Stochastic Games with Finitary Objectives*

Krishnendu Chatterjee[1], Thomas A. Henzinger[2], and Florian Horn[3]

[1] Institute of Science and Technology (IST), Austria
[2] EPFL, Switzerland
[3] CWI, Amsterdam

Abstract. The synthesis of a reactive system with respect to an ω-regular specification requires the solution of a graph game. Such games have been extended in two natural ways. First, a game graph can be equipped with probabilistic choices between alternative transitions, thus allowing the modeling of uncertain behavior. These are called stochastic games. Second, a liveness specification can be strengthened to require satisfaction within an unknown but bounded amount of time. These are called finitary objectives. We study, for the first time, the combination of stochastic games and finitary objectives. We characterize the requirements on optimal strategies and provide algorithms for computing the maximal achievable probability of winning stochastic games with finitary parity or Streett objectives. Most notably, the set of states from which a player can win with probability 1 for a finitary parity objective can be computed in polynomial time, even though no polynomial-time algorithm is known in the nonfinitary case.

1 Introduction

The safety and liveness of reactive systems are usually specified by ω-regular sets of infinite words. Then the *reactive synthesis problem* asks for constructing a winning strategy in a graph game with two players and ω-regular objectives: a player that represents the system and tries to satisfy the specification; and a player that represents the environment and tries to violate the specification. In the presence of uncertain or probabilistic behavior, the graph game is stochastic. Such a *stochastic game* is played on a graph with three kinds of vertices: in player-1 vertices, the first player chooses a successor vertex; in player-2 vertices, the second player chooses a successor vertex; and in probabilistic vertices, a successor vertex is chosen according to a given probability distribution. The result of playing the game ad infinitum is a random walk through the graph. If player 1 has an ω-regular objective ϕ, then she tries to maximize the probability that the infinite path that results from the random walk lies inside the set ϕ. Conversely, player 2 tries to minimize that probability. Since the stochastic games are Borel determined [15], and the ω-regular languages are Borel sets, these games have

* This research was supported in part by the Swiss National Science Foundation under the Indo-Swiss Joint Research Programme, by the European Network of Excellence on Embedded Systems Design (ArtistDesign), and by the European project Combest.

R. Královič and D. Niwiński (Eds.): MFCS 2009, LNCS 5734, pp. 34–54, 2009.

a unique value, i.e., there is a real $v \in [0, 1]$ such that player 1 can ensure ϕ with probability arbitrarily close to v, and at the same time, player 2 can ensure $\neg\phi$ with probability arbitrarily close to $1 - v$. The computation of v is referred to as the *quantitative value problem* for stochastic games; the decision problem of whether $v = 1$ is referred to as the *qualitative value problem*. In the case of parity objectives, both value problems lie in NP $\cap$ coNP [6], but no polynomial-time solutions are known even if there are no probabilistic vertices. The NP $\cap$ coNP characterization results from the existence of pure (i.e., nonrandomized) positional (i.e., memoryless) optimal strategies for both players. In the case of Streett objectives, optimal player-1 strategies may require memory, and both value problems are coNP-complete [3], which is again the same in the absence of probabilistic vertices.

The specification of liveness for a reactive system by ω-regular sets such as parity or Streett languages has the drawback that, while the synthesized system is guaranteed to be live, we cannot put any bound on its liveness behavior. For example, the liveness objective $\Box(r \rightarrow \Diamond q)$ ensures that every request r issued by the environment is eventually followed by a response q of the synthesized system, but the delay between each request and corresponding response may grow without bound from one request to the next. This is an undesirable behavior, especially in synthesis, where one controls the system to be built and where one would like stronger guarantees. At the same time, it may be impossible to put a fixed bound on the desired response time, because the achievable bound usually is not known. For this reason, the time-scale independent notion of *finitary objectives* was introduced [1]. The finitary version of the liveness objective $\Box(r \rightarrow \Diamond q)$ requires that there exists an unknown bound b such that every request r is followed by a response q within b steps. The synthesized system can have any response time, but its response time will not grow from one request to the next without bound. Finitary versions can be defined for both parity and Streett (strong fairness) objectives [4]. It should be noted that finitary objectives are not ω-regular. While in games with ω-regular objectives, both players have finite-memory strategies, to violate a finitary objective, player 2 may require infinite memory even if there are no probabilistic vertices [4]. Nonetheless, finitary objectives are Borel sets, and thus have well-defined values in stochastic games.

Nonstochastic games with finitary parity and Streett objectives were first studied in [4], and the results of [4] were later significantly improved upon by [12]. This work showed that finitary objectives are not only more desirable for synthesis, but also can be far less costly than their infinitary counterparts. In particular, nonstochastic games with finitary parity objectives can be solved in polynomial time. In the present paper, we study for the first time *stochastic* games with *finitary* objectives. As main results, we show that the qualitative value problem for finitary parity objectives remains polynomial in the stochastic case, and the quantitative value problem can be solved in NP $\cap$ coNP. For stochastic games with finitary Streett objectives, we compute values in exponential time. Yet also here we achieve a significant improvement by solving the qualitative value problem with an exponential term of 2^d (where d is the number of Streett pairs) instead of $n^d \cdot d!$ (where n is the number of vertices), which characterizes the best known algorithm for nonstochastic games with infinitary Streett objectives. Our results follow the pattern of extending properties of stochastic games with infinitary parity and Streett objectives to stochastic games with finitary parity and Streett objectives. However, in the finitary case, the proof techniques are more complicated, because we need to consider infinite-memory strategies.

