


The Complexity of Stochastic Rabin and Streett Games

Krishnendu Chatterjee

Luca deAlfaro

Thomas A. Henzinger



Report No. UCB/CSD-4-1355

November 2004

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

The Complexity of Stochastic Rabin and Streett Games ^{*}

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

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

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

[‡] Computer and Communication Sciences, EPFL, Switzerland
{c_krish,tah}@eecs.berkeley.edu, luca@soe.ucsc.edu

November 2004

Abstract

The foundation of modeling and synthesizing reactive processes is the theory of graph games with ω -regular winning conditions. 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 solution problems for stochastic graph games: a *qualitative* problem, calling for the computation of the set of states from which a player can win with probability 1 (*almost-sure winning*), and a *quantitative* problem, calling for the computation of 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 fact was an open problem, and it is of interest in its own right, as it implies that controllers for Rabin objectives have simple implementations. We also prove that for any ω -regular objective optimal winning strategies are no more complex than almost-sure winning strategies.

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

1 Introduction

A stochastic graph game is played on a directed graph with three kinds of states: player-1 states, player-2 states, and probabilistic states. At player-1 states, the first player chooses a successor state; at player-2 states, the second player 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*, and 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, 19, 21]: 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 *stochastic reactive* processes [20, 13].

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 [15]. This paper focuses on the complexity of solving 2^{1/2}-player graph games with respect to two important normal forms of ω -regular conditions: *Rabin* conditions and *Streett* conditions [23]. 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 particular, no blow-up in the system representation is required when encoding fairness as a Streett condition, or dually, in the antecedent of a temporal specification, as a Rabin condition [15].

In the case of 2-player graph games, where no randomization is involved, a fundamental determinacy result ensures that, given an ω -regular (or indeed Borel) 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 never holds [17]. The problem of solving 2-player graph games consists thus in finding the set of *winning states*, from which player 1 can ensure that the condition holds. This problem is known to be in $\text{NP} \cap \text{coNP}$ for parity conditions [11], to be NP-complete for Rabin conditions [12, 11, 23], 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 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.

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: a *qualitative* one, which asks for the computation of the set of states from which player 1 can ensure winning with probability 1, and a *quantitative* one, which asks for the computation of the maximal probability with which player 1 can ensure winning from each state (also called the *value* of the game at a state) [9, 8]. 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 NP-hardness of 2-player games with Rabin conditions [12, 23]; we establish the membership in NP. Both questions have been known to be in $\text{NP} \cap \text{coNP}$ for the more restrictive, self-dual parity conditions [18, 4, 24], 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 [18, 4, 24]; the proofs rely heavily 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; the proof is based on a reduction from $2^{1/2}$ -player games to 2-player games that preserves the ability of player 1 to win with probability 1 (but not, obviously, the maximal probability of winning). The proof technique is different from the techniques for parity games [3] that relies on the notion of *ranking functions* and self-duality of parity conditions. The present proof technique is combinatorial and uses graph theoretic arguments to take care of the fact that Rabin objectives are not closed under complementation. Our reduction establishes the membership in NP of the qualitative solution problem for stochastic Rabin games. To show the existence of pure memoryless *optimal* strategies,

we partition the game graph into value classes, each consisting of states where the value of the game is identical. We show 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 can then leverage the results on almost-winning to show the existence of pure memoryless optimal strategies, and establish the membership in NP also for the quantitative solution problem for stochastic Rabin games. The coNP-completeness of stochastic Streett games follows by duality.

We emphasize that, as mentioned earlier, the existence of pure memoryless strategies is relevant in its own right, as such strategies consist in mappings associating with each player-1 state a unique successor, without need for randomization or memory; such mappings are easily implemented in controllers. The result that a pure memoryless strategy suffices for winning with probability 1 and for optimality in every stochastic Rabin game is far from obvious; recall that Streett conditions in general require memory even in the simpler case of non-stochastic (i.e., 2-player) graph games. Furthermore, our techniques lead us to a far more general result, that states a strong connection between the qualitative and quantitative problems: we show that for any ω -regular objective in a $2^{1/2}$ -player game graph, if a family of strategies suffices for almost-sure winning, it also suffices for optimality. Hence future research about $2^{1/2}$ -player games with ω -regular objectives can focus on qualitative winning strategies, and our result generalizes qualitative winning strategies to quantitative winning strategies.

2 Preliminaries

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).

Probability distribution. For a countable set A , a probability distribution on the set 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 the set A by $\mathcal{D}(A)$.

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 set of states S , 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 states 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 in (S, E) every state has at least one outgoing edge, and we write $t \in E(s)$ for $(s, t) \in E$. For a state s we write $E(s)$ to denote $\{t \in S \mid (s, t) \in E\}$. We denote by n the size of the state space, i.e., $n = |S|$, and by m the number of edges, i.e., $m = |E|$.

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 every state $s \in S$ we write Ω_s for the set of plays that start from the state s .

A set $U \subseteq S$ of states is called *δ -closed* if for every $u \in U \cap S_{\circ}$, we have that $(u, t) \in E$ implies $t \in U$; it is called *δ -live* if for every 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 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*. A game graph which is both deterministic and an MDP is called a *transition system* (*1-player game graph*): a player-1 transition system has only player-1 states; a player-2 transition system has only player-2 states.

Strategies. A *strategy* for player 1 is a function $\sigma: S^* \cdot S_1 \rightarrow \mathcal{D}(S)$ that assigns a probability distribution to every finite sequence $\vec{w} \in S^* \cdot S_1$ of states, which represents the history of the play so far. Player 1 follows the strategy σ if in each move, given that the current history of the play is $\vec{w} \in S^* \cdot S_1$, she chooses the next state according to the probability distribution $\sigma(\vec{w})$. A strategy must prescribe only available moves, i.e., for all $\vec{w} \in S^*$, $s \in S_1$, and $t \in S$, if $\sigma(\vec{w} \cdot s)(t) > 0$, 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. Note that for player-1 MDPs the set Π is a singleton, i.e., player 2 has only a single *trivial* strategy.

