}$, then $(s_k, s_{k+1}) \in E$; (2) if $s_k \in S_1$, then $\sigma(s_0, s_1, \ldots, s_k)(s_{k+1}) > 0$; and (3) if $s_k \in S_2$ then $\pi(s_0, s_1, \ldots, s_k)(s_{k+1}) > 0$. Given two strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a state $s \in S$, we denote by $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ 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 $\mathrm{Pr}_s^{\sigma,\pi}(\mathcal{A})$ 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 $\boldsymbol{w} \in S^*$ and $s \in S_1$, there is a state $t \in S$ such that $\sigma(\boldsymbol{w} \cdot s)(t) = 1$. 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 $\mathtt{M}$ be a set called *memory*. A player-1 strategy can be described as a pair of functions: a *memory-update* function σ_u: $S \times \mathtt{M} \to \mathtt{M}$ and a *next-move* function σ_m: $S_1 \times \mathtt{M} \to \mathcal{D}(S)$. The strategy (σ_u, σ_m) is *finite-memory* if the memory $\mathtt{M}$ is finite. 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 $|\mathtt{M}| = 1$; 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 σ: $S_1 \to \mathcal{D}(S)$. A *pure memoryless strategy* is a pure strategy that is memoryless. A pure memoryless strategy for player 1 can be represented as a function σ: $S_1 \to S$. 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 memoryless player-1 strategy σ, the result G_σ is a player-2 MDP. Similarly, for a player-1 MDP G and a memoryless 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 memoryless 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* $\Phi \subseteq \Omega$ [17]. In this paper we study zero-sum games [9, 15], 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 $\Omega \setminus \Phi$. We consider ω-regular objectives specified in Rabin or Streett normal forms. For a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$, let $\mathrm{Inf}(\omega)$ be the set $\{ s \in S \mid s = s_k \text{ for infinitely many } k \geq 0 \}$ of states that occur infinitely often in ω. We use colors to define objectives independent of game graphs. For a set C of colors, we write $[\![\cdot]\!]$: $C \to 2^S$ for a function that maps each color to a set of states. Inversely, given a set $U \subseteq S$ of states, we write $[U] = \{ c \in C \mid [\![c]\!] \cap U \neq \emptyset \}$ for the set of colors that occur in U. Note that a state can have multiple colors.

A *Rabin objective* is specified as a set $P = \{(e_1, f_1), \ldots, (e_d, f_d)\}$ of pairs of colors $e_i, f_i \in C$. Intuitively, the Rabin condition P requires that for some $1 \leq i \leq d$, all states of color e_i be visited finitely often and some state of color f_i be visited infinitely often. Let $[\![P]\!] = \{(E_1, F_1), \ldots, (E_d, F_d)\}$ be the corresponding set of so-called *Rabin pairs*, where $E_i = [\![e_i]\!]$ and $F_i = [\![f_i]\!]$ for all $1 \leq i \leq d$. Formally, the set of winning plays is $\mathrm{Rabin}(P) = \{ \omega \in \Omega \mid \exists\, 1 \leq i \leq d.\ (\mathrm{Inf}(\omega) \cap E_i = \emptyset \ \wedge\ \mathrm{Inf}(\omega) \cap F_i \neq \emptyset) \}$. Without loss of generality, we require that $\left(\bigcup_{i \in \{1,2,\ldots,d\}} (E_i \cup F_i)\right) = S$. The *parity* (or *Rabin-chain*) objectives are the special case of Rabin objectives such that $E_1 \subset F_1 \subset E_2 \subset F_2 \ldots \subset E_d \subset F_d$. A *Streett objective* is again specified as a set $P = \{(e_1, f_1), \ldots, (e_d, f_d)\}$ of pairs of colors. The Streett condition P requires that for each $1 \leq i \leq d$, if some state of color f_i is visited infinitely often, then some state of color e_i be visited infinitely often. Formally, the set of winning plays is $\mathrm{Streett}(P) = \{ \omega \in \Omega \mid \forall\, 1 \leq i \leq d.\ (\mathrm{Inf}(\omega) \cap E_i \neq \emptyset \ \vee\ \mathrm{Inf}(\omega) \cap F_i = \emptyset) \}$, for the set $[\![P]\!] = \{(E_1, F_1), \ldots, (E_d, F_d)\}$ of so-called *Streett pairs*. Note that the Rabin and Streett objectives are dual; i.e., the complement of a Rabin objective is a Streett objective, and vice versa. Moreover, every parity objective is both a Rabin objective and a Streett objective.

