

Trading Memory for Randomness*

Krishnendu Chatterjee[†]

Luca de Alfaro[§]

Thomas A. Henzinger^{†‡}

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

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

[‡] Computer and Communication Sciences, EPFL, Switzerland

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

Abstract

Strategies in repeated games can be classified as to whether or not they use memory and/or randomization. We consider Markov decision processes and 2-player graph games, both of the deterministic and probabilistic varieties. We characterize when memory and/or randomization are required for winning with respect to various classes of ω -regular objectives, noting particularly when the use of memory can be traded for the use of randomization. In particular, we show that Markov decision processes allow randomized memoryless optimal strategies for all Müller objectives. Furthermore, we show that 2-player probabilistic graph games allow randomized memoryless strategies for winning with probability 1 those Müller objectives which are upward-closed. Upward-closure means that if a set α of infinitely repeating vertices is winning, then all supersets of α are also winning.

1 Introduction

A two-player graph game is played on a directed graph whose vertices are partitioned into player-1 vertices and player-2 vertices. The two players move a token along the edges of the graph. At player-1 vertices, the first player chooses an outgoing edge, and at player-2 vertices the second player moves the token to a neighboring vertex. The outcome of the game is an infinite path through the graph. An objective for a player can be specified as an ω -regular condition on the outcome of the game [27, 23]. These ω -regular games are used in the modeling [1, 14, 11], verification [34, 12, 2, 20], and control (synthesis) [6, 3, 31, 29] of state-based systems, where the vertices represent states and the players represent components or agents of a system.

A strategy for a player is a recipe that describes how the player chooses a move whenever it is her turn. Strategies can be classified as follows. A *pure* strategy always chooses a particular edge to extend the game. In contrast, a *randomized* strategy may choose at a vertex a probability distribution over the outgoing edges. In other words, a randomized strategy instructs the player to toss a coin in order to decide on her move. Randomized strategies are not helpful to win the game with certainty, but they may be useful to win the game with probability 1. To formalize this, notice that every vertex v and every pair (σ, π) consisting of a player-1 strategy σ and a player-2 strategy π determines (1) a set $\text{Outcome}(v, \sigma, \pi)$ of possible outcomes if the two players follow the strategies σ and π starting the game from the initial vertex v , and (2) a probability distribution over $\text{Outcome}(v, \sigma, \pi)$ which indicates the likelihood of each possible outcome. We say that at vertex v player-1 *surely wins* the game with objective Φ if there is a player-1 strategy σ such that for all player-2 strategies π we have $\text{Outcome}(v, \sigma, \pi) \subseteq \Phi$, that is, every possible outcome satisfies Φ . A weaker condition is that player-1 *almost-surely wins* at v with objective Φ , meaning that there is a player-1 strategy σ such that for all player-2 strategies π the set $\text{Outcome}(v, \sigma, \pi) \setminus \Phi$ of undesirable possible outcomes has probability 0.

Strategies can be classified also according to their memory requirements. A *memoryless* strategy depends only on the current position of the token. In contrast, a *memory* strategy may depend on the path the token has taken to obtain its current position. It is well-known that there are ω -regular objectives which can be surely won using a memory strategy, but cannot be surely won using a memoryless strategy. Here is a simple example: there are three vertices, v_0 , v_1 , and v_2 ; at v_0 player 1 moves the token to either v_1 or v_2 , and at both of these vertices, player 2 has no choice but to move the token back to v_0 . The objective to visit both v_1 and v_2 infinitely often cannot be won by player 1 without using memory; for instance, a winning player-1 strategy may alternate the two moves $v_0 \rightarrow v_1$ and $v_0 \rightarrow v_2$. Note,

*This research was supported in part 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.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE Trading Memory for Randomness				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of California at Berkeley, Department of Electrical Engineering and Computer Sciences, Berkeley, CA, 94720-1770				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

however, that in this game player 1 does have a *randomized* memoryless strategy to *almost* surely win, such as the strategy that always chooses the successor of v_0 uniformly at random. In other words, player 1 can trade memory against a random coin. We systematically study this trade-off for all ω -regular objectives.

The results are categorized according to the form of the game graph and the form of the winning condition. For winning conditions, we use the classical classes of parity, Rabin, Streett, and Müller objectives [33]. For game graphs, we distinguish graphs without probabilistic vertices, which are the graphs described above, and graphs that may contain also probabilistic vertices. At a probabilistic vertex, the token is moved according to a fixed probability distribution over the outgoing edges, that is, neither of the two players can choose the successor vertex. Accordingly, we classify the game graphs into 1-*player* graphs (only player-1 vertices), 1 $\frac{1}{2}$ -*player* graphs (player-1 and probabilistic vertices), 2-*player* graphs (player-1 and player-2 vertices), and 2 $\frac{1}{2}$ -*player* graphs (player-1, player-2, and probabilistic vertices). Notice that 1-player graphs are transition systems, and 1 $\frac{1}{2}$ -player graphs are Markov decision processes (MDPs). Instead of almost-sure winning, we consider the more general condition of optimality. For a vertex v , a player-1 strategy σ , and a player-2 strategy π , let $\Pr_v^{\sigma, \pi}(\Phi)$ be the probability of the set $\text{Outcome}(v, \sigma, \pi) \cap \Phi$ of desirable possible outcomes. A player-1 strategy σ is *optimal* at v for Φ if $\inf_{\pi} \Pr_v^{\sigma, \pi}(\Phi) \geq \inf_{\pi} \Pr_v^{\sigma', \pi}(\Phi)$ for all player-1 strategies σ' . It can be shown that player-1 almost surely wins at v for Φ iff she has a strategy σ with $\inf_{\pi} \Pr_v^{\sigma, \pi}(\Phi) = 1$.

For Rabin objectives, it is known that pure memoryless strategies suffice for the sure winning of 2-player games [16, 15], and for the more special case of parity objectives, it is known that there always exist optimal strategies in 2 $\frac{1}{2}$ -player games which are both pure and memoryless [26, 5]. At the other extreme, Streett games are known to require memory for sure winning even in the 1-player case (cf. the above example), and it is easy to see that they also require memory for almost-sure winning in the 2-player case (cf. Example 3). However, for 1 $\frac{1}{2}$ -player Streett games, and more generally, for all 1 $\frac{1}{2}$ -player Müller games, we show that no memory is required for optimal strategies if randomization is available (Theorem 9). In other words, in MDPs the optimal value can be obtained without memory for every objective which cannot distinguish between two paths that visit the same vertices infinitely often. Furthermore, we show that if the objective is Rabin, then optimality in MDPs can be achieved by strategies that are both pure and memoryless (Theorem 8).

We then take a closer look at the general case of 2 $\frac{1}{2}$ -player ω -regular games. We define a Müller objective Φ to be *upward-closed* if for every infinite path $\tau \in \Phi$, if

every vertex that occurs infinitely often in τ also occurs infinitely often in τ' , then $\tau' \in \Phi$. For example, every generalized Büchi objective is upward-closed. We prove that memoryless strategies suffice for the almost-sure winning of upward-closed 2 $\frac{1}{2}$ -player games (Theorem 11). If randomization is not used, then upward-closed objectives (such as the generalized Büchi objective in the above example) may require memory for almost-sure winning; thus, the upward-closed games allow us to trade memory for randomization. Indeed, we give an example of 2-player Müller games with n vertices where sure winning requires $O(n)$ memory but almost-sure winning can be achieved without memory. Moreover, there is a game such that, if a Müller objective is not upward-closed, then randomized memoryless strategies are no better than pure memoryless strategies for almost-sure winning, and they are not as powerful as strategies with memory. This shows that the upward-closed Müller games are the most general games with ω -regular objectives where memory can be traded for randomization.