Pure strategies. 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 $\vec{w} \in S^*$ and $s \in S_1$, there is a state $t \in S$ such that $\sigma(\vec{w} \cdot s)(t) = 1$. The pure strategies for player 2 are defined analogously. We denote by Σ^P and Π^P the sets of pure strategies

for player 1 and player 2, respectively. A strategy that is not necessarily pure is called *randomized*.

Finite memory and memoryless strategies. Let \mathbf{M} be a set called *memory*. A strategy with memory can be described as a pair of functions: (a) *memory update* function $\sigma_u : S \times \mathbf{M} \rightarrow \mathbf{M}$, (b) *next move* function $\sigma_m : S_1 \times \mathbf{M} \rightarrow \mathcal{D}(S)$. A strategy is *finite-memory* if the memory \mathbf{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, $\Sigma^{PF} = \Sigma^P \cap \Sigma^F$. A strategy is *memoryless* if $|\mathbf{M}| = 1$: hence, the next move does not depend on the history but only on the current state. A memoryless strategy for player 1 can be represented as function $\sigma : S_1 \rightarrow \mathcal{D}(S)$ such that for all $s \in S_1$ and $t \in S$, if $\sigma(s)(t) > 0$, then $(s, t) \in E$. A *pure memoryless strategy* is a pure strategy that is memoryless. A pure memoryless strategy for player 1 can be represented as a function $\sigma : S_1 \rightarrow S$ such that $(s, \sigma(s)) \in E$ for all $s \in S_1$. 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 for player 2.

Given a strategy $\sigma \in \Sigma$ for player 1, we write G_σ for the game played on the graph G under the constraint that player 1 follows the strategy σ . The corresponding definition for a player-2 strategy is analogous. 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 given memoryless strategies σ and π , the result $G_{\sigma, \pi}$ is a Markov chain. Given a game graph G and a finite memory strategy σ for player 1 with memory \mathbf{M} , the strategy σ can be interpreted as a memoryless strategy σ_m in the usual synchronous product game graph G with the memory \mathbf{M} , i.e., $G \times \mathbf{M}$. Analogous observations hold for player 2 strategies π . These observations will be useful in the analysis of $2^{1/2}$ -player games.

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* $\mathcal{A} \subseteq \Omega_s$ 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* in a $2^{1/2}$ -player game graph 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 strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a state s , we denote by $\text{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ the set of feasible plays

that start from s given strategies σ and π . For a state $s \in S$ and an event \mathcal{A} , we write $\Pr_s^{\sigma, \pi}(\mathcal{A})$ 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.

Objectives. An *objective* for a player consists in an ω -regular set of *winning plays* $\Phi \subseteq \Omega$ for each player [22]. In this paper we study only zero-sum games [20, 13], where the objectives of the two players are complementary. In other words, it is implicit that if the objective of one player is Φ , then the objective of the other player is $\Omega \setminus \Phi$. Given a game graph G and an objective $\Phi \subseteq \Omega$, we write (G, Φ) for the game played on the graph G with the objective Φ for player 1.

In this paper we consider ω -regular objectives specified as Rabin and Streett objectives. For a play $\omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega$, we define $\text{Inf}(\omega) = \{s \in S \mid s_k = s \text{ for infinitely many } k \geq 0\}$ to be the set of states that occur infinitely often in ω . We use colors to define objectives independent of game graphs. For a set C of colors, we write $\llbracket \cdot \rrbracket: C \rightarrow 2^S$ for a function that maps each color to a set of states. Inversely, given a set $U \subseteq S$ of states, we write $\llbracket U \rrbracket = \{c \in C \mid \llbracket c \rrbracket \cap U \neq \emptyset\}$ for the set of colors that occur in U . Note that a state can have multiple colors.

1. *Reachability and safety objectives.* Given a color c , the reachability objective requires that some state of color c be visited. Let $T = \llbracket c \rrbracket$ be the set of so-called *target* states. Formally, we write $\text{Reach}(T) = \{\omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0\}$ for the set of winning plays. Given c , the safety objective requires that only states of color c be visited. Let $F = \llbracket c \rrbracket$ be the set of so-called *safe* states. Formally, the set of winning plays is $\text{Safe}(F) = \{\omega = \langle s_0, s_1, s_2, \dots \rangle \in \Omega \mid s_k \in F \text{ for all } k \geq 0\}$.
2. *Rabin, parity, and Streett objectives.* Given a set $P = \{(e_1, f_1), \dots, (e_d, f_d)\}$ of pairs of colors, the Rabin objective 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 $\hat{P} = \{(E_1, F_1), \dots, (E_d, F_d)\}$ be the corresponding set of so-called *Rabin pairs*, where $E_i = \llbracket e_i \rrbracket$ and $F_i = \llbracket f_i \rrbracket$ for all $1 \leq i \leq d$. Formally, the set of winning plays is $\text{Rabin}(\hat{P}) = \{\omega \in \Omega \mid \exists 1 \leq i \leq d. (\text{Inf}(\omega) \cap E_i = \emptyset \wedge \text{Inf}(\omega) \cap F_i \neq \emptyset)\}$. Without loss of generality, we require that $(\bigcup_{i \in \{1, 2, \dots, d\}} (E_i \cup F_i)) = S$. The *parity* (or *Rabin-chain*) objectives are the special case of Rabin objectives

where $E_1 \subset F_1 \subset E_2 \subset F_2 \cdots \subset E_d \subset F_d$. The Rabin-chain objective can be represented as a parity objective defined as follows: define a priority function p that labels each state in $E_i \setminus F_{i-1}$ by a priority $2i - 1$ and each state in $F_i \setminus E_i$ by a priority $2i$. The parity objective requires that minimum priority state that is visited infinitely often is even. Formally, the set of winning plays is $\text{Parity}(p) = \{ \omega \in \Omega \mid \min(p(\text{Inf}(\omega))) \text{ is even} \}$. Given P , the Streett objective 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 is visited infinitely often. Formally, for the set $\hat{P} = \{(E_1, F_1), \dots, (E_d, F_d)\}$ of so-called *Streett pairs*, the set of winning plays is $\text{Streett}(\hat{P}) = \{ \omega \in \Omega \mid \forall 1 \leq i \leq d. (\text{Inf}(\omega) \cap E_i \neq \emptyset \vee \text{Inf}(\omega) \cap F_i = \emptyset) \}$. Note that the Rabin and Streett objectives are dual. Moreover, every parity objective is both a Rabin objective and a Streett objective. Hence, parity objectives are closed under complementation.