We now summarize our results in more detail and draw precise comparisons with the two simpler cases of (i) stochastic games with infinitary (rather than finitary) objectives and (ii) nonstochastic (rather than stochastic) games with finitary objectives.

Comparison of finitary and infinitary parity objectives. In case of parity objectives, pure memoryless optimal strategies exist for both players in both nonstochastic (2-player) game graphs [9] and stochastic ($2\,1/2$-player) game graphs [6,17]. For finitary parity objectives on 2-player game graphs, a pure memoryless optimal strategy exists for the player with the finitary parity objective, while the optimal strategy of the other player (with the complementary objective) in general requires infinite memory [5]. We show in this work that the same class of strategies that suffices in 2-player game graphs also suffices for optimality in $2\,1/2$-player game graphs for finitary parity objectives and their complements. The best known complexity bound for 2- and $2\,1/2$-player games with parity objectives is NP $\cap$ coNP [9,6]. In case of $2\,1/2$-player games, the best known complexity bound for the qualitative analysis is also NP $\cap$ coNP. The solution of 2-player game graphs with finitary parity objectives can be achieved in polynomial time (in $O(n^2 \cdot m)$ time [12,5] for game graphs with n states and m edges). In this work we show that the quantitative analysis of $2\,1/2$-player game graphs with finitary parity objectives lies in NP $\cap$ coNP, and the qualitative analysis can be done in $O(n^4 \cdot m)$ time. To obtain a polynomial time solution for the quantitative analysis of $2\,1/2$-player game graphs with finitary parity objectives, one must obtain a polynomial-time solution for the quantitative analysis of $2\,1/2$-player game graphs with Büchi objectives (which is a major open problem).

Comparison of finitary and infinitary Streett objectives. In case of Streett objectives with d pairs, strategies with $d!$ memory is necessary and sufficient for both 2-player game graphs and $2\,1/2$-player game graphs, and for the complementary player pure memoryless optimal strategies exist [8,11,3]. For finitary Streett objectives on 2-player game graphs, an optimal strategy with $d \cdot 2^d$ memory exists for the player with the finitary Streett objective, while the optimal strategy of the other player (with the complementary objective) in general requires infinite memory [5]. We show that the same class of strategies that suffices for 2-player game graphs also suffices for optimality in $2\,1/2$-player game graphs for finitary Streett objectives and their complements. The decision problems for 2- and $2\,1/2$-player games with Streett objectives are coNP-complete. The solution of 2-player game graphs with finitary Streett objectives can be achieved in EXPTIME. In this work we show that both the qualitative and quantitative analysis of $2\,1/2$-player game graphs with finitary Streett objectives can be achieved in EXPTIME. The best known algorithm for 2-player game graphs with Streett objectives is $O(n^d \cdot d!)$ [16], where as in case of $2^1\!/\!_2$-player game graphs with finitary Streett objectives, we show that the qualitative analysis can be achieved in time $O(n^4 \cdot m \cdot d \cdot 2^d)$. For the quantitative analysis, we present our results for the more general class of *tail* (i.e., prefix-independent) objectives, and obtain the results for finitary parity and Streett objectives as a special case.

2 Definitions

We consider several classes of turn-based games, namely, two-player turn-based probabilistic games ($2^1\!/\!_2$-player games), two-player turn-based deterministic games (2-player games), and Markov decision processes ($1^1\!/\!_2$-player games).

Notation. For a finite set A, a *probability distribution* on A is a function $\delta \colon A \to [0,1]$ such that $\sum_{a \in A} \delta(a) = 1$. We denote the set of probability distributions on A by $\mathcal{D}(A)$. Given a distribution $\delta \in \mathcal{D}(A)$, we denote by $\mathrm{Supp}(\delta) = \{x \in A \mid \delta(x) > 0\}$ the *support* of δ.

Game graphs. A *turn-based probabilistic game graph* (*$2^1\!/\!_2$-player game graph*) $G = ((S,E),(S_1,S_2,S_\bigcirc),\delta)$ consists of a directed graph (S,E), a partition (S_1, S_2, $S_\bigcirc$) of the finite set S of states, and a probabilistic transition function δ: $S_\bigcirc \to \mathcal{D}(S)$, where $\mathcal{D}(S)$ denotes the set of probability distributions over the state space S. The states in S_1 are the *player*-1 states, where player 1 decides the successor state; the states in S_2 are the *player*-2 states, where player 2 decides the successor state; and the states in $S_\bigcirc$ are the *probabilistic* states, where the successor state is chosen according to the probabilistic transition function δ. We assume that for $s \in S_\bigcirc$ and $t \in S$, we have $(s,t) \in E$ iff $\delta(s)(t) > 0$, and we often write $\delta(s,t)$ for $\delta(s)(t)$. For technical convenience we assume that every state in the graph (S,E) has at least one outgoing edge. For a state $s \in S$, we write $E(s)$ to denote the set $\{t \in S \mid (s,t) \in E\}$ of possible successors. The size of a game graph $G = ((S,E),(S_1,S_2,S_\bigcirc),\delta)$ is

$$|G| = |S| + |E| + \sum_{t \in S} \sum_{s \in S_\bigcirc} |\delta(s)(t)|;$$

where $|\delta(s)(t)|$ denotes the space to represent the transition probability $\delta(s)(t)$ in binary.