Sure winning, almost-sure 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 $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Phi$. The 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 $\mathrm{Pr}_s^{\sigma,\pi}(\Phi) = 1$. The sure and almost-sure winning strategies for player 2 are defined analogously. Given an objective Φ, the *sure winning set* $\langle\langle 1 \rangle\rangle_{sure}(\Phi)$ for player 1 is the set of states from which player 1 has a sure winning strategy. The *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning strategy. The sure winning set $\langle\langle 2 \rangle\rangle_{sure}(\Omega \setminus \Phi)$ and the almost-sure winning set $\langle\langle 2 \rangle\rangle_{almost}(\Omega \setminus \Phi)$ for player 2 are defined analogously. It follows from the definitions that for all $2\frac{1}{2}$-player game graphs and all objectives Φ, we have $\langle\langle 1 \rangle\rangle_{sure}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\Phi)$. Computing sure and almost-sure winning sets and strategies is referred to as the *qualitative* analysis of $2\frac{1}{2}$-player games [7].

Given ω-regular objectives $\Phi \subseteq \Omega$ for player 1 and $\Omega \setminus \Phi$ for player 2, we define the *value* functions $\langle\langle 1 \rangle\rangle_{val}$ and $\langle\langle 2 \rangle\rangle_{val}$ 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 $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \mathrm{Pr}_s^{\sigma,\pi}(\Omega \setminus \Phi)$. In other words, the value $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$ 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 $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi)$. The optimal strategies for player 2 are defined analogously. Computing values is referred to as the *quantitative* analysis of $2\frac{1}{2}$-player games. The set of states with value 1 is called the *limit-sure winning set* [7]. For $2\frac{1}{2}$-player game graphs with ω-regular objectives the almost-sure and limit-sure winning sets coincide [3].

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 $s \in \langle\langle 1 \rangle\rangle_{sure}(\Phi)$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi \in \Pi$, we have $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Phi$. Similarly, the family $\Sigma^{\mathcal{C}}$ *suffices* with respect to the objective Φ on the class $\mathcal{G}$ of game graphs for *almost-sure winning* if for every game graph $G \in \mathcal{G}$ and state $s \in \langle\langle 1 \rangle\rangle_{almost}(\Phi)$, there is a player-1 strategy $\sigma \in \Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi \in \Pi$, we have $\mathrm{Pr}_s^{\sigma,\pi}(\Phi) = 1$; and for *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 $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi)$.

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 ω-regular objectives.

Theorem 1 (Qualitative determinacy [8, 10]). *For all 2-player game graphs and Rabin or Streett objectives Φ, we have $\langle\langle 1\rangle\rangle_{sure}(\Phi) \cap \langle\langle 2\rangle\rangle_{sure}(\Omega \setminus \Phi) = \emptyset$ and $\langle\langle 1\rangle\rangle_{sure}(\Phi) \cup \langle\langle 2\rangle\rangle_{sure}(\Omega \setminus \Phi) = S$. Moreover, on 2-player game graphs, the family of pure memoryless strategies suffices for sure winning with respect to Rabin objectives, and the family of pure finite-memory strategies suffices for sure winning with respect to Streett objectives.*

Theorem 2 (Quantitative determinacy [12]). *For all $2\,^1\!/_2$-player game graphs, all Rabin or Streett objectives Φ, and all states s, we have $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) + \langle\langle 2\rangle\rangle_{val}(\Omega \setminus \Phi)(s) = 1$.*

3 Qualitative Analysis

We show that the pure memoryless strategies suffice for almost-sure winning with respect to Rabin objectives on $2\,^1\!/_2$-player game graphs. The result is achieved by a reduction to 2-player Rabin games. The reduction also allows us to apply algorithms for solving 2-player Rabin games to the qualitative analysis of $2\,^1\!/_2$-player Rabin games. Furthermore, in the next section, we will use the existence of pure memoryless almost-sure winning strategies to prove the existence of pure memoryless optimal strategies.