We commonly use terminology like the following: a $2^{1/2}$ -player Rabin game $(G, \text{Rabin}(\hat{P}))$ consists of a $2^{1/2}$ -player game graph G and a Rabin objective for player 1.

Values of a game. 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 follows:

$$\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \text{Pr}_s^{\sigma, \pi}(\Phi) \qquad \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \text{Pr}_s^{\sigma, \pi}(\Omega \setminus \Phi)$$

A strategy σ for player 1 is *optimal* from state s for objective Φ if

$$\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \inf_{\pi \in \Pi} \text{Pr}_s^{\sigma, \pi}(\Phi).$$

The optimal strategies for player 2 are defined analogously. The quantitative determinacy of $2^{1/2}$ -player games with Rabin objectives follows from the result of Martin [16].

Theorem 1 (Quantitative determinacy [16]) *For all $2^{1/2}$ -player game graphs, all Rabin objectives Φ , and all states s ,*

$$\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) + \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) = 1.$$

Sure, almost-sure and limit-sure winning strategies. Given an objective Φ , a strategy σ is a *sure winning strategy* for player 1 from a state

s if for every strategy π of player 2 we have $\text{Outcome}(s, \sigma, \pi) \subseteq \Phi$. A strategy σ is an *almost-sure winning strategy* for player 1 from a state s for the objective Φ if for every strategy π of player 2 we have $\Pr_s^{\sigma, \pi}(\Phi) = 1$. A family of strategies Σ^C are *limit-sure winning* for player 1 from a state s if $\sup_{\sigma \in \Sigma^C} \inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi) = 1$. The sure, almost-sure and limit-sure 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. The *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning strategy. The *limit-sure winning set* $\langle\langle 1 \rangle\rangle_{\text{limit}}(\Phi)$ for player 1 is the set of states from which player 1 has limit-sure winning strategies. 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 limit-sure winning set $\langle\langle 2 \rangle\rangle_{\text{limit}}(\Omega \setminus \Phi)$ 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 $\langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{\text{limit}}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi) \subseteq \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi) \subseteq \langle\langle 2 \rangle\rangle_{\text{limit}}(\Omega \setminus \Phi)$. Computing sure winning, almost-sure winning and limit-sure winning sets and strategies is referred to as the *qualitative analysis* of $2^{1/2}$ -player games [8]. The following result is the classical determinacy result for 2-player deterministic games.

Theorem 2 (Qualitative determinacy [17]) *For all 2-player game graphs and all Rabin objectives Φ , we have*

$$\begin{aligned} \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \cap \langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi) &= \emptyset; & \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi) &= \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi); \\ \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi) \cup \langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi) &= S; & \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi) &= \langle\langle 2 \rangle\rangle_{\text{sure}}(\Omega \setminus \Phi). \end{aligned}$$

Sufficiency of a family of strategies. Let $\mathcal{C} \in \{P, M, F, PM, PF\}$ and consider the family Σ^C of special strategies for player 1. We say that the family Σ^C *suffices* with respect to an objective Φ on a class \mathcal{G} of game graphs for

- *sure winning* if for every game graph $G \in \mathcal{G}$, for every $s \in \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi)$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that for every player-2 strategy $\pi \in \Pi$ we have $\text{Outcome}(s, \sigma, \pi) \subseteq \Phi$;
- *almost-sure winning* if for every game graph $G \in \mathcal{G}$, for every state $s \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that for every player-2 strategy $\pi \in \Pi$ we have $\Pr_s^{\sigma, \pi}(\Phi) = 1$;
- *limit-sure winning* if for every game graph $G \in \mathcal{G}$, for every state $s \in \langle\langle 1 \rangle\rangle_{\text{limit}}(\Phi)$ we have $\sup_{\sigma \in \Sigma^C} \inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi) = 1$;

- *optimality* if for every game graph $G \in \mathcal{G}$, for every state $s \in S$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(s) = \inf_{\pi \in \Pi} \text{Pr}_s^{\sigma, \pi}(\Phi)$.

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. This is formalized by the proposition below.

Proposition 1 *If a family Σ^C of strategies suffices for sure winning with respect to an ω -regular objective Φ on all 2-player game graphs, then the family Σ^C suffices for sure winning with respect to Φ also on all $1^{1/2}$ -player and $2^{1/2}$ -player game graphs.*

The following result is the classical determinacy result for 2-player deterministic graph games with Rabin and Streett objectives.

Theorem 3 (Pure memoryless and finite-memory strategies) 1.

[12, 10] *The family Σ^{PM} of pure memoryless strategies suffices for sure winning with respect to all Rabin objectives on 2-player game graphs.*

2. [14] *The family Σ^{PF} of pure finite-memory strategies suffices for sure winning with respect to all Streett objectives on 2-player game graphs.*

3 MDPs, End Components, and Streett objectives

In this section we develop some facts on *end components* [7] that are needed for the further developments of the paper. We consider player-1 MDPs and hence strategies for player 1. Let $G = ((S, E), (S_1, S_2, S_\circ), \delta)$ with $S_2 = \emptyset$ be a $1^{1/2}$ -player game graph.