2 Preliminaries

Game graphs. A *turn-based probabilistic game graph* (2 $\frac{1}{2}$ -*player game graph*) $G = ((V, E), V_0, V_1, V_2, p)$ consists of a directed graph (V, E) , a partition V_0, V_1, V_2 of the vertex set V , and a probabilistic transition function $p: V_0 \rightarrow \mathcal{D}(V)$, where $\mathcal{D}(V)$ denotes the set of probability distributions over the vertex set V . The vertices in V_1 are the *player-1* vertices, where player 1 decides the successor vertex; the vertices in V_2 are the *player-2* vertices, where player 2 decides the successor vertex; and the vertices in V_0 are the *probabilistic* vertices, where the successor vertex is chosen according to the probabilistic transition function p . We assume that, for $u \in V_0$ and $v \in V$, we have $(u, v) \in E$ iff $p(u)(v) > 0$, and we often write $p(u, v)$ for $p(u)(v)$. For technical convenience we assume that in (V, E) every vertex has at least one outgoing edge, and we write $v \in E(u)$ for $(u, v) \in E$.

An infinite path, or *play*, of the game graph G is an infinite sequence $\langle v_0, v_1, v_2, \dots \rangle$ of vertices such that $(v_k, v_{k+1}) \in E$ for all $k \in \mathbb{N}$. We write Ω for the set of all plays, and for every vertex $v \in V$ we write Ω_v for the set of plays that start from the vertex v . A set $U \subseteq V$ of vertices is called *p-closed* if for every $u \in U \cap V_0$, we have $(u, v) \in E$ implies $v \in U$. A *p-closed* subset of V induces a *subgame graph* of G , indicated by $G \upharpoonright U$, if for every vertex $u \in U \cap (V_1 \cup V_2)$ there is a vertex $v \in U$ such that $(u, v) \in E$.

The *turn-based deterministic game graphs* (2-*player game graphs*) are the special case of the 2 $\frac{1}{2}$ -player game graphs with $V_0 = \emptyset$. The *Markov decision processes* (1 $\frac{1}{2}$ -*player game graphs*) are the special case of the 2 $\frac{1}{2}$ -player game graphs with $V_2 = \emptyset$ or $V_1 = \emptyset$. We refer to the

MDPs with $V_2 = \emptyset$ as *player-1 MDPs*, and to the MDPs with $V_1 = \emptyset$ 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 vertices; a player-2 transition system has only player-2 vertices.

Strategies. A *strategy* for player 1 is a function $\sigma: V^* \cdot V_1 \rightarrow \mathcal{D}(V)$ that assigns a probability distribution to every finite sequence $\vec{w} \in V^* \cdot V_1$ of vertices, 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 $\vec{w}$, she chooses the next vertex according to the probability distribution $\sigma(\vec{w})$. A strategy must prescribe only available moves, i.e., for all $\vec{w} \in V^*$, $v \in V_1$, and $u \in V$, if $\sigma(\vec{w} \cdot v)(u) > 0$, then $(v, u) \in E$. 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.

Once a starting vertex $v \in V$ and strategies $\sigma \in \Sigma$ and $\pi \in \Pi$ for the two players are fixed, the outcome of the game is a random path $\omega_v^{\sigma, \pi}$ for which the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega_v$ is a measurable set of paths. Given strategies σ for player 1 and π for player 2, a play $\langle v_0, v_1, v_2, \dots \rangle$ is *feasible* if for every $k \in \mathbb{N}$ the following three conditions hold: (1) if $v_k \in V_0$, then $(v_k, v_{k+1}) \in E$; (2) if $v_k \in V_1$, then $\sigma(v_0, v_1, \dots, v_k)(v_{k+1}) > 0$; and (3) if $v_k \in V_2$ then $\pi(v_0, v_1, \dots, v_k)(v_{k+1}) > 0$. Given strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a vertex v , we denote by $\text{Outcome}(v, \sigma, \pi) \subseteq \Omega_v$ the set of feasible plays that start from v given strategies σ and π . For a vertex $v \in V$ and an event $\mathcal{A} \subseteq \Omega_v$, we write $\Pr_v^{\sigma, \pi}(\mathcal{A})$ for the probability that a path belongs to $\mathcal{A}$ if the game starts from the vertex v 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. Objectives for the players in nonterminating games are specified by providing the set of *winning plays* $\Phi \subseteq \Omega$ for each player. In this paper we study only zero-sum games [30, 18], where the objectives of the two players are strictly competitive. In other words, it is implicit that if the objective of one player is Φ , then the objective of the other player is $\Omega \setminus \Phi$. Given a game graph G and an objective $\Phi \subseteq \Omega$, we write (G, Φ) for the game played on the graph G with the objective Φ for player 1.

A general class of objectives are the Borel objectives [21]. A *Borel objective* $\Phi \subseteq V^\omega$ is a Borel set in the Cantor topology on V^ω . In this paper we consider ω -regular objectives [33], which lie in the first $2^{1/2}$ levels of the Borel hierarchy (i.e., in the intersection of Σ_3 and Π_3). The ω -regular objectives, and subclasses thereof, can be specified

in the following forms.

For a play $\omega = \langle v_0, v_1, v_2, \dots \rangle \in \Omega$, we define $\text{Inf}(\omega) = \{v \in V \mid v_k = v \text{ for infinitely many } k \geq 0\}$ 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 $\llbracket \cdot \rrbracket: C \rightarrow 2^V$ for a function that maps each color to a set of vertices. Inversely, given a set $U \subseteq V$ of states, we write $[U] = \{c \in C \mid \llbracket c \rrbracket \cap U \neq \emptyset\}$ for the set of colors that occur in U .

1. *Reachability and safety objectives.* Given a color c , the reachability objective requires that some vertex of color c be visited. Let $T = \llbracket c \rrbracket$ be the set of so-called *target* vertices. Formally, we write $\text{Reach}(T) = \{\langle v_0, v_1, v_2, \dots \rangle \in \Omega \mid v_k \in T \text{ for some } k \geq 0\}$ for the set of winning plays. Given c , the safety objective requires that only vertices of color c be visited. Let $F = \llbracket c \rrbracket$ be the set of so-called *safe* vertices. Formally, the set of winning plays is $\text{Safe}(F) = \{\langle v_0, v_1, v_2, \dots \rangle \in \Omega \mid v_k \in F \text{ for all } k \geq 0\}$.
2. *Büchi and generalized Büchi objectives.* Given a color c , the Büchi objective requires some vertex of color c be visited infinitely often. Let $B = \llbracket c \rrbracket$ be the set of so-called *Büchi* vertices. Formally, the set of winning plays is $\text{Büchi}(B) = \{\omega \in \Omega \mid \text{Inf}(\omega) \cap B \neq \emptyset\}$. Given a set $C = \{c_1, \dots, c_m\}$ of colors, the generalized Büchi objective requires that all m Büchi objectives in C be satisfied. Formally, the set of winning plays is $\bigcap_{1 \leq i \leq m} \text{Büchi}(\llbracket c_i \rrbracket)$.
3. *Rabin, parity, and Streett objectives.* Given a set $P = \{(e_1, f_1), \dots, (e_m, f_m)\}$ of pairs of colors, the Rabin objective requires that for some $1 \leq i \leq m$, all vertices of color e_i be visited finitely often and some vertex of color f_i be visited infinitely often. Let $R = \{(E_1, F_1), \dots, (E_m, F_m)\}$ 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 $1 \leq i \leq m$. Formally, the set of winning plays is $\text{Rabin}(R) = \{\omega \in \Omega \mid \exists 1 \leq i \leq m. (\text{Inf}(\omega) \cap E_i = \emptyset \wedge \text{Inf}(\omega) \cap F_i \neq \emptyset)\}$. The *parity* (or *Rabin-chain*) objectives are the special case of Rabin objectives where $E_1 \subset F_1 \subset \dots \subset E_m \subset F_m$. Given P , the Streett objective requires that for each $1 \leq i \leq m$, if some vertex of color f_i is visited infinitely often, then some vertex of color e_i is visited infinitely often. Formally, for the set $S = \{(E_1, F_1), \dots, (E_m, F_m)\}$ of so-called *Streett pairs*, the set of winning plays is $\text{Streett}(S) = \{\omega \in \Omega \mid \forall 1 \leq i \leq m. (\text{Inf}(\omega) \cap E_i \neq \emptyset \vee \text{Inf}(\omega) \cap F_i = \emptyset)\}$. Note that the Rabin and Streett objectives are dual.
4. *Müller and upward-closed objectives.* Given a set C of colors, and a set $\Gamma \subseteq 2^C$ of subsets of the colors, the Müller objective requires that the set of colors that appear infinitely often in a play is exactly one of the sets

in Γ . Formally, for the set $M_\Gamma = \{U \subseteq V \mid [U] \in \Gamma\}$ of so-called *Müller sets* of vertices, the set of winning plays is $\text{Müller}(M_\Gamma) = \{\omega \in \Omega \mid \text{Inf}(\omega) \in M_\Gamma\}$. We call Γ a (game graph independent) *specification* of the objective $\text{Müller}(M_\Gamma)$, because Γ does not refer to the vertex names of G . The specification Γ is *upward-closed* if for all $\alpha \subseteq \beta \subseteq C$, if $\alpha \in \Gamma$, then $\beta \in \Gamma$.