End components of MDPs. We review some facts about *end components* [6] which are needed for the further development of the paper. We consider player-1 MDPs and hence strategies for player 1. Let $G = ((S, E), (S_1, S_2, S_{\bigcirc}), \delta)$ with $S_2 = \emptyset$ be a $1\,^1\!/_2$-player game graph.

Definition 1 (End components). *A set $U \subseteq S$ of states is an* end component *if U is δ-closed and the subgame graph $G \upharpoonright U$ is strongly connected.*

We denote by $\mathcal{E} \subseteq 2^S$ the set of all end-components of G. The next lemma states that, under every strategy (memoryless or not), with probability 1 the set of states visited infinitely often along a play is an end component. This lemma allows us to derive conclusions on the (infinite) set of plays in an MDP by analyzing the (finite) set of end components in the MDP. In particular, the lemma implies that to show that a set $\{(E_1, F_1), \ldots, (E_d, F_d)\}$ of Rabin pairs is satisfied with probability 1, it suffices to show that for each reachable end component U, there exists an $1 \leq i \leq d$ such that $U \cap E_i = \emptyset$ and $U \cap F_i \neq \emptyset$. To state the lemma, for $s \in S$ and $U \subseteq S$, we define $\Omega_s^U = \{\, \omega \in \Omega_s \mid \mathrm{Inf}(\omega) = U \,\}$.

Lemma 1. [6] *For all states $s \in S$ and strategies $\sigma \in \Sigma$, $\Pr_s^\sigma(\bigcup_{U \in \mathcal{E}} \Omega_s^U) = 1$.*

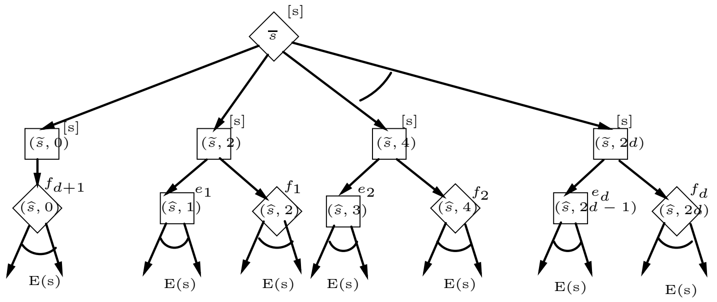

Fig. 1. Gadget for the reduction of $2{}^1\!/\!{}_2$-player Rabin games to 2-player Rabin games.

Reduction. Given a $2{}^1\!/\!{}_2$-player game graph $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$, a set $C = \{e_1, f_1, \ldots, e_d, f_d\}$ of colors, and a color map $[\cdot]\colon S \to 2^C \setminus \emptyset$, we construct a 2-player game graph $\overline{G} = ((\overline{S}, \overline{E}), (\overline{S}_1, \overline{S}_2), \overline{\delta})$ together with a color map $[\cdot]\colon \overline{S} \to 2^{\overline{C}} \setminus \emptyset$ for the extended color set $\overline{C} = C \cup \{e_{d+1}, f_{d+1}\}$. The construction is specified as follows. For every nonprobabilistic state $s \in S_1 \cup S_2$, there is a corresponding state $\overline{s} \in \overline{S}$ such that (1) $\overline{s} \in \overline{S}_1$ iff $s \in S_1$, and (2) $[\overline{s}] = [s]$, and (3) $(\overline{s}, \overline{t}) \in \overline{E}$ iff $(s, t) \in E$. Every probabilistic state $s \in S_\bigcirc$ is replaced by the gadget shown in Figure 1. In the figure, diamond-shaped states are player-2 states (in $\overline{S}_2$), and square-shaped states are player-1 states (in $\overline{S}_1$). From the state $\overline{s}$ with $[\overline{s}] = [s]$, the players play the following 3-step game in $\overline{G}$. First, in state $\overline{s}$ player 2 chooses a successor $(\widetilde{s}, 2k)$, for $k \in \{0, 1, \ldots, d\}$. For every state $(\widetilde{s}, 2k)$, we have $[(\widetilde{s}, 2k)] = [s]$. For $k > 1$, in state $(\widetilde{s}, 2k)$ player 1 chooses from two successors: state $(\widehat{s}, 2k-1)$ with $[(\widehat{s}, 2k-1)] = e_k$, or state $(\widehat{s}, 2k)$ with $[(\widehat{s}, 2k)] = f_k$. The state $(\widetilde{s}, 0)$ has only one successor $(\widehat{s}, 0)$, with $[(\widehat{s}, 0)] = f_{d+1}$. Note that no state in $\overline{S}$ is labeled by the new color e_{d+1}, that is, $[\![e_{d+1}]\!] = \emptyset$. Finally, in each state $(\widehat{s}, j)$ the choice is between all states $\overline{t}$ such that $(s, t) \in E$, and it belongs to player 1 if k is odd, and to player 2 if k is even.