Definition 1 (End component) *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 any 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 Streett

(resp. Rabin) condition $\{(e_1, f_1), \dots, (e_d, f_d)\}$ is satisfied with probability 1, it suffices to show that for all reachable end components U , we have that $\forall i \in [1..d].(U \cap E_i \neq \emptyset \vee U \cap F_i = \emptyset)$ (resp., for Rabin conditions, $\exists i \in [1..d].(U \cap E_i = \emptyset \wedge U \cap F_i \neq \emptyset)$). To state the lemma, for $s \in S$ and $U \subseteq S$, we define $\Omega_s^U = \{\omega \in \Omega_s \mid \text{Inf}(\omega) = U\}$.

Lemma 1 [7] *For all state $s \in S$ and strategies $\sigma \in \Sigma$, we have $\Pr_s^\sigma(\bigcup_{U \in \mathcal{E}} \Omega_s^U) = 1$.*

Next, we present a polynomial-time algorithm for computing the maximal probability of satisfying a Streett condition in an MDP; the algorithm will be used in later sections to argue that certain witnesses can be checked in polynomial time. Consider a set $\hat{P} = \{(E_1, F_1), \dots, (E_d, F_d)\}$ of Streett pairs. Let $U \in \mathcal{U}$ iff $U \in \mathcal{E}$ and for all $1 \leq i \leq d$, we have either $U \cap E_i \neq \emptyset$ or $U \cap F_i = \emptyset$. The following lemma states that the maximal probability of satisfying Streett(\hat{P}) is equal to the maximal probability of reaching $T_{end} = \bigcup_{U \in \mathcal{U}} U$.

Lemma 2 [2] $\langle\langle 1 \rangle\rangle_{val} \text{Streett}(\hat{P}) = \langle\langle 1 \rangle\rangle_{val} \text{Reach}(T_{end})$.

We present a polynomial-time algorithm for computing T_{end} ; the computation of the value then reduces to computing values of a MDP with a reachability objective which can be achieved by linear programming [6]. To state the algorithm, we say that an end-component $U \subseteq S$ is *maximal* in $V \subseteq S$ if $U \subseteq V$, and if there is no end-component U' with $U \subset U' \subseteq V$. Given a set $V \subseteq S$, we denote by $\text{MaxEC}(V)$ the set consisting in all maximal end components U such that $U \subseteq V$. This set can be computed in quadratic time with standard graph algorithms; see, e.g., [7]. The set T_{end} can be computed with the following algorithm.

```

L := MaxEC(S); D := ∅
while L ≠ ∅ do
  choose U ∈ L and let L := L \ {U}
  if ∀i ∈ [1..d].(U ∩ Ei ≠ ∅ ∨ U ∩ Fi = ∅)
    then D := D ∪ {U}
    else choose i ∈ [1..d] such that U ∩ Fi ≠ ∅, and
  let L := L ∪ MaxEC(U \ Fi)
  end if
end while
Return: Tend = ⋃U ∈ D U.

```

It is easy to see that every state $s \in S$ is considered as part of an end-component in the **else**-part of the above algorithm at most once for every $1 \leq i \leq d$; hence, the algorithm runs in time polynomial in $|G|$ and $|P|$.

4 Almost-sure winning strategies in Rabin games

In this section we show that pure memoryless strategies suffices 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. This also gives a direct proof of the fact that the limit-sure and almost-sure winning sets coincide in $2^{1/2}$ -player games with Rabin objectives. Since any ω -regular objective can be expressed as a Rabin objective the result holds for all ω -regular objectives in $2^{1/2}$ -player games. Moreover, the reduction allows us to apply the algorithms for 2-player Rabin games for qualitative analysis of $2^{1/2}$ -player games with Rabin objectives. In the next section, we use the existence of pure memoryless almost-sure winning strategies to prove existence of pure memoryless optimal strategies.

4.1 Reduction

Given a $2^{1/2}$ -player Rabin game $(G = ((S, E), (S_1, S_2, S_\circ), \delta), [\cdot] : S \rightarrow 2^P \setminus \emptyset)$, where $P = \{(e_1, f_1), (e_2, f_2), \dots, (e_d, f_d)\}$ is a set of d pairs of colors, we construct a 2-player Rabin game $(\overline{G} = ((\overline{S}, \overline{E}), (\overline{S}_1, \overline{S}_2, \overline{S}_\circ), [\cdot] : \overline{S} \rightarrow 2^P \setminus \emptyset)$. The construction is described as follows: for every state $s \in S_1 \cup S_2$, there is a state $\overline{s} \in \overline{S}$ with “the same” outgoing edges, i.e., $(s, t) \in E$ if and only if $(\overline{s}, \overline{t}) \in \overline{E}$. Each probabilistic state $s \in S_\circ$ is substituted by the gadget presented in Figure 1. More formally, the players play the following 3-step game in \overline{G} from a probabilistic state \overline{s} . For the state \overline{s} we have $[\overline{s}] = [s]$. First, in vertex \overline{s} player 2 chooses a successor $(\tilde{s}, 2k)$, for $k \in \{0, 1, 2, \dots, d\}$. For every state $(\tilde{s}, 2k)$ we have $[(\tilde{s}, 2k)] = [s]$. For $k > 1$, in state $(\tilde{s}, 2k)$ player 1 chooses from two successors: state $(\hat{s}, 2k - 1)$ with $[(\hat{s}, 2k - 1)] = e_k$; state $(\hat{s}, 2k)$ with $[(\hat{s}, 2k)] = f_k$. In state $(\tilde{s}, 0)$ there is only one successor $(\hat{s}, 0)$ with $[(\hat{s}, 0)] = \{f_1, f_2, \dots, f_d\}$. Finally, in a state (\hat{s}, k) 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.