The generalized Büchi objectives, Rabin objectives, and Streett objectives are special cases of Müller objectives. In particular, all Büchi and generalized Büchi objectives are upward-closed. Moreover, reachability and safety objectives can be turned into Büchi objectives on slightly modified game graphs. However, a parity, Rabin, or Streett objective need not be upward-closed.

We commonly use terminology like the following: a $2^{1/2}$ -player *Müller game* $(G, \text{Müller}(M_\Gamma))$ consists of a $2^{1/2}$ -player game graph G and a Müller objective for player 1, where M_Γ is a set of Müller sets.

Values of a game and optimal strategies. Given objectives Φ 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:

$$\begin{aligned} \langle\langle 1 \rangle\rangle_{val}(\Phi)(v) &= \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \Pr_v^{\sigma, \pi}(\Phi) \\ \langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(v) &= \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \Pr_v^{\sigma, \pi}(\Omega \setminus \Phi) \end{aligned}$$

A strategy σ for player 1 is *optimal* from vertex v for objective Φ if $\langle\langle 1 \rangle\rangle_{val}(\Phi)(v) = \inf_{\pi \in \Pi} \Pr_v^{\sigma, \pi}(\Phi)$. The optimal strategies for player 2 are defined analogously.

Theorem 1 (Quantitative determinacy [24]). *For all $2^{1/2}$ -player game graphs, all Borel objectives Φ , and all vertices v ,*

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

Every ω -regular objective can be expressed as a parity objective [28, 33]. The existence of optimal strategies for $2^{1/2}$ -player games with parity objectives follows from [26, 5]. This gives the following theorem.

Theorem 2 (Optimal strategies). *For all $2^{1/2}$ -player game graphs with Müller objectives, optimal strategies exist for both players.*

Sure and almost-sure winning strategies. Given an objective Φ , a strategy σ is a *sure winning strategy* for player 1 from a vertex v if for every strategy π of player 2 we have $\text{Outcome}(v, \sigma, \pi) \subseteq \Phi$. Similarly, a strategy σ is an *almost-sure winning strategy* for player 1 from a vertex v for the objective Φ if for every strategy π of player 2 we have $\Pr_v^{\sigma, \pi}(\Phi) = 1$. The sure and almost-sure winning strategies for player 2 are defined analogously. Given an objective Φ ,




Figure 1. An MDP with a reachability objective.

the *sure winning set* $\langle\langle 1 \rangle\rangle_{sure}(\Phi)$ for player 1 is the set of vertices 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 vertices 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 $2^{1/2}$ -player game graphs and all objectives Φ , we have $\langle\langle 1 \rangle\rangle_{sure}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{sure}(\Omega \setminus \Phi) \subseteq \langle\langle 2 \rangle\rangle_{almost}(\Omega \setminus \Phi)$.

Computing sure winning and almost-sure winning sets and strategies is referred to as the *qualitative* analysis of $2^{1/2}$ -player games. It follows from Theorem 2 that $\langle\langle 1 \rangle\rangle_{almost}(\Phi) = \{v \in V \mid \langle\langle 1 \rangle\rangle_{val}(\Phi)(v) = 1\}$. The following result is the classical determinacy result for 2-player deterministic games.

Theorem 3 (Qualitative determinacy [25]). *For all 2-player game graphs and all Borel objectives Φ , we have*

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

The following example shows that Theorem 3 cannot be extended to $1^{1/2}$ -player and $2^{1/2}$ -player games.

Example 1 Consider the MDP with a reachability objective shown in Fig. 1. In all our figures we use $\square$ to denote player-1 vertices, $\diamond$ to denote player-2 vertices, and $\circ$ to denote probabilistic vertices. The objective Φ of player 1 is to reach the vertex v_2 . Given the strategy σ that chooses $v_0 \rightarrow v_1$ at vertex v_0 , the target v_2 is reached with probability 1. However, there is an infinite paths that is consistent with the player-1 strategy σ but only visits the vertices v_0 and v_1 . Hence, $\langle\langle 1 \rangle\rangle_{sure}(\Phi) = \{v_2\}$ and $\langle\langle 1 \rangle\rangle_{almost}(\Phi) = \{v_0, v_1, v_2\}$. This shows that in general for MDPs and $2^{1/2}$ -player games $\langle\langle 1 \rangle\rangle_{sure}(\Phi) \subsetneq \langle\langle 1 \rangle\rangle_{almost}(\Phi)$. ■

3 Special Families of Strategies

Pure, finite-memory, and memoryless strategies. We classify strategies according to their use of randomization and memory. The strategies that do not use randomization and memory are called pure. A player-1 strategy σ is *pure* if for all $w \in V^*$ and $v \in V_1$, there is a vertex $u \in V$ such that

$\sigma(\vec{w} \cdot v)(u) = 1$. 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*.

A strategy is *finite-memory* if it depends on the current vertex and on a finite number of bits from the history of the play so far. 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 *memoryless* strategy does not depend on the history but only on the current vertex. A *memoryless* strategy for player 1 can be represented as function $\sigma: V_1 \rightarrow \mathcal{D}(V)$ such that for all $v \in V_1$ and $u \in V$, if $\sigma(v)(u) > 0$, then $(v, u) \in E$. A *pure memoryless strategy* is a pure strategy that is memoryless. A pure memoryless strategy for player 1 can be represented as a function $\sigma: V_1 \rightarrow V$ such that $(v, \sigma(v)) \in E$ for all $v \in V_1$. 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. These observations will be useful in the analysis of $2^{1/2}$ -player games.

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 vertex $v \in \langle\langle 1 \rangle\rangle_{\text{sure}}(\Phi)$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that for every player-2 strategy $\pi \in \Pi$ we have $\text{Outcome}(v, \sigma, \pi) \subseteq \Phi$;
- *almost-sure winning* if for every game graph $G \in \mathcal{G}$, for every vertex $v \in \langle\langle 1 \rangle\rangle_{\text{almost}}(\Phi)$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that for every player-2 strategy $\pi \in \Pi$ we have $\Pr_v^{\sigma, \pi}(\Phi) = 1$;
- *optimality* if for every game graph $G \in \mathcal{G}$, for every vertex $v \in V$ there is a player-1 strategy $\sigma \in \Sigma^C$ such that $\langle\langle 1 \rangle\rangle_{\text{val}}(\Phi)(v) = \inf_{\pi \in \Pi} \Pr_v^{\sigma, \pi}(\Phi)$.

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

vertices) 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 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 proposition states that randomization is not necessary for sure winning.

Proposition 2 *If a family Σ^C of strategies suffices for sure winning with respect to a Borel objective Φ on all $2^{1/2}$ -player game graphs, then the family $\Sigma^C \cap \Sigma^P$ of pure strategies suffices for sure winning with respect to Φ on all $2^{1/2}$ -player game graphs.*

The following result is the classical determinacy result for 2-player deterministic graph games.

Theorem 4 (Pure and finite-memory strategies).

1. [25] *The family Σ^P of pure strategies suffices for sure winning with respect to all Borel objectives on 2-player game graphs.*
2. [19] *The family Σ^{PF} of pure finite-memory strategies suffices for sure winning with respect to all Müller objectives on 2-player game graphs.*

It is easy to see that for any 2-player game a sure winning strategy is also an almost-sure winning strategy. Hence the almost-sure winning strategies need not be more complex than the sure winning strategies in 2-player games. This is formalized by the proposition below.