We consider the $2{}^1\!/\!{}_2$-player game played on the graph G with the Rabin condition $P = \{(e_1, f_1), \ldots, (e_d, f_d)\}$ for player 1. Let $\overline{U}_1$ and $\overline{U}_2$ be the sure winning sets for players 1 and 2, respectively, in the constructed 2-player game graph $\overline{G}$ with the modified Rabin condition $\overline{P} = \{(e_1, f_1), \ldots, (e_{d+1}, f_{d+1})\}$ for player 1. Define the sets U_1 and U_2 in the original $2{}^1\!/\!{}_2$-player game graph G by $U_1 = \{\, s \in S \mid \overline{s} \in \overline{U}_1 \,\}$ and $U_2 = \{\, s \in S \mid \overline{s} \in \overline{U}_2 \,\}$. From the determinacy of 2-player Rabin games (Theorem 1), it follows that $\overline{U}_1 = \overline{S} \setminus \overline{U}_2$, and hence $U_1 = S \setminus U_2$.

Lemma 2. *In the $2{}^1\!/\!{}_2$-player game graph G with the Rabin condition P for player 1, there exists a pure memoryless strategy σ for player 1 such that for all player-2 strategies π and all states $s \in U_1$, we have $\Pr_s^{\sigma,\pi}(\mathrm{Rabin}(P)) = 1$.*

Proof. We define a pure memoryless strategy σ for player 1 in the game G from a strategy $\overline{\sigma}$ in the game $\overline{G}$ as follows: for all states $s \in S_1$, if $\overline{\sigma}(\overline{s}) = \overline{t}$, then set $\sigma(s) = t$. Consider a pure memoryless sure winning strategy $\overline{\sigma}$ in the game $\overline{G}$ from every state $\overline{s} \in \overline{U}_1$. Our goal is to establish that σ is an almost-sure winning strategy from every state in U_1.

For the Rabin objective Rabin(P), let the set Rabin pairs be $[\![P]\!] = \{ (E_1, F_1), (E_2, F_2), \ldots, (E_d, F_d) \}$. A strongly connected component (s.c.c.) W in a graph G_1 is winning for player 1, if there exists $i \in \{ 1, 2, \ldots, d \}$ such that $W \cap F_i \neq \emptyset$ and $W \cap E_i = \emptyset$; otherwise W is winning for player 2. If G_1 is a MDP, then an end component W in G_1 is winning for player 1, if there exists $i \in \{ 1, 2, \ldots, d \}$ such that $W \cap F_i \neq \emptyset$ and $W \cap E_i = \emptyset$; otherwise W is winning for player 2.

We prove that every end component in the player-2 MDP $(G \upharpoonright U_1)_\sigma$ is winning for player 1. It would follow from Lemma 1 that σ is an almost-sure winning strategy. We argue that if there is an end component W in $(G \upharpoonright U_1)_\sigma$ that is winning for player 2, then we can construct an s.c.c. in the subgraph $(\overline{G} \upharpoonright \overline{U}_1)_{\overline{\sigma}}$ that is winning for player 2, which is impossible because $\overline{\sigma}$ is a sure winning strategy for player 1 from the set $\overline{U}_1$ in the 2-player Rabin game $\overline{G}$. Let W be an end component in $(G \upharpoonright U_1)_\sigma$ that is winning for player 2. We denote by $\overline{W}$ the set of states in the gadget of states in W. Hence for all $i \in \{ 1, 2, \ldots, d \}$ we have if $F_i \cap W \neq \emptyset$, then $W \cap E_i \neq \emptyset$. Let us define the set $I = \{ i_1, i_2, \ldots, i_j \}$ such that $E_{i_k} \cap W \neq \emptyset$. Thus for all $i \in (\{ 1, 2, \ldots, d \} \setminus I)$ we have $F_i \cap W = \emptyset$. Note that $I \neq \emptyset$, as every state has at least one color. We now construct a sub-game in $\overline{G}_{\overline{\sigma}}$ as follows:

1. For a state $\overline{s} \in \overline{W} \cap \overline{S}_2$ keep all the edges $(\overline{s}, \overline{t})$ such that $\overline{t} \in \overline{W}$.
2. For a state $\overline{s} \in \overline{W} \cap \overline{S}_{\bigcirc}$ the sub-game is defined as follows:
   - At state $\overline{s}$ choose the edges to state $(\widetilde{s}, 2i)$ such that $i \in I$.
   - For a state $s \in W$, let $dis(s, W \cap E_i)$ denote the shortest distance (BFS distance) from s to $W \cap E_i$ in the graph of $(G \upharpoonright W)_\sigma$. At state $(\widehat{s}, 2i)$, which is a player 2 state, player 2 chooses a successor $\widehat{s}_1$ such that $dis(s_1, W \cap E_i) < dis(s, W \cap E_i)$ (i.e., shorten distance to the set $\overline{W} \cap E_i$ in $\overline{G}$).

We now prove that every terminal s.c.c. is winning for player 2 in the subgame thus constructed in $(\overline{G} \upharpoonright \overline{W})_{\overline{\sigma}}$, where $\overline{W}$ is the set of states in the gadget of states in W. Consider any arbitrary terminal s.c.c. $\overline{Y}$ in the subgame constructed in $(\overline{G} \upharpoonright \overline{W})_{\overline{\sigma}}$. It follows from the construction that for every $i \in (\{ 1, 2, \ldots, d \} \setminus I)$, we have $F_i \cap \overline{Y} = \emptyset$. Suppose for a $i \in I$ we have $F_i \cap \overline{Y} \neq \emptyset$, we show that $E_i \cap \overline{Y} \neq \emptyset$. There are two cases:

1. If there is at least one state $(\widetilde{s}, 2i)$ such that the strategy $\overline{\sigma}$ chooses the successor $(\widehat{s}, 2i-1)$, then $E_i \cap \overline{Y} \neq \emptyset$, since $[(\widetilde{s}, 2i-1)] = e_i$.
2. Else for every state $(\widetilde{s}, 2i)$ the strategy for player 1 chooses the successor $(\widehat{s}, 2i)$. At state $(\widehat{s}, 2i)$, which is a player 2 state, player 2 chooses a successor $\widehat{s}_1$ that shortens distance to the set $\overline{Y} \cap E_i$. Hence the terminal s.c.c. $\overline{Y}$ must contain a state $\overline{s}$ such that $[\overline{s}] = e_i$. Hence $E_i \cap \overline{Y} \neq \emptyset$.

We argue that for every probabilistic state $s \in S_{\bigcirc} \cap U_1$, all of its successors are in U_1. Otherwise, player 2 in the state $\overline{s}$ of the game $\overline{G}$ can choose the successor $(\widetilde{s}, 0)$ and then a successor to its winning set $\overline{U}_2$, which contradicts the assumption that the strategy $\overline{\sigma}$ is a sure winning strategy for player 1 in the game $\overline{G}$ from $\overline{U}_1$. It follows from Lemma 1 that for all strategies π, for all states $s \in U_1$, with probability 1 the set of states visited infinitely often along the play $\omega_s^{\sigma,\pi}$ is an end component in U_1. Since every end component in $(G \upharpoonright U_1)_\sigma$ is winning for player 1 the strategy σ is an almost-sure winning strategy for player 1 from U_1. ∎

Lemma 3. *In the $2^1\!/_2$-player game graph G with the Rabin condition P for player 1, there exists a finite-memory strategy π for player 2 such that for all player-1 strategies σ and all states $s \in U_2$, we have $\Pr_s^{\sigma,\pi}(\Omega \setminus \mathrm{Rabin}(P)) > 0$.*