Let \overline{U}_1 and \overline{U}_2 be the sure-winning sets for players 1 and 2, respectively, in the 2-player Rabin game \overline{G} . Define sets U_1 and U_2 of states in the $2^{1/2}$ -player Rabin game G by $U_1 = \{s \in S \mid \overline{s} \in \overline{U}_1\}$, and $U_2 = \{s \in S \mid \overline{s} \in \overline{U}_2\}$. By the determinacy of 2-player Rabin games [11] (Theorem 3) we have that $\overline{U}_1 \cup \overline{U}_2 = \overline{S}$, and hence $U_1 \cup U_2 = S$.




Figure 1: The gadget for the reduction of a $2^{1/2}$ -player parity game to a 2-player parity game.

Definition 2 (Winning strongly connected component and end components)

Let G be a 1-player game graph with a $\text{Rabin}(\hat{P})$ objective for player 1 and $P = \{ (e_1, f_1), (e_2, f_2), \dots, (e_d, f_d) \}$ of d pairs of colors. A strongly connected component (s.c.c) C in G is winning for player 1 if there is a $i \in \{ 1, 2, \dots, d \}$ such that $C \cap F_i \neq \emptyset$ and $C \cap E_i = \emptyset$; otherwise C is winning for player 2. If G is a MDP with the set P of colors, then an end component C in G is winning for player 1 if there is an $i \in \{ 1, 2, \dots, d \}$ such that $C \cap F_i \neq \emptyset$ and $C \cap E_i = \emptyset$; otherwise C is winning for player 2. ■

Lemma 3 Let G be a $2^{1/2}$ -player game graph, and let $P = \{ (e_1, f_1), (e_2, f_2), \dots, (e_d, f_d) \}$ be a set of pairs of colors, and let $\hat{P} = \{ (E_1, F_1), \dots, (E_d, F_d) \}$ be the corresponding sets of pairs of states. There exists pure memoryless strategy σ for player 1 in the game G , such that for all strategy π for player 2 we have $\text{Pr}_s^{\sigma, \pi}(\text{Rabin}(\hat{P})) = 1$, for all states $s \in U_1$. Hence $U_1 \subseteq \langle\langle 1 \rangle\rangle_{\text{almost}}(\text{Rabin}(\hat{P}))$.

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

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




Figure 2: The strategy sub-graph in \overline{G}_σ .

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

- For a state $\overline{s} \in \overline{C} \cap \overline{S}_2$ keep all the edges $(\overline{s}, \overline{t})$ such that $\overline{t} \in \overline{C}$.
- For a state $\overline{s} \in \overline{C} \cap \overline{S}_0$ the sub-game is defined as follows:
 - At state \overline{s} choose the edges to state $(\tilde{s}, 2i)$ such that $i \in I$.
 - For a state $(\hat{s}, 2i)$, player 2 chooses a successor that shortens the distance to the vertex set $\overline{C} \cap E_i$ in the game G .

The construction is illustrated in Fig. 2.

We now prove that every terminal s.c.c. is winning for player 2 in the subgame thus constructed in $(\overline{G} \upharpoonright \overline{C})_\sigma$, where \overline{C} is the set of states in the gadget of states in C . Consider any arbitrary terminal s.c.c \overline{Y} in the subgame constructed in $(\overline{G} \upharpoonright \overline{C})_\sigma$. It follows from the construction that for every $i \in \{1, 2, \dots, d\} \setminus I$ 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 $(\tilde{s}, 2i)$ such that the strategy $\bar{\sigma}$ chooses the successor $(\hat{s}, 2i - 1)$ then the $E_i \cap \bar{Y} \neq \emptyset$ since $[(\tilde{s}, 2i - 1)] = e_i$.
2. If for every state $(\tilde{s}, 2i)$ the strategy for player 1 chooses the successor $(\hat{s}, 2i)$ then since $(\hat{s}, 2i)$ is a player 2 state player 2 chooses an successor to shorten distance to the vertex set E_i and hence the terminal s.c.c. \bar{Y} must contain a state \bar{s} such that $[\bar{s}] = e_i$. Hence $E_i \cap \bar{Y} \neq \emptyset$.

Now we argue that for every probabilistic state $s \in S_{\circ} \cap U_1$, all of its successors are in U_1 . Otherwise, player 2 in the state \bar{s} of the game \bar{G} can choose the successor $(\tilde{s}, 0)$ and then a successor to its winning set \bar{U}_2 , which contradicts the assumption that the strategy $\bar{\sigma}$ is a sure winning strategy for the player 1 in the game \bar{G} . It follows from Lemma 1 that for any strategy π with probability 1 the set of states visited infinitely often along the play $\omega_{\sigma, \pi}$ 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. ■

Lemma 4 *Let G be a $2^{1/2}$ -player game graph with a set $P = \{ (e_1, f_1), (e_2, f_2), \dots, (e_d, f_d) \}$ of d pairs of colors and winning objective $\text{Rabin}(\hat{P})$ for player 1. There exists finite-memory strategy π for player 2 in the game G such that for all strategy σ for player 1 we have $\text{Pr}_s^{\sigma, \pi}(\text{Streett}(\hat{P})) > 0$, for all states $s \in U_2$. Hence $S \setminus U_1 \subseteq S \setminus \langle\langle 1 \rangle\rangle_{\text{almost}}(\text{Rabin}(\hat{P}))$.*

Proof. The proof idea is similar to the proof of Lemma 3. Consider a finite-memory sure winning strategy $\bar{\pi}$ for player 2 in the game $\bar{G} \upharpoonright \bar{U}_2$; and π be the corresponding strategy in G . Let \mathbf{M} be the memory of the strategy $\bar{\pi}$. We argue that every end component in the game $(G \upharpoonright U_2)_{\pi}$ is winning for player 2. Consider the product game $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})$ and the corresponding memoryless strategy π_m of π in the game $G \times \mathbf{M}$. It suffices to argue that every end component in $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$ is winning for player 2. Let C be a end component in $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$ that is winning for player 1, then we construct an s.c.c. that is winning for player 1 in $(\bar{G} \times \mathbf{M} \upharpoonright \bar{U}_2 \times \mathbf{M})_{\bar{\pi}_m}$, which is a contradiction since $\bar{\pi}$ is a sure winning strategy for player 2 in $\bar{G} \upharpoonright \bar{U}_2$. We describe the key steps to construct a winning s.c.c. \bar{C} from a winning end component C ; mainly we describe the strategy corresponding to a probabilistic state. If C is a winning end component for player 1 and let i be the witness Rabin pair that C is winning, i.e., $C \cap F_i \neq \emptyset$ and $C \cap E_i = \emptyset$. The strategy for player 1 is as follows:

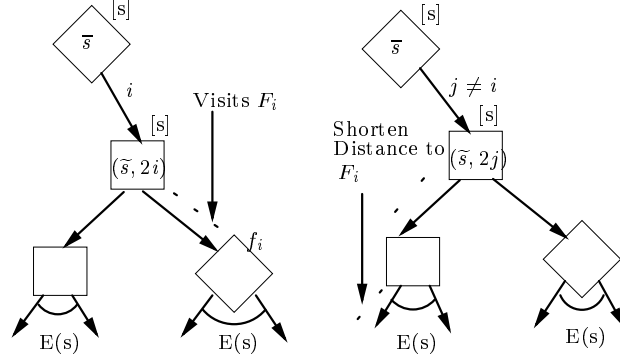


Figure 3: The strategy sub-graph in $\overline{G}_{\overline{\pi}}$.

- If the strategy for player 2 at a state $(\overline{s}, \mathbf{m})$ chooses successor $((\tilde{s}, 0), \mathbf{m}')$ then the following successor state is $((\hat{s}, 0), \mathbf{m}')$ and since $[(\hat{s}, 0)] = \{f_1, f_2, \dots, f_d\}$ player 1 ensures that a state in F_i is visited.
- If the strategy for player 2 at a state $(\overline{s}, \mathbf{m})$ chooses a successor $((\tilde{s}, 2i), \mathbf{m}')$ then player 1 chooses a successor $((\hat{s}, 2i - 1), \mathbf{m}')$, where $\mathbf{m}, \mathbf{m}' \in \mathbf{M}$. Since $[(\tilde{s}, 2i)] = f_i$ player 1 ensures that a state in F_i is visited.
- If the strategy for player 2 at a state $(\overline{s}, \mathbf{m})$ chooses a successor $((\tilde{s}, 2j), \mathbf{m}')$, for $j \neq i$, then player 1 chooses a successor $((\hat{s}, 2j - 1), \mathbf{m}')$ and then a successor to shorten distance to the set F_i , where $\mathbf{m}, \mathbf{m}' \in \mathbf{M}$. Since $[(\tilde{s}, 2j - 1)] = e_j \neq e_i$, player 1 ensures that a state in E_i is not visited.

The construction is illustrated in Fig. 3.

Consider any terminal s.c.c. \overline{Y} in the sub-game thus constructed. The strategy for player 1 ensures that in the sub-game \overline{C} whenever a state \overline{s} is visited such that $s \in S_{\circ}$, no state in E_i is visited. Since $C \cap E_i = \emptyset$ it follows that $\overline{Y} \cap E_i = \emptyset$. Moreover, the strategy for player 1 ensures that a state in F_i is always visited, i.e., $\overline{Y} \cap F_i \neq \emptyset$. Hence in the sub-game of $(\overline{G} \times \mathbf{M} \upharpoonright \overline{C} \times \mathbf{M})_{\pi_m}$ every terminal s.c.c. \overline{Y} is winning for player 1, i.e., $F_i \cap \overline{Y} \neq \emptyset$ and $E_i \cap \overline{Y} = \emptyset$. However, this is a contradiction since $\overline{\pi}$ is a sure winning strategy for player 2. Hence, all the end-components in $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$ are winning for player 2.

Note that $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$ is a finite-state player-1 MDP and if player 1 can win almost-surely she can win by a pure memoryless strategy. Hence,

it suffices to argue that player 2 wins with probability greater than 0 from every state $s \in (G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$ against all pure memoryless strategy σ for player 1 in $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\pi_m}$. For every probabilistic state $s \in S_{\circ} \cap U_2$, at least one successor must be in the set U_2 . Otherwise if both the successors of s are in U_1 it follows from the construction of the gadget that $\bar{s} \in \langle\langle 1 \rangle\rangle_{\text{sure}}(\text{Reach}(\overline{U_1}))$ in the game \overline{G} . In other words, there is a strategy for player 1 in the 2-player game to reach the set $\overline{U_1}$ from \bar{s} ; this leads to $\bar{s} \in \overline{U_1}$, which is a contradiction. Hence for any pure memoryless strategy σ consider the Markov chain $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\sigma, \pi_m}$. From every state $s \in U_2 \times \mathbf{M}$ there is a path from to a terminal strongly connected component in $U_2 \times \mathbf{M}$, i.e., there is a path to a closed recurrent class that is a subset of $U_2 \times \mathbf{M}$. Every end component is winning for player 2 in $U_2 \times \mathbf{M}$. Hence, for every state $s \in U_2 \times \mathbf{M}$ there is a path to a closed recurrent class that is winning for player 2. Therefore for any pure memoryless strategy σ , in the Markov chain, $(G \times \mathbf{M} \upharpoonright U_2 \times \mathbf{M})_{\sigma, \pi_m}$, if the play starts at any state $s \in U_2 \times \mathbf{M}$ there is a positive probability that it reaches a terminal strongly connected component that is winning for player 2. Hence the desired result follows. ■

It follows from Lemma 3 and Lemma 4 that $U_1 = \langle\langle 1 \rangle\rangle_{\text{almost}}(\text{Rabin}(\hat{P}))$. Moreover, pure memoryless almost-sure winning strategies exist for $2^{1/2}$ -player Rabin games.