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

The *turn-based deterministic game graphs* (*2-player game graphs*) are the special case of the $2^1\!/\!_2$-player game graphs with $S_\bigcirc = \emptyset$. The *Markov decision processes* (*$1^1\!/\!_2$-player game graphs*) are the special case of the $2^1\!/\!_2$-player game graphs with $S_1 = \emptyset$ or $S_2 = \emptyset$. We refer to the MDPs with $S_2 = \emptyset$ as *player*-1 *MDPs*, and to the MDPs with $S_1 = \emptyset$ as *player*-2 *MDPs*.

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

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

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

We classify strategies according to their use of randomization and memory. The strategies that do not use randomization are called pure. A player-1 strategy σ is *pure* if for all $\boldsymbol{w} \in S^*$ and $s \in S_1$, there is a state $t \in S$ such that $\sigma(\boldsymbol{w} \cdot s)(t) = 1$. We denote by $\Sigma^P \subseteq \Sigma$ the set of pure strategies for player 1. A strategy that is not necessarily pure is called *randomized*. Let $\mathtt{M}$ be a set called *memory*, that is, $\mathtt{M}$ is a set of memory elements. A player-1 strategy σ can be described as a pair of functions $\sigma = (\sigma_u, \sigma_m)$: a *memory-update* function σ_u: $S \times \mathtt{M} \to \mathtt{M}$ and a *next-move* function $\sigma_m\colon S_1 \times \mathtt{M} \to \mathcal{D}(S)$. We can think of strategies with memory as input/output automaton computing the strategies (see [8] for details). A strategy $\sigma = (\sigma_u, \sigma_m)$ is *finite-memory* if the memory $\mathtt{M}$ is finite, and then the size of the strategy σ, denoted as $|\sigma|$, is the size of its memory $\mathtt{M}$, i.e., $|\sigma| = |\mathtt{M}|$. We denote by Σ^F the set of finite-memory strategies for player 1, and by Σ^{PF} the set of *pure finite-memory* strategies; that is, $\Sigma^{PF} = \Sigma^P \cap \Sigma^F$. The strategy (σ_u, σ_m) is *memoryless* if $|\mathtt{M}| = 1$; that is, the next move does not depend on the history of the play but only on the current state. A memoryless player-1 strategy can be represented as a function $\sigma\colon S_1 \to \mathcal{D}(S)$. A *pure memoryless strategy* is a pure strategy that is memoryless. A pure memoryless strategy for player 1 can be represented as a function $\sigma\colon S_1 \to S$. We denote by Σ^M the set of memoryless strategies for player 1, and by Σ^{PM} the set of pure memoryless strategies; that is, $\Sigma^{PM} = \Sigma^P \cap \Sigma^M$. Analogously we define the corresponding strategy families Π^P, Π^F, Π^{PF}, Π^M, and Π^{PM} for player 2.

Counting strategies. We call an infinite memory strategy σ *finite-memory counting* if there is a finite-memory strategy σ' such that for all $j \geq 0$ there exists $k \leq j$ such that the following condition hold: for all $w \in S^*$ such that $|w| = j$ and for all $s \in S_1$ we have $\sigma(w \cdot s) = \sigma'(\text{suffix}(w,k) \cdot s)$, where for $w \in S^*$ of length j and $k \leq j$ we denote by $\text{suffix}(w,k)$ the suffix of w of length k. In other words, the strategy σ repeatedly plays the finite-memory strategy σ' in different segments of the play and the switch of the strategy in different segments only depends on the length of the play. We denote by $\text{nocount}(|\sigma|)$ the size of the memory of the finite-memory strategy σ' (the memory that is used not for counting), i.e., $\text{nocount}(|\sigma|) = |\sigma'|$. We use similar notations for player 2 strategies.

Objectives. An *objective* for a player consists of a Borel set of *winning plays* $\Phi \subseteq \Omega$. In this paper we consider ω-regular objectives, and *finitary parity* and *finitary Streett* objectives (all the objectives we consider in this paper are Borel objectives).

Classical ω-regular objectives. We first present the definitions of various canonical forms of ω-regular objectives and sub-classes of ω-regular objectives. For a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$, let $\text{Inf}(\omega)$ be the set $\{s \in S \mid s = s_k \text{ for infinitely many } k \geq 0\}$ of states that appear infinitely often in ω.