From Lemma 2, it follows that $U_1 \subseteq \langle\!\langle 1 \rangle\!\rangle_{almost}\mathrm{Rabin}(P)$. From Lemma 3, it follows that $\langle\!\langle 1 \rangle\!\rangle_{almost}\mathrm{Rabin}(P) \subseteq U_1$. Therefore $U_1 = \langle\!\langle 1 \rangle\!\rangle_{almost}\mathrm{Rabin}(P)$. The proof of Lemma 2 also establishes the existence of pure memoryless almost-sure winning strategies for Rabin objectives.

Theorem 3. *The family of pure memoryless strategies suffices for almost-sure winning with respect to Rabin objectives on $2^1\!/_2$-player game graphs.*

4 Quantitative Analysis

We extend sufficiency results for families of strategies from almost-sure winning to optimality with respect to all ω-regular objectives. In the following, we fix a $2^1\!/_2$-player game graph G. Given an ω-regular objective Φ, for every real $r \in \mathbb{R}$ the *value class* with value r is $\mathrm{VC}(r) = \{s \in S \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s) = r\}$. Proposition 1 states that there exist optimal strategies for player 1 such that they never choose an edge to a lower value class.

Proposition 1. *For all ω-regular objectives Φ, there exists an optimal strategy σ for player 1 such that for all $\boldsymbol{w} \in S^*$, $s \in S_1$, and $t \in S$, if $\langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(t) < \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s)$, then $\sigma(\boldsymbol{w} \cdot s)(t) = 0$.*

Definition 2 (Boundary probabilistic states). *Given an ω-regular objective Φ, a probabilistic state $s \in S_\bigcirc$ is a* boundary probabilistic state *if there exists a successor $t \in E(s)$ such that $\langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(t) \neq \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s)$. Observe that for every boundary probabilistic state s, there exist $t_1, t_2 \in E(s)$ such that $\langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(t_1) < \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s)$ and $\langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(t_2) > \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s)$.*

Lemma 4. *Consider a $2^1\!/_2$-player game G with an ω-regular objective Φ. Given a value class $\mathrm{VC}(r)$ with $0 < r < 1$, let $B(r)$ be the set of boundary probabilistic states in the value class $\mathrm{VC}(r)$. Convert each state in $B(r)$ into a sink state that is winning for player 1. Let the new game be G'. Then player 1 wins almost-surely from all states in the subgame with game graph $G' \upharpoonright \mathrm{VC}(r)$ and objective Φ.*

Proof. Assume that player 1 does not win almost-surely from every state in $G' \upharpoonright \mathrm{VC}(r)$. Then there exists a state where player 2 wins with positive bounded probability. It follows from Corollary 1 of [7] that there exist a non-empty set $U \subseteq \mathrm{VC}(r)$ such that that player 2 wins almost-surely from U in $G' \upharpoonright \mathrm{VC}(r)$. Consider an optimal strategy σ that never chooses an edge with positive probability to a lower value class (such a strategy exists from Proposition 1). Since player 2 wins almost-surely from U it follows that for every state $s \in U \cap S_1$, for every successor t of s in $\mathrm{VC}(r)$ we have $t \in U$. It follows that every move of the strategy σ exists in U. Hence player 2 wins almost-surely from U against σ. This is a contradiction to the assumption that $r > 0$ and that σ is an optimal strategy. ∎

Definition 3 (Qualitatively optimal strategies). *A strategy σ is* qualitatively optimal *for player 1, for an ω-regular objective Φ, if the following conditions hold: (a) for every state $s \in \langle\langle 1 \rangle\rangle_{almost}(\Phi)$, the strategy σ is almost-sure winning, and (b) for every state $s \in \mathrm{VC}(r)$ such that $0 < r < 1$, there is a constant $c > 0$ such that $\inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi) \geq c$.*

Lemma 4 shows that in every value class, if the boundary probabilistic states are assumed to be winning for player 1, then player 1 wins almost-surely. We call such an almost-sure winning strategy a *conditional* almost-sure winning strategy. We compose conditional almost-sure winning strategies in value classes to obtain an optimal strategy. If a strategy σ is conditional almost-sure winning, it follows that for all player-2 strategies π that are optimal against σ, the play $\omega_s^{\sigma,\pi}$ reaches the boundary probabilistic states with positive probability, for $s \in \mathrm{VC}(r)$ and $r > 0$. From every boundary probabilistic state the game proceeds to a higher value class with positive probability. An induction on the number of value classes yields Lemma 5.