Theorem 4 *The family Σ^{PM} of pure memoryless strategies suffices for almost-sure winning with respect to Rabin objectives on $2^{1/2}$ -player game graphs.*

5 From almost-sure to optimal

In this section we show how to extend the sufficiency result for a family of strategies from almost-sure to optimality for any ω -regular objective.

Definition 3 (Value Class) *Given an ω -regular objective Φ , for any real $r \in \mathbb{R}$, we denote by $\text{VC}(r)$ the value class with value r , i.e., $\text{VC}(r) = \{s \in S \mid \langle\langle 1 \rangle\rangle_{\text{val}} \Phi(s) = r\}$. ■*

The following Proposition states that there exists optimal strategies for player 1 such that they never choose an edge to a lower value class.

Proposition 2 *For all ω -regular objectives Φ there exists optimal strategy σ for player 1 such that for any sequence $\vec{w} \in S^*$ and $s \in S_1$ we have $\sigma(\vec{w} \cdot s)(t) = 0$ if $\langle\langle 1 \rangle\rangle_{\text{val}} \Phi(t) < \langle\langle 1 \rangle\rangle_{\text{val}} \Phi(s)$.*

The following Proposition follows from Theorem 4.

Proposition 3 ([3]) *For all ω -regular objectives Φ and for all $2^{1/2}$ -player game graphs, the limit-sure and almost-sure winning sets coincide: $\langle\langle 1 \rangle\rangle_{\text{limit}}(\Phi) = \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{\text{limit}}(\Omega \setminus \Phi) = \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi)$.*

Definition 4 (Boundary probabilistic states) *Given a value class $\text{VC}(r)$ a probabilistic state s is a boundary probabilistic vertex if there exists a successor t of s such that $\langle\langle 1 \rangle\rangle_{\text{val}}\Phi(t) \neq \langle\langle 1 \rangle\rangle_{\text{val}}\Phi(s)$. It may be noted that for every boundary probabilistic state s , there exists a successors t_1, t_2 of s such that $\langle\langle 1 \rangle\rangle_{\text{val}}\Phi(t_1) < \langle\langle 1 \rangle\rangle_{\text{val}}\Phi(s)$ and $\langle\langle 1 \rangle\rangle_{\text{val}}\Phi(t_2) > \langle\langle 1 \rangle\rangle_{\text{val}}\Phi(s)$. ■*

Lemma 5 *Consider a $2^{1/2}$ -player game with an ω -regular objective Φ . Given a value class $\text{VC}(r)$, with $0 < r < 1$, let $B(r)$ be the set of boundary probabilistic states of the value class $\text{VC}(r)$. Convert each of the state in $B(r)$ to a sink state that is winning for player 1. Let the new game be G' . Then player 1 wins almost-surely in the sub-game $G' \upharpoonright \text{VC}(r)$.*

Proof. Assume that player 1 does not win almost-surely from every state in $G' \upharpoonright \text{VC}(r)$. Then there exists a state where player 2 wins with positive bounded probability. It follows from Corollary 1 of [8] and Proposition 3 that there exist a non-empty set $U \subseteq \text{VC}(r)$ such that that player 2 wins almost-surely from U . Consider a optimal strategy σ that never chooses an edge with positive probability to a lower value class (such a strategy exist from Proposition 2). 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 $\text{VC}(r)$ we have $t \in U$. Note that it follows that every move of strategy σ exists in U . Hence player 2 wins almost-surely from U against σ . However, this is a contradiction to the assumption that $r > 0$ and that σ is an optimal strategy. ■

It follows, from Lemma 5 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 as conditional almost-sure winning strategy.

Definition 5 (Qualitative optimal strategy) *A strategy σ is qualitative optimal for player 1, for an ω -regular objective Φ , if the following conditions hold:*

- *For every state $s \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ the strategy σ is almost-sure winning.*

- For every state $s \in \text{VC}(r)$ such that $0 < r < 1$, there is a constant $c > 0$ such that

$$\inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi) \geq c. \blacksquare$$

Lemma 6 Consider a strategy σ , and an ω -regular objective Φ , such that σ is almost-sure winning from every state $s \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$, and σ is a conditional almost-sure winning strategy from every state s in $S \setminus \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi)$, then σ is qualitative optimal for Φ .

Proof. Since the strategy σ is conditional almost-sure winning it follows that any strategy π that is optimal against σ the play $\omega_s^{\sigma, \pi}$ reaches the boundary probabilistic states with positive probability, for $s \in \text{VC}(r)$ and $r > 0$. From every boundary probabilistic state the game proceeds to a higher value class with positive probability. By an easy induction on the number of value classes it follows that from every state in $S \setminus \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi)$ the game reaches $\langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ with positive probability. Since σ is almost-sure winning for every state $s \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ it follows that σ is qualitative optimal. \blacksquare

Definition 6 (Locally optimal strategies) A strategy σ is locally optimal for an ω -regular objective Φ if for all $\vec{w} \in S^*$ and $s \in S_1$ we have $\sigma(\vec{w} \cdot s)(t) = 0$ if $\langle\langle 1 \rangle\rangle_{\text{val}}\Phi(t) < \langle\langle 1 \rangle\rangle_{\text{val}}\Phi(s)$. \blacksquare

Note that by definition a conditional almost-sure winning strategy is locally optimal. The following proof is similar to the proof of Lemma 5.3 of [4].

Lemma 7 Consider a $2^{1/2}$ -player game G with an ω -regular objective Φ for player 1. Let σ be a memoryless strategy such that σ is qualitative optimal and locally optimal for Φ . Then σ is an optimal strategy for Φ .