1. *Reachability and safety objectives.* Given a set $F \subseteq S$ of states, the reachability objective $\text{Reach}(F)$ requires that some state in F be visited, and dually, the safety objective $\text{Safe}(F)$ requires that only states in F be visited. Formally, the sets of winning plays are $\text{Reach}(F) = \{\langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid \exists k \geq 0.\ s_k \in F\}$ and $\text{Safe}(F) = \{\langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid \forall k \geq 0.\ s_k \in F\}$.
2. *Büchi and co-Büchi objectives.* Given a set $F \subseteq S$ of states, the Büchi objective $\text{Buchi}(F)$ requires that some state in F be visited infinitely often, and dually, the co-Büchi objective $\text{coBuchi}(F)$ requires that only states in F be visited infinitely often. Thus, the sets of winning plays are $\text{Buchi}(F) = \{\omega \in \Omega \mid \text{Inf}(\omega) \cap F \neq \emptyset\}$ and $\text{coBuchi}(F) = \{\omega \in \Omega \mid \text{Inf}(\omega) \subseteq F\}$.
3. *Rabin and Streett objectives.* Given a set $P = \{(E_1, F_1), \ldots, (E_d, F_d)\}$ of pairs of sets of states (i.e, for all $1 \leq j \leq d$, both $E_j \subseteq S$ and $F_j \subseteq S$), the Rabin objective $\text{Rabin}(P)$ requires that for some pair $1 \leq j \leq d$, all states in E_j be visited finitely often, and some state in F_j be visited infinitely often. Hence, the winning plays are $\text{Rabin}(P) = \{\omega \in \Omega \mid \exists 1 \leq j \leq d.\ (\text{Inf}(\omega) \cap E_j = \emptyset \text{ and } \text{Inf}(\omega) \cap F_j \neq \emptyset)\}$. Dually, given $P = \{(E_1, F_1), \ldots, (E_d, F_d)\}$, the Streett objective $\text{Streett}(P)$ requires that for all pairs $1 \leq j \leq d$, if some state in F_j is visited infinitely often, then some state in E_j be visited infinitely often, i.e., $\text{Streett}(P) = \{\omega \in \Omega \mid \forall 1 \leq j \leq d.\ (\text{Inf}(\omega) \cap E_j \neq \emptyset \text{ or } \text{Inf}(\omega) \cap F_j = \emptyset)\}$.
4. *Parity objectives.* Given a function $p\colon S \to \{0, 1, 2, \ldots, d-1\}$ that maps every state to an integer *priority*, the parity objective $\text{Parity}(p)$ requires that of the states that are visited infinitely often, the least priority be even. Formally,

the set of winning plays is $\text{Parity}(p) = \{\omega \in \Omega \mid \min\{p(\text{Inf}(\omega))\} \text{ is even}\}$. The dual, co-parity objective has the set $\text{coParity}(p) = \{\omega \in \Omega \mid \min\{p(\text{Inf}(\omega))\} \text{ is odd}\}$ of winning plays. Parity objectives are closed under complementation: given a function $p : S \to \{0, 1, \ldots, d-1\}$, consider the function $p+1 : S \to \{1, 2, \ldots, d\}$ defined as $p+1(s) = p(s)+1$, for all $s \in S$, and then we have $\text{Parity}(p+1) = \text{coParity}(p)$.

Every parity objective is both a Rabin objective and a Streett objective. The Büchi and co-Büchi objectives are special cases of parity objectives with two priorities, namely, p: $S \to \{0, 1\}$ for Büchi objectives with $F = p^{-1}(0)$, and p: $S \to \{1, 2\}$ for co-Büchi objectives with $F = p^{-1}(2)$. The reachability and safety objectives can be turned into Büchi and co-Büchi objectives, respectively, on slightly modified game graphs.

Finitary objectives. We now define a stronger notion of winning, namely, *finitary winning*, in games with parity and Streett objectives.

Finitary winning for parity objectives. For parity objectives, the finitary winning notion requires that for each visit to an odd priority that is visited infinitely often, the distance to a stronger (i.e., lower) even priority be bounded. To define the winning plays formally, we need the concept of a distance sequence.

Distance sequences for parity objectives. Given a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ and a priority function p: $S \to \{0, 1, \ldots, d-1\}$, we define a sequence of distances $dist_k(\omega, p)$, for all $k \geq 0$, as follows:

$$dist_k(\omega, p) = \begin{cases} 0 & \text{if } p(s_k) \text{ is even;} \\ \inf\{k' \geq k \mid p(s_{k'}) \text{ is even and } p(s_{k'}) < p(s_k)\} & \text{if } p(s_k) \text{ is odd.} \end{cases}$$

Intuitively, the distance for a position k in a play with an odd priority at position k, denotes the shortest distance to a stronger even priority in the play. We assume the standard convention that the infimum of the empty set is ∞.