Proposition 3 *If a family Σ^C of strategies suffices for sure winning with respect to a Borel objective Φ on all 2-player game graphs, then the family Σ^C suffices also for optimality with respect to Φ on all 2-player game graphs.*

4 Reachability and Safety Objectives

Pure memoryless strategies suffice for sure winning and optimality (and therefore for almost-sure winning) with respect to reachability and safety objectives.

Theorem 5

1. *The family Σ^{PM} of pure memoryless strategies suffices for sure winning with respect to reachability and safety objectives on $2^{1/2}$ -player game graphs.*
2. [7] *The family Σ^{PM} of pure memoryless strategies suffices for optimality with respect to reachability and safety objectives on $2^{1/2}$ -player game graphs.*

Theorem 5(1) for 2-player games is classical. It is an easy consequence of the alternating reachability analysis of AND-OR graphs; see [33] for details. Due to Proposition 1, the result carries over to $2^{1/2}$ -player games. Theorem 5(2) follows from the results of [7]. However, the proof given there is analytical; it analyzes the behavior of discounted games as the discount factor converges to 1. As in the following sections we will make frequent use of this result for MDPs, we provide here an elementary proof that pure memoryless strategies suffice for optimality with respect to reachability objectives on MDPs. Our proof uses only facts from graph theory and matrix algebra.

Consider a player-1 MDP $G = ((V, E), V_0, V_1, V_2, p)$ (where $V_2 = \emptyset$), together with a set $T \subseteq V$ of target vertices. Let $T_1 = T$, and let $T_0 \subseteq V$ be the set of vertices that cannot reach T_1 in the graph (V, E) ; let also $U = V \setminus (T_0 \cup T_1)$. From $T_0 \cup T_1$, all strategies are optimal with respect to the objective $\text{Reach}(T)$, so we can fix a pure memoryless strategy on $T_0 \cup T_1$ arbitrarily. Consider the following fixpoint equation for $x: V \rightarrow [0, 1]$, where for all $v \in V$:

$$x(v) = \begin{cases} 0 & \text{if } v \in T_0; \\ 1 & \text{if } v \in T_1; \\ \max_{u \in E(v)} x(u) & \text{if } v \in V_1 \setminus T; \\ \sum_{u \in E(v)} x(u) \cdot p(v, u) & \text{if } v \in V_0 \setminus T. \end{cases} \quad (1)$$

This system of equations in general has many fixpoints, and it is well-known that the least fixpoint x^* equals $\langle\langle 1 \rangle\rangle_{\text{val}} \text{Reach}(T)$; see, e.g., [13]. For $v \in U \cap V_1$, define the set of *optimal successors* of v by $A(v) = \{u \in E(v) \mid x^*(u) = x^*(v)\}$. Clearly, an optimal strategy must select only optimal successors of vertices in $U \cap V_1$. Thus, we cut from the MDP all the edges $(v, u) \in E$ with $u \in V_1 \cap U$ and $u \notin A(v)$. It is immediate to check that x^* is still a fixpoint of (1) in the resulting MDP.

To obtain a memoryless strategy, we can choose all optimal successors of vertices in $U \cap V_1$ uniformly at random. To obtain a memoryless pure strategy, we observe that if a vertex $v \in U \cap V_1$ has multiple optimal successors, i.e., $|A(v)| > 1$, and we cut an edge (v, u) with $u \in A(v)$, then x^* is still a fixpoint of (1) in the resulting MDP. However, we cannot arbitrarily fix one optimal successor for each vertex in $U \cap V_1$ and cut the edges to all other successors: doing so could create new fixpoints *below* x^* . This occurs, for instance, whenever there are mutually reachable vertices with equal x^* , and the selected successors create a cycle that prevents reaching T . Our goal is to pick optimal successors, and cut the edges to other successors, so that x^* is the *only* fixpoint of (1) in the resulting MDP. This will guarantee that $x^* = \langle\langle 1 \rangle\rangle_{\text{val}} \text{Reach}(T)$ for the resulting pure memoryless strategy.

To ensure the uniqueness of the fixpoint, we cut edges from $V_1 \cap U$ while maintaining the invariant that every ver-

tex in U can reach T_1 in the graph $(V \setminus T_0, E)$. Note that this invariant holds initially by the definition of T_0 . Let $E' \subseteq E$ be a subset of edges that consists of shortest paths from U to T such that every vertex has only one outgoing edge, i.e., for all $v, u_1, u_2 \in V$, if $(v, u_1), (v, u_2) \in E'$, then $u_1 = u_2$. Then, prune from player-1 vertices all edges that are not in E' ; precisely, for all $v \in U \cap V_1$ and $(v, u) \in E$, keep (v, u) if $(v, u) \in E'$, and prune it otherwise. The MDP corresponds thus to a Markov chain. We define the transition probability matrix $[P_{v,u}]_{v,u \in U}$ and the vector $[q_v]_{v \in U}$ as follows, for all $v, u \in U$:

$$P_{v,u} = \begin{cases} 1 & \text{if } v \in V_1 \text{ and } (v, u) \in E; \\ 0 & \text{if } v \in V_1 \text{ and } (v, u) \notin E; \\ p(v, u) & \text{if } v \in V_0; \end{cases}$$

$$q_v = \begin{cases} 1 & \text{if } v \in V_1 \text{ and } \exists u \in T. (v, u) \in E; \\ 0 & \text{if } v \in V_1 \text{ and } \forall u \in T. (v, u) \notin E; \\ \sum_{u \in T} p(v, u) & \text{if } v \in V_0. \end{cases}$$

Then x^* , as a fixpoint of (1), is a solution of $x = Px + q$. Since every vertex in U has a path to T in the graph $(V \setminus T_0, E)$, the matrix P corresponds to a transient chain, and $\det(I - P) \neq 0$ [22]. Hence, $x^* = (I - P)^{-1}q$ is the unique fixpoint of (1), showing the optimality of the pure memoryless strategy thus constructed.

5 Parity Objectives

Pure memoryless strategies suffice for sure winning and optimality (and therefore for almost-sure winning) with respect to parity objectives.

Theorem 6

1. The family Σ^{PM} of pure memoryless strategies suffices for sure winning with respect to parity objectives on $2^{1/2}$ -player game graphs.
2. [26, 5] The family Σ^{PM} of pure memoryless strategies suffices for optimality with respect parity objectives on $2^{1/2}$ -player game graphs.

Theorem 6(1) for 2-player games is a classical result of [16]; an alternative proof is presented in [32]. Due to Proposition 1, the result carries over to $2^{1/2}$ -player games. Theorem 6(2) follows from two independent results: an analytical proof using results on recursive games of Everett [17] is presented in [26]; a combinatorial proof using graph-theoretic arguments is presented in [5].

6 Rabin Objectives

Pure memoryless strategies suffice for sure winning with respect to Rabin objectives in $1^{1/2}$ -player and 2-player

games, and for optimality (and therefore for almost-sure winning) in $1^{1/2}$ -player games (MDPs). It is an open problem whether the family Σ^{PM} of pure memoryless strategies suffices for almost-sure winning on $2^{1/2}$ -player game graphs.

Theorem 7 *The family Σ^{PM} of pure memoryless strategies suffices for sure winning with respect to Rabin objectives on $2^{1/2}$ -player game graphs.*

Theorem 7 for 2-player games is a classical result of [16]; an alternative proof is presented in [15]. Due to Proposition 1, the result carries over to $2^{1/2}$ -player games.

It follows from Theorem 4 and Proposition 3 that the family Σ^{PF} of pure finite-memory strategies suffices for optimality (and almost-sure winning) with respect to Rabin objectives on 2-player game graphs. On the other hand, pure memoryless strategies suffice for optimality with respect to Rabin objectives on MDPs, as stated by the following theorem. This result does not follow from the preceding results, as the case for $2^{1/2}$ -player games is open, as noted above.