Proof. Given σ is a memoryless the game G_σ is a player-2 MDP. Since σ is a qualitative optimal strategy it follows that for every state $s \in \text{VC}(r)$, for $r > 0$, for all strategy π of player 2 we have $\Pr_s^{\sigma, \pi}(\Phi) > c$, for some constant c . Hence, the set of almost-sure winning states for player 2 in G_σ coincide with the set of almost-sure winning states in G . Let us denote by W_2 the set of almost-sure winning states for player 2 in G and G_σ , i.e., $W_2 = \langle\langle 2 \rangle\rangle_{\text{almost}}(\Omega \setminus \Phi)$. It follows from the analysis of MDPs that in the game G_σ , for all state s we have $\langle\langle 2 \rangle\rangle_{\text{val}}(\Omega \setminus \Phi)(s) = \langle\langle 2 \rangle\rangle_{\text{val}}(\text{Reach}(W_2))(s)$.

By [6, 5], the values are the unique solution to the linear program consisting in minimizing $\sum_{s \in S} x_s$ subject to:

$$\begin{aligned} \forall s \in S_{\circ} : & & x_s &= \sum_{t \in E(s)} x_t \cdot \delta(s, t) & & \forall s \in S : & & x_s \geq 0 \\ \forall s \in S_1 : & & x_s &= \sum_{t \in E(s)} x_t \cdot \sigma(s)(t) & & \forall s \in W_2 : & & x_s = 1 \\ \forall s \in S_2, \forall t \in E(s) : & & x_s &\geq x_t \end{aligned}$$

Let us denote by x^* the optimal solution of the above liner program. The local optimality of the strategy σ ensures that for every state $s \in S_1$, $x_s = \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s)$, satisfy the constraints of the linear program. Moreover, $x_s = \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)$ satisfy the constraints for all state $s \in S_2 \cup S_{\circ}$. Hence, $x_s = \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)$ is a feasible solution of the linear program. Since the above linear program is a minimization problem we have $x^* \leq x_s = \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s)$ for all $s \in S$. It follows that in the MDP G_{σ} we have $\sup_{\pi \in \Pi} \Pr_s^{\pi}(\Omega \setminus \Phi) \leq \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s)$. Hence it follows that $\inf_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Phi) = 1 - \sup_{\pi \in \Pi} \Pr_s^{\sigma, \pi}(\Omega \setminus \Phi) \geq 1 - \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) = \langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$. This implies that σ is an optimal strategy for player 1 in G . ■

Observe that arguments similar to the arguments of Lemma 7 can be extended to the synchronous product of the game graph G with any finite memory M . Hence, the proof of Lemma 7 can be easily extended for finite-memory strategy σ in place of memoryless strategy σ . This gives us the following general Theorem.

Theorem 5 *If a family $\Sigma^C \subseteq \Sigma^F$ of strategies suffices for almost-sure winning with respect to an ω -regular objective Φ on $2^{1/2}$ -player game graphs, then Σ^C suffices for optimality with respect to objective Φ on $2^{1/2}$ -player game graphs.*

Since pure memoryless suffices for almost-sure winning with respect to Rabin objectives on $2^{1/2}$ -player game graphs (Theorem 4) the following Theorem is immediate from Theorem 5.

Theorem 6 *The family Σ^{PM} of pure memoryless strategy suffices for optimality with respect to all Rabin objectives on $2^{1/2}$ -player game graphs.*

Theorem 7 *Given a $2^{1/2}$ -player game graph G , an objective Φ for player 1, a state $s \in S$ and a rational $r \in \mathbb{R}$, the complexity of determining whether $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) \geq r$ is as follows:*

1. NP-complete if Φ is a Rabin objective.

2. *coNP-complete if Φ is a Streett objective.*
3. *[4, 18] NP \cap coNP if Φ is a parity objective.*

Proof.

1. Let G be a $2^{1/2}$ -player game with a Rabin objective $\text{Rabin}(\hat{P})$ for player 1. Given a pure memoryless optimal strategy σ for player 1 the game G_σ is a player-2 MDP with Streett objective for player 2. Since the values of MDPs with Streett objective can be computed in polynomial time (Section 3) the problem is in NP. The NP-hardness proof follows from the fact the 2-player games with Rabin objectives are NP-hard [12, 23].
2. Follows immediately from the fact that Streett objectives are complementary to Rabin objectives.
3. Follows from the previous two completeness result, as a parity objective is both a Rabin objective and a Streett 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. Trading memory for randomness. In *Proceedings of the First Annual Conference on Quantitative Evaluation of Systems*. IEEE Computer Society Press, 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] C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *ICALP'90*, volume 443 of *LNCS*, pages 336–349. Springer, 1990.

- [7] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- [8] L. de Alfaro and T. A. Henzinger. Concurrent ω -regular games. In *LICS'00*, pages 141–154. IEEE Computer Society Press, 2000.
- [9] L. de Alfaro, T. A. Henzinger, and O. Kupferman. Concurrent reachability games. In *FOCS'98*, pages 564–575. IEEE Computer Society Press, 1998.
- [10] 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.
- [11] E. A. Emerson and C. Jutla. Tree automata, μ -calculus and determinacy. In *FOCS'91*, pages 368–377. IEEE Computer Society Press, 1991.
- [12] E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *FOCS'88*, pages 328–337. IEEE Computer Society Press, 1988.
- [13] J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
- [14] Y. Gurevich and L. Harrington. Trees, automata, and games. In *STOC'82*, pages 60–65. ACM Press, 1982.
- [15] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer-Verlag, 1992.
- [16] D. A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
- [17] D.A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- [18] 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.
- [19] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM Press, 1989.

- [20] T. E. S. Raghavan and J. A. Filar. Algorithms for stochastic games — a survey. *ZOR — Methods and Models of Operations Research*, 35:437–472, 1991.
- [21] P.J. Ramadge and W.M. Wonham. Supervisory control of a class of discrete-event processes. *SIAM Journal of Control and Optimization*, 25(1):206–230, 1987.
- [22] W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 133–191. Elsevier, 1990.
- [23] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
- [24] W.Zielonka. Perfect-information stochastic parity games. In *FoS-SaCS'2004*, pages 499–513, 2004.