Finitary parity objectives. The finitary parity objective $\text{finParity}(p)$ for a priority function p requires that the sequence of distances for the positions with odd priorities that occur infinitely often be bounded. This is equivalent to requiring that the sequence of all distances be bounded in the limit, and captures the notion that the "good" (even) priorities that appear infinitely often do not appear infinitely rarely. Formally, the sets of winning plays for the finitary parity objective and its complement are $\text{finParity}(p) = \{\omega \in \Omega \mid \limsup_{k\to\infty} dist_k(\omega, p) < \infty\}$ and $\text{cofinParity}(p) = \{\omega \in \Omega \mid \limsup_{k\to\infty} dist_k(\omega, p) = \infty\}$, respectively. Observe that if a play ω is winning for a co-parity objective, then the lim sup of the distance sequence for ω is ∞, that is, $\text{coParity}(p) \subseteq \text{cofinParity}(p)$. However, if a play ω is winning for a (classical) parity objective, then the lim sup of the distance sequence for ω can be ∞ (as shown in Example 1), that is, $\text{finParity}(p) \subsetneq \text{Parity}(p)$.

Example 1. Consider the game shown in Figure 1. The square-shaped states are player 1 states, where player 1 chooses the successor state, and the diamond-shaped states are player 2 states (we will follow this convention throughout this

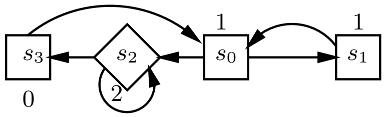

Fig. 1. A simple game graph

paper). The priorities of states are shown next to each state in the figure. If player 1 follows a memoryless strategy σ that chooses the successor s_2 at state s_0, this ensures that against all strategies π for player 2, the minimum priority of the states that are visited infinitely often is even (either state s_3 is visited infinitely often, or both states s_0 and s_1 are visited finitely often). However, consider the strategy π_w for player 2: the strategy π_w is played in rounds, and in round $k \geq 0$, whenever player 1 chooses the successor s_2 at state s_0, player 2 stays in state s_2 for k transitions, and then goes to state s_3 and proceeds to round $k+1$. The strategy π_w ensures that for all strategies σ for player 1, either the minimum priority visited infinitely often is 1 (i.e., both states s_0 and s_1 are visited infinitely often and state s_3 is visited finitely often); or states of priority 1 are visited infinitely often, and the distances between visits to states of priority 1 and subsequent visits to states of priority 0 increase without bound (i.e., the limit of the distances is ∞). Hence it follows that in this game, although player 1 can win for the parity objective, she cannot win for the finitary parity objective. ∎

Finitary winning for Streett objectives. The notion of distance sequence for parity objectives has a natural extension to Streett objectives.

Distance sequences for Streett objectives. Given a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ and a set $P = \{(E_1, F_1), \ldots, (E_d, F_d)\}$ of Streett pairs of state sets, the d sequences of distances $dist_k^j(\omega, P)$, for all $k \geq 0$ and $1 \leq j \leq d$, are defined as follows:

$$dist_k^j(\omega, P) = \begin{cases} 0 & \text{if } s_k \notin F_j; \\ \inf\{k' \geq k \mid s_{k'} \in E_j\} & \text{if } s_k \in F_j. \end{cases}$$

Let $dist_k(\omega, P) = \max\{dist_k^j(\omega, P) \mid 1 \leq j \leq d\}$ for all $k \geq 0$.

Finitary Streett objectives. The finitary Streett objective $\text{finStreett}(P)$ for a set P of Streett pairs requires that the distance sequence be bounded in the limit, i.e., the winning plays are $\text{finStreett}(P) = \{\omega \in \Omega \mid \limsup_{k\to\infty} dist_k(\omega, P) < \infty\}$. We use the following notations for the complementary objective: $\text{cofinStreett}(P) = \Omega \setminus \text{finStreett}(P)$.

Tail objectives. An objective Φ is a *tail* objective if the objective is independent of finite prefixes. Formally, an objective Φ is a tail objective if for all $\omega \in \Omega$, we have $\omega \in \Phi$ iff for all ω' obtained by adding or deleting a finite prefix with ω we have $\omega' \in \Phi$ (see [2] for details). The parity, Streett, finitary parity and finitary Streett are independent of finite prefixes and are all tail objectives. Since tail objectives are closed under complementation, it follows that the complementary objectives to finitary parity and Streett are tail objectives as well.