Theorem 8 *The family Σ^{PM} of pure memoryless strategies suffices for optimality with respect to Rabin objectives on $1^{1/2}$ -player game graphs.*

This theorem can be proved using the techniques developed in [8, 9] to compute the maximal probability of satisfying an ω -regular specification. We consider player-1 MDPs and hence strategies for player 1. Let $G = ((V, E), V_0, V_1, V_2, p)$ with $V_2 = \emptyset$ be a $1^{1/2}$ -player game graph. The key concept underlying the proof is that of an end-component. A set $U \subseteq V$ of vertices is an *end-component* if U is p -closed and the subgame graph $G \upharpoonright U$ is strongly connected. We denote by $\mathcal{E} \subseteq 2^V$ the set of all end-components of G .

We will use two facts about end-components. The first fact states that, under any strategy, with probability 1 the set of vertices visited infinitely often along a play is an end-component. This theorem parallels the well-known property of *closed recurrent classes* in Markov chains [22]. To state the lemma, for $v \in V$ and $U \subseteq V$, we define $\Omega_v^U = \{\omega \in \Omega_v \mid \text{Inf}(\omega) = U\}$.

Lemma 1 [8] *For all vertices $v \in V$ and strategies $\sigma \in \Sigma$, we have $\Pr_v^\sigma(\bigcup_{U \in \mathcal{E}} \Omega_v^U) = 1$.*

For an end-component $U \in \mathcal{E}$, we denote by ρ_U the randomized memoryless strategy that at each vertex $v \in U \cap V_1$ selects uniformly at random one of the edges $(v, u) \in E$ having $u \in U$. The following lemma is immediate, as U under strategy ρ_U forms a closed recurrent class of a Markov chain.

Lemma 2 [8] *For all end-components $U \in \mathcal{E}$ and all vertices $v \in U$, we have $\Pr_v^{\rho_U}(\Omega_v^U) = 1$.*

Consider a set $R = \{(E_1, F_1), \dots, (E_m, F_m)\}$ of Rabin pairs. For convenience, set $\bar{E}_i = V \setminus E_i$ for $1 \leq i \leq m$. With this notation, the Rabin objective can be read as follows: a play is winning if there is some $1 \leq i \leq m$ such that (1) the play is eventually confined in $\bar{E}_i$, and (2) the play visits F_i infinitely often. We denote by $\mathcal{U} \subseteq \mathcal{E}$ the set consisting of the end-components $U \in \mathcal{E}$ such that there is an $1 \leq i \leq m$ for which $U \subseteq \bar{E}_i$ and $U \cap F_i \neq \emptyset$. The set $\mathcal{U}$ consists thus of the end-components that satisfy the Rabin objective. Let $T_{\text{end}} = \bigcup_{U \in \mathcal{U}} U$ be union of all such winning end-components. From Lemmas 1 and 2 above, it follows that the maximal probability of satisfying $\text{Rabin}(R)$ is equal to the maximal probability of reaching the union of the winning end-components. We present a proof of this fact, as it will be useful in the construction of a pure memoryless strategy.

Lemma 3 [8] $\langle\langle 1 \rangle\rangle_{\text{val}} \text{Rabin}(R) = \langle\langle 1 \rangle\rangle_{\text{val}} \text{Reach}(T_{\text{end}})$.

Proof. Given any strategy σ , let σ' be a strategy that behaves like σ outside of T_{end} , and that upon entering T_{end} at a state v , follows the strategy ρ_U , for some end-component $U \in \mathcal{U}$ with $v \in U$. Then, from Lemma 2 it follows that for all vertices $v \in V$, we have $\Pr_v^\sigma(\text{Reach}(T_{\text{end}})) = \Pr_v^{\sigma'}(\text{Rabin}(R))$, and thus, $\langle\langle 1 \rangle\rangle_{\text{val}} \text{Rabin}(R)(v) \geq \langle\langle 1 \rangle\rangle_{\text{val}} \text{Reach}(T_{\text{end}})(v)$. For the reverse inequality, consider again an arbitrary strategy σ , and notice that from Lemma 1 we have:

$$\begin{aligned} \Pr_v^\sigma(\text{Rabin}(R)) &= \sum_{U \subseteq V} \Pr_v^\sigma(\text{Rabin}(R) \mid \Omega_v^U) \cdot \Pr_v^\sigma(\Omega_v^U) \\ &= \sum_{U \in \mathcal{U}} \Pr_v^\sigma(\text{Rabin}(R) \mid \Omega_v^U) \cdot \Pr_v^\sigma(\Omega_v^U) \\ &\leq \sum_{U \in \mathcal{U}} \Pr_v^\sigma(\Omega_v^U) \leq \Pr_v^\sigma(\text{Reach}(T_{\text{end}})). \end{aligned}$$

As pure memoryless strategies suffice with respect to reachability in MDPs, the above proof is a first step in showing that there are pure memoryless optimal strategies. However, the strategy σ' constructed above is not necessarily pure memoryless, because it needs to remember one of the winning end-components (corresponding to the entrance in T_{end}), and it follows a randomized strategy inside that end-component. We can construct a suitable pure memoryless strategy as follows. Let $\mathcal{U} = \{U_1, \dots, U_k\}$, thus fixing an arbitrary order among the winning end-components. For $1 \leq j \leq k$, let $\text{pair}(j)$ be any fixed $i \in \{1, \dots, m\}$ such that $U_j \subseteq \bar{E}_i$ and $U_j \cap F_i \neq \emptyset$. In other words, $(E_{\text{pair}(j)}, F_{\text{pair}(j)}) \in R$ is a Rabin pair that witnesses the winning of the end-component U_j . With this notation, for $1 \leq j \leq k$ let $\hat{\sigma}_j$ be the pure memoryless strategy defined over U_j which chooses only successors in U_j such that:

- in $U_j \setminus F_{pair(j)}$, it coincides with a pure memoryless strategy for reaching $F_{pair(j)}$;
- in $F_{pair(j)}$, it chooses an arbitrary destination in U_j .

The existence of such a strategy follows from the existence of pure memoryless strategies with respect to reachability (Theorem 5). For $v \in T_{end}$, let $rank(v) = \max\{1 \leq j \leq k \mid v \in U_j\}$ be the rank of the vertex v . Now define the strategy $\hat{\sigma}$ as follows:

- outside T_{end} , the strategy $\hat{\sigma}$ coincides with a pure memoryless optimal strategy with respect to the objective $\text{Reach}(T_{end})$;
- at each vertex $v \in T_{end}$, the strategy $\hat{\sigma}$ coincides with $\hat{\sigma}_{rank(v)}$.

Once such a memoryless strategy is fixed, the MDP becomes a Markov chain $MC_{\hat{\sigma}}$, with transition probabilities defined by $P_{u,v} = \hat{\sigma}(u)(v)$ for $u \in V_1$, and by $P_{u,v} = p(u,v)$ for $u \in V_0$. The following lemma characterizes the closed recurrent classes of this Markov chain in the set T_{end} , stating that they satisfy the Rabin objective.

Lemma 4 *If C is a closed recurrent class of the Markov chain $MC_{\hat{\sigma}}$ with $C \cap T_{end} \neq \emptyset$, then there is an $1 \leq i \leq m$ such that $C \subseteq \overline{E}_i$ and $C \cap F_i \neq \emptyset$.*

Proof. Let $E' = \{(u,v) \in T_{end}^2 \mid P_{u,v} > 0\}$. The closed recurrent classes of $MC_{\hat{\sigma}}$ are the terminal strongly connected components (SCCs) of the graph (T_{end}, E') . By the construction of $\hat{\sigma}$, the rank of the vertices along all paths in (T_{end}, E') is nondecreasing. Hence, each terminal SCC C of (T_{end}, E') must consist of vertices with the same rank; we indicate this rank by $rank(C)$. Then, at all states of C the strategy $\hat{\sigma}_{rank(C)}$ is used. Thus, it immediately follows that $C \subseteq U_{rank(C)}$. Moreover, since from every state of $U_{rank(C)} \setminus F_{pair(rank(C))}$ the strategy $\hat{\sigma}_{rank(C)}$ aims at reaching $F_{pair(rank(C))}$, and as C has no outgoing edges in E' , it follows that $C \cap F_{pair(rank(C))} \neq \emptyset$. ■

The optimality of the strategy $\hat{\sigma}$ is a simple consequence of Lemma 4.