Lemma 5. *For every ω-regular objective Φ, if a player-1 strategy σ is almost-sure winning from every state $s \in \langle\langle 1 \rangle\rangle_{almost}(\Phi)$, and is conditionally almost-sure winning from every state $s \notin \langle\langle 2 \rangle\rangle_{almost}(\Omega \setminus \Phi)$, then σ is qualitatively optimal for Φ.*

Definition 4 (Locally optimal strategies). *A strategy σ is* locally optimal *for player 1, for an ω-regular objective Φ, if for all $\boldsymbol{w} \in S^*$, $s \in S_1$, and $t \in S$, if $\langle\langle 1 \rangle\rangle_{val}(\Phi)(t) < \langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$, then $\sigma(\boldsymbol{w} \cdot s)(t) = 0$.*

Note that by definition, a conditional almost-sure winning strategy is locally optimal. The following Lemma generalizes Lemma 5.3 of [4]. Theorem 4 follows from Lemma 6. Since pure memoryless strategies suffice for almost-sure winning with respect to Rabin objectives on $2^1\!/_2$-player game graphs (Theorem 3), Theorem 5 is immediate from Theorem 4.

Lemma 6. *Consider a $2^1\!/_2$-player game G with an ω-regular objective Φ for player 1. Let σ be a finite-memory strategy such that σ is both qualitatively optimal and locally optimal for Φ. Then σ is an optimal strategy for Φ from all states of G.*

Theorem 4. *If a family $\Sigma^{\mathcal{C}}$ of strategies suffices for almost-sure winning with respect to an ω-regular objective Φ on $2^1\!/_2$-player game graphs, then $\Sigma^{\mathcal{C}}$ suffices for optimality with respect to Φ on $2^1\!/_2$-player game graphs.*

Theorem 5. *The family of pure memoryless strategy suffices for optimality with respect to Rabin objectives on $2^1\!/_2$-player game graphs.*

The existence of pure memoryless optimal strategies for $2^1\!/_2$-player game graphs with Rabin objectives, and of polynomial-time algorithms for computing the values of MDPs with Streett objectives [2], establishes that the $2^1\!/_2$-player games with Rabin objectives can be decided (qualitatively and quantitatively) in NP. The NP-hardness follows from the hardness of 2-player Rabin games.

Theorem 6. *Given a $2^1\!/_2$-player game graph G, an objective Φ for player 1, a state s of G, and a rational r, the complexity of determining whether $\langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s) \geq r$ is as follows: NP-complete if Φ is a Rabin objective; coNP-complete if Φ is a Streett objective; and in NP $\cap$ coNP if Φ is a parity objective.*

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, L. de Alfaro, and T.A. Henzinger. The complexity of stochastic Rabin and Streett games. Technical Report UCB/CSD-3-1355, UC Berkeley, 2004.
3. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
4. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA'04*, pages 114–123. SIAM, 2004.
5. A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
6. L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD Thesis, Stanford University, 1997.
7. L. de Alfaro and T.A. Henzinger. Concurrent ω-regular games. In *LICS'00*, pages 141–154. IEEE Computer Society, 2000.
8. E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *FOCS'88*, pages 328–337. IEEE Computer Society, 1988.
9. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
10. Y. Gurevich and L. Harrington. Trees, automata, and games. In *STOC'82*, pages 60–65. ACM, 1982.
11. Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1992.
12. D.A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63:1565–1581, 1998.
13. 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.
14. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM, 1989.
15. 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.
16. 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.
17. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3 (*Beyond Words*), pages 389–455. Springer, 1997.
18. W. Zielonka. Perfect-information stochastic parity games. In *FoSSaCS'04*, volume 2987 of *LNCS*, pages 499–513. Springer, 2004.