Sure, almost-sure, positive winning, and optimality. Given a player-1 objective Φ, a strategy $\sigma \in \Sigma$ is *sure winning* for player 1 from a state $s \in S$ if for every strategy $\pi \in \Pi$ for player 2, we have $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Phi$. A strategy σ is *almost-sure winning* for player 1 from the state s for the objective Φ if for every player-2 strategy π, we have $\mathrm{Pr}_s^{\sigma,\pi}(\Phi) = 1$. A strategy σ is *positive winning* for player 1 from the state s for the objective Φ if for every player-2 strategy π, we have $\mathrm{Pr}_s^{\sigma,\pi}(\Phi) > 0$. The sure, almost-sure and positive winning strategies for player 2 are defined analogously. Given an objective Φ, the *sure winning set* $\langle\langle 1 \rangle\rangle_{sure}(\Phi)$ for player 1 is the set of states from which player 1 has a sure winning strategy. Similarly, the *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ and the *positive winning set* $\langle\langle 1 \rangle\rangle_{pos}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning and a positive winning strategy, respectively. The sure winning set $\langle\langle 2 \rangle\rangle_{sure}(\Omega \setminus \Phi)$, the almost-sure winning set $\langle\langle 2 \rangle\rangle_{almost}(\Omega \setminus \Phi)$, and the positive winning set $\langle\langle 2 \rangle\rangle_{pos}(\Omega \setminus \Phi)$ for player 2 are defined analogously. It follows from the definitions that for all $2^1\!/_2$-player game graphs and all objectives Φ, we have $\langle\langle 1 \rangle\rangle_{sure}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{pos}(\Phi)$. Computing sure, almost-sure and positive winning sets and strategies is referred to as the *qualitative* analysis of $2^1\!/_2$-player games [7].

Given objectives $\Phi \subseteq \Omega$ for player 1 and $\Omega \setminus \Phi$ for player 2, we define the *value* functions $\langle\langle 1 \rangle\rangle_{val}$ and $\langle\langle 2 \rangle\rangle_{val}$ for the players 1 and 2, respectively, as the following functions from the state space S to the interval $[0, 1]$ of reals: for all states $s \in S$, let $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \mathrm{Pr}_s^{\sigma,\pi}(\Omega \setminus \Phi)$. In other words, the value $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$ gives the maximal probability with which player 1 can achieve her objective Φ from state s, and analogously for player 2. The strategies that achieve the value are called optimal: a strategy σ for player 1 is *optimal* from the state s for the objective Φ if $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\Phi)$. The optimal strategies for player 2 are defined analogously. Computing values and optimal strategies is referred to as the *quantitative* analysis of $2^1\!/_2$-player games. The set of states with value 1 is called the *limit-sure winning set* [7]. For $2^1\!/_2$-player game graphs with ω-regular objectives the almost-sure and limit-sure winning sets coincide [3].

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

$\sigma \in \Sigma^{\mathcal{C}}$ such that $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) = \inf_{\pi\in\Pi} \Pr_s^{\sigma,\pi}(\Phi)$. The notion of sufficiency for size of finite-memory strategies is obtained by referring to the size of the memory $\mathtt{M}$ of the strategies. The notions of sufficiency of strategies for player 2 is defined analogously.

Determinacy. For sure winning, the $1^1\!/_2$-player and $2^1\!/_2$-player games coincide with 2-player (deterministic) games where the random player (who chooses the successor at the probabilistic states) is interpreted as an adversary, i.e., as player 2. We present the result formally as a Lemma. We use the following notation: given a $2^1\!/_2$-player game graph $G = ((S,E),(S_1,S_2,S_{\bigcirc}),\delta)$, we denote by $\widehat{G} = Tr_2(G)$ the 2-player game graph defined as follows: $\widehat{G} = ((S,E),(S_1,S_2 \cup S_{\bigcirc}))$.

Lemma 1. *For all $2^1\!/_2$-player game graphs, for all Borel objectives Φ, the sure winning sets for objective Φ for player 1 in the game graphs G and $Tr_2(G)$ coincide.*

Theorem 1 and Theorem 2 state the classical determinacy results for 2-player and $2^1\!/_2$-player game graphs with Borel objectives. It follows from Theorem 2 that for all Borel objectives Φ, for all $\varepsilon > 0$, there exists an ε-optimal strategy σ_ε for player 1 such that for all π and all $s \in S$ we have $\Pr_s^{\sigma,\pi}(\Phi) \geq \langle\langle 1\rangle\rangle_{val}(\Phi)(s) - \varepsilon$.

Theorem 1 (Qualitative determinacy). *The following assertions hold.*

1. *For all 2-player game graphs with state set S, and for all Borel objectives Φ, we have $\langle\langle 1\rangle\rangle_{sure}(\Phi) = S \setminus \langle\langle 2\rangle\rangle_{sure}(\overline{\Phi})$, i.e., the sure winning sets for the two players form a partition of the state space [14].*
2. *The family of pure memoryless strategies suffices for sure winning with respect to Rabin objectives for 2-player game graphs [9]; and the family of pure finite-memory strategies suffices for sure winning with respect to Streett objectives for $2^1\!/_2$-player game graphs [10], and sure winning strategies for Streett objectives in general require memory.*

Theorem 2 (Quantitative determinacy). *The following assertions hold.*

1. *For all $2^1\!/_2$-player game graphs, for all Borel objectives Φ, and for all states s, we have $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) + \langle\langle 2\rangle\rangle_{val}(\overline{\Phi})(s) = 1$ [15].*
2. *The family of pure memoryless strategies suffices for optimality with respect to Rabin objectives for $2^1\!/_2$-player game graphs [3]; and the family of pure finite-memory strategies suffices for optimality with respect to Streett objectives for $2^1\!/_2$-player game graphs [3], and optimal strategies for Streett objectives in general require memory.*

We now present the main results of 2-player games with finitary parity and Streett objectives.