Corollary 1 *For all vertices $v \in V$, we have $\langle\langle 1 \rangle\rangle_{val} \text{Rabin}(R)(v) = \Pr_v^{\hat{\sigma}}(\text{Rabin}(R))$.*

Proof. In view of Lemma 3, we need to show that $\langle\langle 1 \rangle\rangle_{val} \text{Reach}(T_{end})(v) = \Pr_v^{\hat{\sigma}}(\text{Rabin}(R))$. To this end, it suffices to note that by the construction of $\hat{\sigma}$, we have $\langle\langle 1 \rangle\rangle_{val} \text{Reach}(T_{end}) = \Pr_v^{\hat{\sigma}}(\text{Reach}(T_{end}))$ and $\Pr_v^{\hat{\sigma}}(\text{Rabin}(R) \mid \text{Reach}(T_{end})) = 1$. The second equality follows from the fact that under strategy $\hat{\sigma}$, once a play ω enters T_{end} , with probability 1 we have $\text{Inf}(\omega) = C$ for some closed recurrent class C of $MC_{\hat{\sigma}}$. Lemma 4 then leads to the conclusion. ■




Figure 2. A Streett game.

7 Streett Objectives

Sure winning requires memory for Streett objectives even in the case of 1-player games. This follows from the example given in the introduction, which is repeated here.

Example 2 Consider the 1-player game graph shown in Fig. 2. The objective is a Streett objective with two Streett pairs: $S = \{(E_1, F_1), (E_2, F_2)\}$ for $F_1 = F_2 = \{v_0, v_1, v_2\}$ and $E_1 = \{v_1\}$ and $E_2 = \{v_2\}$. We consider the two possible pure memoryless strategies: (1) for the strategy that always chooses $v_0 \rightarrow v_1$, the Streett pair (E_2, F_2) is not satisfied; and (2) for the strategy that always chooses $v_0 \rightarrow v_2$, the Streett pair (E_1, F_1) is not satisfied. Hence there is no pure memoryless sure winning strategy for player 1. It follows from Proposition 2 that there is no randomized memoryless sure winning strategy either. ■

It will follow from Theorem 10 that memoryless strategies suffice for almost-sure winning with respect to Streett objectives on $1^{1/2}$ -player (and hence on 1-player) game graphs. We now show that almost-sure winning 2-player Streett games does require memory.

Example 3 Consider the 2-player game graph shown in Fig. 3. The objective is a Streett objective with two Streett pairs: $S = \{(E_1, F_1), (E_2, F_2)\}$ for $E_1 = \{v_2, v_4\}$, $E_2 = \{v_3\}$, $F_1 = \{v_3\}$, and $F_2 = \{v_4\}$. Consider the two possible pure memoryless strategies for player 1: (1) for the player-1 strategy that always chooses $v_0 \rightarrow v_1$, the player-2 strategy that chooses $v_1 \rightarrow v_3$ ensures that the Streett pair (E_1, F_1) is not satisfied; and (2) for the player-1 strategy that always chooses $v_0 \rightarrow v_4$, the Streett pair (E_2, F_2) is not satisfied. For any randomized memoryless strategy that chooses both $v_0 \rightarrow v_1$ and $v_0 \rightarrow v_4$ with positive probabilities, the player-2 strategy that chooses $v_1 \rightarrow v_2$ ensures that the vertex set $\{v_0, v_1, v_2, v_4\}$ is visited infinitely often. Hence the Streett pair (E_2, F_2) is not satisfied. Note, however, that the pure memory strategy that chooses $v_0 \rightarrow v_4$ once whenever player 2 chooses $v_1 \rightarrow v_3$, and otherwise chooses $v_0 \rightarrow v_1$, is a sure winning strategy (and hence also an almost-sure winning strategy) for player 1. ■

The results on Streett games are summarized in the following theorem.

Theorem 9

1. The family Σ^M of memoryless strategies does not suffice for sure winning with respect to Streett objectives on 1-player game graphs.

2. The family Σ^M of memoryless strategies suffices for almost-sure winning with respect to Streett objectives on $1\frac{1}{2}$ -player game graphs.
3. The family Σ^M of memoryless strategies does not suffice for almost-sure winning with respect to Streett objectives on 2-player game graphs.

8 Müller Objectives

It follows from Example 2 that sure winning strategies require memory for Müller objectives even in 1-player games. Moreover, Example 3 shows that in 2-player games with Müller objectives almost-sure winning requires memory. We now show that for $1\frac{1}{2}$ -player Müller games memoryless almost-sure winning strategies exist.

Theorem 10 *The family Σ^M of memoryless strategies suffices for optimality with respect to Müller objectives on $1\frac{1}{2}$ -player game graphs.*

Given a set $M_\Gamma \subseteq 2^V$ of Müller sets, we denote by $\mathcal{U} = \mathcal{E} \cap M_\Gamma$ the set of end-components that are Müller sets (see Section 6 for a definition of end-components); these are the *winning* end-components. Let $T_{end} = \bigcup_{U \in \mathcal{U}} U$ be their union. From Lemmas 1 and 2, it follows that the maximal probability of satisfying the objective $\text{Müller}(M_\Gamma)$ is equal to the maximal probability of reaching the union of the winning end-components.

Lemma 5 $\langle\langle 1 \rangle\rangle_{val} \text{Müller}(M_\Gamma) = \langle\langle 1 \rangle\rangle_{val} \text{Reach}(T_{end})$.

The proof of this lemma is analogous to the proof of Lemma 3. To construct a memoryless winning strategy, we again let $\mathcal{U} = \{U_1, \dots, U_k\}$, thus fixing an arbitrary order among the winning end-components, and we define the rank of a vertex $v \in T_{end}$ by $\text{rank}(v) = \max\{1 \leq j \leq k \mid v \in U_j\}$. We define a randomized memoryless strategy $\hat{\rho}$ as follows:

- In $V \setminus T_{end}$, the strategy $\hat{\rho}$ coincides with an optimal memoryless strategy to reach T_{end} .
- At each vertex $v \in T_{end} \cap V_1$, the strategy $\hat{\rho}$ coincides with the strategy $\rho_{U_{\text{rank}(v)}}$ defined in Section 6; that is, it selects uniformly at random one of the edges $(v, u) \in E$ having $u \in U_{\text{rank}(v)}$.

Once such a memoryless strategy is fixed, the MDP becomes a Markov chain $MC_{\hat{\rho}}$, with transition probabilities defined by $P_{u,v} = \hat{\rho}(u)(v)$ for $u \in V_1$, and by $P_{u,v} = p(u,v)$ for $u \in V_0$. The following lemma characterizes the closed recurrent classes of this Markov chain in the set T_{end} , stating that they are all winning end-components.

Lemma 6 *If C is a closed recurrent class of the Markov chain $MC_{\hat{\rho}}$, then either $C \cap T_{end} = \emptyset$ or $C \in \mathcal{U}$.*

Proof. Let $E' = \{(u, v) \in T_{end}^2 \mid P_{u,v} > 0\}$. The closed recurrent classes of $MC_{\hat{\rho}}$ are the terminal SCCs of the graph (T_{end}, E') . As the rank of the vertices along all paths in (T_{end}, E') is nondecreasing, each terminal SCC C of (T_{end}, E') must consist of vertices with the same rank, denoted $\text{rank}(C)$. Clearly, $C \subseteq U_{\text{rank}(C)}$. To see that $C = U_{\text{rank}(C)}$ note that in C player 1 follows the strategy $\rho_{U_{\text{rank}(C)}}$, which causes the whole of $U_{\text{rank}(C)}$ to be visited. Hence, as C is terminal, we have $C = U_{\text{rank}(C)}$. ■

The optimality of the strategy $\hat{\rho}$ is a simple consequence of Lemma 6. The following corollary is proved in a fashion analogous to Corollary 1.

Corollary 2 *For all vertices $v \in V$, we have $\langle\langle 1 \rangle\rangle_{val} \text{Müller}(M_\Gamma)(v) = \text{Pr}_v^{\hat{\rho}}(\text{Müller}(M_\Gamma))$.*

9 Upward-closed Objectives

We show that memoryless almost-sure winning strategies exist for all $2\frac{1}{2}$ -player Müller games if the objective can be specified in an upward-closed way.

Theorem 11 *The family Σ^M of memoryless strategies suffices for almost-sure winning on $2\frac{1}{2}$ -player game graphs with respect to Müller objectives that have upward-closed specifications.*

Proof. Consider an upward-closed specification Γ of an objective $\text{Müller}(M_\Gamma)$ and a $2\frac{1}{2}$ -player game graph $G = ((V, E), V_0, V_1, V_2, p)$. Let $W_1 \subseteq V$ be the almost-sure winning set for player 1. It is easy to argue that for every vertex $u \in W_1 \cap V_1$, there is a vertex $v \in W_1$ with $(u, v) \in E$, and for every vertex $u \in W_1 \cap (V_0 \cup V_2)$, for all edges $(u, v) \in E$ we have $v \in W_1$. Hence, $G \upharpoonright W_1$ is a subgame graph. By the definition of W_1 , player 1 has a winning strategy σ_w (memoryless or not) such that $\text{Pr}_v^{\sigma_w, \pi}(\text{Müller}(M_\Gamma)) = 1$ for all vertices $v \in W_1$ and player-2 strategies π . Moreover, the strategy σ_w can choose only edges in $G \upharpoonright W_1$, as it cannot leave W_1 . Hence, from now on we concentrate on the subgame graph $G \upharpoonright W_1$.

Let $\hat{\sigma}$ be the memoryless player-1 strategy that plays uniformly at random in $G \upharpoonright W_1$. Precisely, for a vertex $u \in W_1 \cap V_1$, let $E_u = \{(u, v) \in E \mid v \in W_1\}$, and let $\hat{\sigma}$ be the player-1 strategy that at $u \in W_1 \cap V_1$ plays each edge in E_u uniformly at random. We claim that $\hat{\sigma}$ is winning, that is, $\text{Pr}_v^{\hat{\sigma}, \pi}(\text{Müller}(M_\Gamma)) = 1$ for all vertices $v \in W_1$ and player-2 strategies π , thus showing the existence of a memoryless almost-winning strategy for player 1.

Assume, towards a contradiction, that player 2 has a strategy π_w such that $\text{Pr}_v^{\hat{\sigma}, \pi_w}(\text{Müller}(M_\Gamma)) < 1$ for some vertex $v \in W_1$. Note that $G \upharpoonright W_1$ is a player-2 MDP under strategy $\hat{\sigma}$; we denote this player-2 MDP by $(G \upharpoonright W_1)_{\hat{\sigma}}$. From our results on Müller MDPs, there must be an end-component $A_2 \subseteq W_1$ of $(G \upharpoonright W_1)_{\hat{\sigma}}$ which is winning for




Figure 3. A Müller game.

player 2, that is, $[A_2] \notin \Gamma$. Moreover, player 2 has a memoryless strategy $\hat{\pi}$ that enables it to win with maximal probability in $(G \upharpoonright W_1)_{\hat{\sigma}}$, and A_2 is a closed recurrent class of the Markov chain $(G \upharpoonright W_1)_{\hat{\sigma}, \hat{\pi}}$.

Now consider the situation arising when player 1 uses its original winning strategy σ_w against $\hat{\pi}$. Under strategy $\hat{\pi}$, the game graph $G \upharpoonright W_1$ is a player-1 MDP, which we denote by $(G \upharpoonright W_1)_{\hat{\pi}}$. As A_2 is closed under $\hat{\sigma}$ and $\hat{\pi}$, it has no outgoing player-1 edge in $(G \upharpoonright W_1)_{\hat{\pi}}$. By the definitions of σ_w and A_2 , player 1 can win with probability 1 from A_2 . Therefore, again from our results on Müller MDPs, there must be an end-component $A_1 \subset A_2$ of $(G \upharpoonright W_1)_{\hat{\pi}}$ which is winning for player 1, that is, $[A_1] \in \Gamma$. This contradicts the upward-closure of Γ . ■

There are games with Müller objectives such that sure winning with a pure strategy requires $O(n)$ memory, where n is the number of vertices, but almost-sure winning can be achieved by a randomized memoryless strategy. To see this, for arbitrary $n > 0$, consider the set $C = \{c_1, \dots, c_n\}$ of colors and the Müller specification $\Gamma = \{C\}$. It follows from the split-tree construction of [15] that there is a 2-player game graph G_n with n vertices, each of which is labeled by a unique color from C , such that a pure sure winning strategy on G_n for the objective $\text{Müller}(M_\Gamma)$ requires $O(n)$ memory. On the other hand, since Γ is upward-closed, by Theorem 11 a randomized memoryless almost-sure winning strategy exists.

We now show that there exists a 2-player game graph such that for every Müller objective that is not upward-closed, randomization does not help, i.e., memoryless almost-sure winning strategies exist if pure memoryless almost-sure winning strategies exist, whereas strategies with memory may be almost-sure winning.

Example 4 Let C be a set of colors, and let Γ be a specification of a Müller objective over C which is not upward-closed. Let $\alpha \subset \beta \subseteq C$ such that $\alpha \in \Gamma$ and $\beta \notin \Gamma$ witness that Γ is not upward-closed. Consider the 2-player game graph shown in Fig. 3, where the $\square$ vertices are the player-1 vertices, and the $\diamond$ vertices are the player-2 vertices. The colors of each vertex are defined by $[v_2] = \alpha$, $[v_4] = \beta$, and $[v_0] = [v_1] = [v_3] = \emptyset$.

We show that every memoryless strategy that is not pure is not an almost-sure winning strategy. Consider the randomized memoryless strategy σ for player 1 which plays at

v_0 both edges $v_0 \rightarrow v_1$ and $v_0 \rightarrow v_4$ with positive probability. Let π be the strategy for player 2 which chooses $v_1 \rightarrow v_2$ at v_1 . Given the strategies σ and π , the game is a Markov chain and the vertex set $\{v_0, v_1, v_2, v_4\}$ is a closed recurrent class of the Markov chain; hence it is visited infinitely often. Thus, the set of colors that are visited infinitely often is $\alpha \cup \beta = \beta$, because $\alpha \subseteq \beta$. Since $\beta \notin \Gamma$, there is no randomized memoryless almost-sure winning strategy.

We now show that on the game graph of Fig. 3, for every set $\delta \subseteq C$, if $\beta \subset \delta$ and $\delta \in \Gamma$, then almost-sure winning strategies exist for player 1. The vertex colors are now defined by $[v_2] = \alpha$, $[v_4] = \beta$, $[v_3] = \delta \setminus \beta$, and $[v_0] = [v_1] = \emptyset$. We construct a sure winning strategy (and hence an almost-sure winning strategy) that uses memory. Consider the following strategy $\hat{\sigma}$ for player 1: given any sequence of vertices $\vec{w} \in V^*$, let

$$\hat{\sigma}(\vec{w} \cdot v_0) = \begin{cases} v_1 & \text{if the last vertex of } \vec{w} \text{ is not } v_3; \\ v_4 & \text{otherwise.} \end{cases}$$

Intuitively, the strategy $\hat{\sigma}$ can be described as follows: if at vertex v_1 the edge $v_1 \rightarrow v_2$ is played, then player 1 plays $v_0 \rightarrow v_1$ at v_0 ; if at vertex v_1 the edge $v_1 \rightarrow v_3$ is played, then player 1 chooses $v_0 \rightarrow v_4$ followed by $v_0 \rightarrow v_1$. We prove that $\hat{\sigma}$ is a sure winning strategy for player 1 by considering the following three cases:

1. For every play ω such that $v_1 \rightarrow v_2$ occurs infinitely often and $v_1 \rightarrow v_3$ occurs finitely often, we have $\text{Inf}(\omega) = \{v_0, v_1, v_2\}$ and $[\text{Inf}(\omega)] = \alpha \in \Gamma$.
2. For every play ω such that $v_1 \rightarrow v_3$ occurs infinitely often and $v_1 \rightarrow v_2$ occurs finitely often, we have $\text{Inf}(\omega) = \{v_0, v_1, v_3, v_4\}$ and $[\text{Inf}(\omega)] = \beta \cup (\delta \setminus \beta) = \delta \in \Gamma$.
3. For every play ω such that $v_1 \rightarrow v_3$ occurs infinitely often and $v_1 \rightarrow v_2$ occurs infinitely often, we have $\text{Inf}(\omega) = \{v_0, v_1, v_2, v_3, v_4\}$ and $[\text{Inf}(\omega)] = \alpha \cup \beta \cup (\delta \setminus \beta) = \beta \cup (\delta \setminus \beta) = \delta \in \Gamma$, because $\alpha \subseteq \beta$.