Theorem 3 (Finitary parity games [12,5]). *For all 2-player game graphs with n states and m edges, and all priority functions p the following assertions hold.*

1. *The family of pure memoryless strategies suffices for sure winning with respect to finitary parity objectives. There exist infinite-memory winning strategies* π *for player 2 for the objective* $\text{cofinParity}(p)$ *such that* π *is finite-memory counting with* $\text{nocount}(|\pi|) = 2$. *In general no finite-memory winning strategies exist for player 2 for the objective* $\text{cofinParity}(p)$.
2. *The sure winning sets* $\langle\langle 1\rangle\rangle_{sure}(\text{finParity}(p))$ *and* $\langle\langle 2\rangle\rangle_{sure}(\text{cofinParity}(p))$ *can be computed in* $O(n^2 \cdot m)$ *time.*

Theorem 4 (Finitary Streett games [12,5]). *For all 2-player game graphs with* n *states and* m *edges, and for all sets* $P = \{(E_1, F_1), \ldots, (E_d, F_d)\}$ *with* d *Streett pairs, the following assertions hold.*

1. *There exist finite-memory winning strategies* σ *for player 1 for the objective* $\text{finStreett}(P)$ *such that* $|\sigma| = d \cdot 2^d$. *In general winning strategies for player 1 for the objective* $\text{finStreett}(P)$ *require* $2^{\lfloor \frac{d}{2} \rfloor}$ *memory. There exist infinite-memory winning strategies* π *for player 2 for the objective* $\text{cofinStreett}(P)$ *such that* π *is finite-memory counting with* $\text{nocount}(|\pi|) = d \cdot 2^d$. *In general no finite-memory winning strategies exist for player 2 for the objective* $\text{cofinStreett}(P)$.
2. *The sure winning sets* $\langle\langle 1\rangle\rangle_{sure}(\text{finStreett}(P))$ *and* $\langle\langle 2\rangle\rangle_{sure}(\text{cofinStreett}(P))$ *can be computed in* $O(n^2 \cdot m \cdot d^2 \cdot 4^d)$ *time.*

Remark 1. Recall that Büchi and co-Büchi objectives correspond to parity objectives with two priorities. A finitary Büchi objective is in general a strict subset of the corresponding classical Büchi objective; a finitary co-Büchi objective coincides with the corresponding classical co-Büchi objective. However, it can be shown that for parity objectives with two priorities, the value functions for the classical parity objectives and the finitary parity objectives are the same; that is, for all $2^1\!/\!_2$-player game graphs G and all priority functions p with two priorities, we have $\langle\langle 1\rangle\rangle_{val}(\text{finParity}(p)) = \langle\langle 1\rangle\rangle_{val}(\text{Parity}(p))$ and $\langle\langle 2\rangle\rangle_{val}(\text{cofinParity}(p)) = \langle\langle 2\rangle\rangle_{val}(\text{coParity}(p))$. Note that in Example 1, we have $s_0 \in \langle\langle 1\rangle\rangle_{sure}(\text{Parity}(p))$ and $s_0 \notin \langle\langle 1\rangle\rangle_{sure}(\text{finParity}(p))$. This shows that for priority functions with three or more priorities, the sure winning set for a finitary parity objective can be a strict subset of the sure winning set for the corresponding classical parity objective on 2-player game graphs, that is, $\langle\langle 1\rangle\rangle_{sure}(\text{finParity}(p)) \subsetneq \langle\langle 1\rangle\rangle_{sure}(\text{Parity}(p))$, and in general for $2^1\!/\!_2$-player game graphs we have $\langle\langle 1\rangle\rangle_{val}(\text{finParity}(p)) \leq \langle\langle 1\rangle\rangle_{val}(\text{Parity}(p))$. ∎

3 Qualitative Analysis of Stochastic Finitary Games

In this section we present algorithms for qualitative analysis of $2^1\!/\!_2$-player games with finitary parity and finitary Streett objectives. We first present a few key lemmas that would be useful to prove the correctness of the algorithms.

Lemma 2. *Let* G *be a* $2^1\!/\!_2$*-player game graph with the set* S *of states, and let* $P = \{(E_1, F_1), (E_2, F_2), \ldots, (E_d, F_d)\}$ *be a set of* d *Streett pairs. If* $\langle\langle 1\rangle\rangle_{sure}(\text{finStreett}(F)) = \emptyset$, *then the following assertions hold:*

1. *$\langle\langle 2\rangle\rangle_{almost}(\text{cofinStreett}(F)) = S$; and*
2. *there is an almost-sure winning strategy π for player 2 with* $\text{nocount}(|\pi|) = d \cdot 2^d$.

Proof. Let $\widehat{G} = Tr_2(G)$ be the 2-player game graph obtained from G. If $\langle\langle 1\rangle\rangle_{sure}(\text{finStreett}(F)) = \emptyset$ in G, then by Lemma 1 it follows that $\langle\langle 1\rangle\rangle_{sure}(\text{finStreett}(F)) = \emptyset$ in $\widehat{G}$, and then by Theorem 1 we have $\langle\langle 2\rangle\rangle_{sure}(\text{cofinStreett}(F)) = S$ for the game graph $\widehat{G}$. If $\langle\langle 2\rangle\rangle_{sure}(\text{cofinStreett}(F)) = S$ in $\widehat{G}$, then it follows from the results of [5] that there is a pure strategy $\widehat{\pi}$ in $\widehat{G}$ that satisfies the following conditions.