Since $\alpha, \delta \in \Gamma$, it follows that $\hat{\sigma}$ is a sure winning strategy. ■

The following example shows that sure winning may require memory for 1-player games with upward-closed objectives. It follows that Theorem 11 cannot be strengthened to sure winning strategies.

Example 5 Recall the 1-player game graph shown in Fig. 2. The set of colors is $C = \{c_1, c_2\}$, the vertex v_1 is labeled with color c_1 , and v_2 is labeled with c_2 . The specification of the Müller objective is $\Gamma = \{\{c_1, c_2\}\}$; that is, the objective of the player is to visit both v_1 and v_2 infinitely often. We have already seen that there is no pure memoryless sure or almost-sure strategy to achieve this objective. Note,

Table 1. AS - Almost Sure, PM - Pure Memoryless, F - Finite Memory, RM - Randomized Memoryless.

Players	Parity		Rabin		Streett		Müller		Upward-closed	
	Sure	Optimal	Sure	Optimal	Sure	Optimal	Sure	Optimal	Sure	AS
$2^{1/2}$	PM	PM	PM	F	F	F	F	F	F	RM
2	PM	PM	PM	PM	F	F	F	F	F	RM
$1^{1/2}$	PM	PM	PM	PM	F	RM	F	RM	F	RM
1	PM	PM	PM	PM	F	RM	F	RM	F	RM

however, that a strategy that alternately chooses between $v_0 \rightarrow v_1$ and $v_0 \rightarrow v_2$ is a sure winning strategy. Now consider the randomized memoryless strategy σ^M that chooses the edges $v_0 \rightarrow v_1$ and $v_0 \rightarrow v_2$ each with probability $1/2$. Then, with probability 1 all vertices are visited infinitely often. Thus σ^M is an almost-sure winning strategy. ■

10 Conclusion

The memory and randomization requirements of sure winning and optimal (or almost-sure winning) strategies for $2^{1/2}$ -, 2-, $1^{1/2}$ -, and 1-player game graphs are summarized in Table 1. We showed that in $2^{1/2}$ -player games with upward-closed objectives randomized memoryless almost-sure winning strategies exist. Moreover, the randomized memoryless strategies are always simple, in the sense that they use only uniform randomization over given sets of edges. Several important classes of Müller objectives, such as generalized Büchi objectives, are upward-closed. In particular, in 2-player games with generalized Büchi objectives the classical pure sure winning strategies require memory, but randomized memoryless optimal strategies exist.

In the case of $2^{1/2}$ -player games with parity objectives pure memoryless sure winning, almost-sure winning, and optimal strategies exist [4, 5]. It is an open problem whether pure memoryless almost-sure winning strategies exist for $2^{1/2}$ -player games with Rabin objectives. We also leave open the problem whether memoryless *optimal* strategies exist $2^{1/2}$ -player games with upward-closed objectives.

We considered *turn-based* games, where at each (non-probabilistic) vertex one of the two players chooses a successor vertex. A more general class of games are the *concurrent games*, where at each vertex both players simultaneously and independently choose moves, and the combination of the chosen moves results either deterministically or probabilistically in a specific successor vertex. The following results are known for concurrent games [10]: memoryless strategies suffice for optimality with respect to safety objectives; memoryless strategies suffice for optimality with respect to reachability objectives only in the limit; and Büchi objectives require both infinite memory and randomization for almost-sure winning. In the case of concurrent games, sure winning is always simpler than almost-sure

winning, in terms of the requirements of winning strategies. In contrast, for MDPs with Müller objectives sure winning strategies require memory but memoryless strategies suffice for almost-sure winning.

References

- [1] M. Abadi, L. Lamport, and P. Wolper. Relizable and unrelizable concurrent program specifications. In *ICALP'89*, volume 372 of *LNCS*, pages 1–17. Springer, 1989.
- [2] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49:672–713, 2002.
- [3] J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the AMS*, 138:295–311, 1969.
- [4] K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Simple stochastic parity games. In *CSL'03*, volume 2803 of *LNCS*, pages 100–113. Springer, 2003.
- [5] K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA'04*, pages 114–123. SIAM, 2004.
- [6] A. Church. Logic, arithmetic, and automata. In *Proceedings of the International Congress of Mathematicians*, pages 23–35. Institut Mittag-Leffler, 1962.
- [7] A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
- [8] L. de Alfaro. *Formal Verification of Probabilistic Systems*. PhD thesis, Stanford University, 1997.
- [9] L. de Alfaro. Temporal logics for the specification of performance and reliability. In *STACS'97*, volume 1200 of *LNCS*, pages 165–176. Springer, 1997.
- [10] L. de Alfaro and T.A. Henzinger. Concurrent ω -regular games. In *LICS'00*, pages 141–154. IEEE Computer Society, 2000.
- [11] L. de Alfaro and T.A. Henzinger. Interface automata. In *FSE'01*, pages 109–120. ACM, 2001.
- [12] L. de Alfaro, T.A. Henzinger, and F.Y.C. Mang. Detecting errors before reaching them. In *CAV'00*, volume 1855 of *LNCS*, pages 186–201. Springer, 2000.

- [13] L. de Alfaro and R. Majumdar. Quantitative solution of ω -regular games. In *STOC'01*, pages 675–683. ACM, 2001.
- [14] D.L. Dill. *Trace Theory for Automatic Hierarchical Verification of Speed-independent Circuits*. MIT Press, 1989.
- [15] 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, 1997.
- [16] E.A. Emerson and C. Jutla. Tree automata, μ -calculus, and determinacy. In *FOCS'91*, pages 368–377. IEEE Computer Society, 1991.
- [17] H. Everett. Recursive games. In *Contributions to the Theory of Games III*, volume 39 of *Annals of Mathematical Studies*, pages 47–78. Princeton University Press, 1957.
- [18] J.A. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
- [19] Y. Gurevich and L. Harrington. Trees, automata, and games. In *STOC'82*, pages 60–65. ACM, 1982.
- [20] T.A. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Information and Computation*, 173:64–81, 2002.
- [21] A. Kechris. *Classical Descriptive Set Theory*. Springer, 1995.
- [22] J.H. Kemeny. *Finite Markov Chains*. Springer, 1983.
- [23] B. Khousainov and A. Nerode. *Automata Theory and its Applications*. Birkhäuser, 2001.
- [24] D.A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63:1565–1581, 1998.
- [25] D.A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.
- [26] A.K. McIver and C.C. Morgan. Games, probability, and the quantitative μ -calculus. In *LPAR'02*, volume 2514 of *LNAI*, pages 292–310. Springer, 2002.
- [27] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65:149–184, 1993.
- [28] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In *Computation Theory*, volume 208 of *LNCS*, pages 157–168. Springer, 1984.
- [29] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190. ACM, 1989.
- [30] 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.
- [31] 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.
- [32] W. Thomas. On the synthesis of strategies in infinite games. In *STACS'95*, volume 900 of *LNCS*, pages 1–13. Springer, 1995.
- [33] W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, pages 389–455. Springer, 1997.
- [34] M.Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:1–37, 1994.