1. For every integer $b \geq 0$, for every strategy $\widehat{\sigma}$ of player 1 in $\widehat{G}$, and from all states s, the play from s given strategies $\widehat{\pi}$ and $\widehat{\sigma}$ satisfies the following condition: there exists position k and $1 \leq j \leq d$, such that the state s_k at the k-th position is in F_j, and for all $k \leq k' < k + b$ the state in k'-th position does not belong to E_j, and $k + b \leq |S| \cdot d \cdot 2^d \cdot (b + 1)$.
2. $\text{nocount}(|\widehat{\pi}|) = d \cdot 2^d$.

We obtain an almost-sure winning strategy π^* for player 2 in G as follows: set $b = 1$, the strategy π^* is played in rounds, and in round b the strategy is played according to the following rule:

1. (Step 1). Start play according to $\widehat{\pi}$
   (a) if at any random state the chosen successor is different from $\widehat{\pi}$, then go to the start of step 1 (i.e., start playing like the beginning of round b);
   (b) if for $|S| \cdot d \cdot 2^d \cdot (b + 1)$ steps at all random states the chosen successor matches $\widehat{\pi}$, then increment b and proceed to beginning of round $b + 1$.

We argue that the strategy π^* is almost-sure winning. Observe that since π^* follows $\widehat{\pi}$ in round b unless there is a deviation at a random state, it follows that if the strategy proceeds from round b to $b+1$, then at round b, there exists a position where the distance is at least b. Hence if the strategy π^* proceeds for infinitely many rounds, then cofinStreett(F) is satisfied. To complete the proof we argue that π^* proceeds through infinitely many rounds with probability 1. For a fixed b, the probability that step 1.(b). succeeds at a given trial is at least $\left(\frac{1}{\delta_{\min}}\right)^{|S| \cdot d \cdot 2^d \cdot (b+1)} > 0$, where $\delta_{\min} = \min\{\delta(s)(t) \mid s \in S_{\bigcirc}, t \in E(s)\} > 0$. Hence it follows that the probability that the strategy gets stuck in step 1.(a). for a fixed b is zero. Since the probability of a countable union of measure zero set is zero, it follows that the probability that the strategy gets stuck in step 1.(a). of any round b is zero. Hence with probability 1 the strategy π^* proceeds through infinitely many rounds, and the desired result follows. ∎

Lemma 2 states that for a finitary Streett objective if the sure winning set for player 1 is empty, then player 2 wins almost-surely everywhere in the game graph. Since parity objectives and finitary parity objectives are a special case of Streett and finitary Streett objectives, respectively, the result of Lemma 2 also holds for finitary parity objectives. This is formalized as the following lemma.

Lemma 3. *Let G be a $2^1\!/_2$-player game graph with the set S of states, and let p be a priority function. If $\langle\langle 1 \rangle\rangle_{sure}(\text{finParity}(p)) = \emptyset$, then the following assertions hold:*

1. *$\langle\langle 2 \rangle\rangle_{almost}(\text{cofinParity}(p)) = S$; and*
2. *there is an almost-sure winning strategy π for player 2 with $\text{nocount}(|\pi|) = 2$.*

We now present the notions of attractors in $2^1\!/_2$-player games and the basic properties of such attractors.

Definition 1 (Attractors). *Given a $2^1\!/_2$-player game graph G and a set $U \subseteq S$ of states, such that $G \upharpoonright U$ is a subgame, and $T \subseteq S$ we define $Attr_{1,\bigcirc}(T,U)$ as follows:*

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

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

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

Lemma 4 (Attractor properties). *Let G be a $2^1\!/_2$-player game graph and $U \subseteq S$ be a set of states such that $G \upharpoonright U$ is a subgame. For a set $T \subseteq S$ of states, let $Z = Attr_{1,\bigcirc}(T,U)$. Then the following assertions hold.*

1. *$G \upharpoonright (U \setminus Z)$ is a subgame.*
2. *Let σ^Z be a pure memoryless attractor strategy for player 1. There exists a constant $c > 0$, such that for all strategies π for player 2 in the subgame $G \upharpoonright U$ and for all states $s \in U$*
   - *(a) We have $\Pr_s^{\sigma^Z,\pi}(\text{Reach}(T)) \geq c \cdot \Pr_s^{\sigma^Z,\pi}(\text{Reach}(Z))$; and*
   - *(b) if $\Pr_s^{\sigma^Z,\pi}(\text{Buchi}(Z)) > 0$, then $\Pr_s^{\sigma^Z,\pi}(\text{Buchi}(T) \mid \text{Buchi}(Z)) = 1$.*

We now present the second key lemma for the algorithms for the qualitative analysis of $2^1\!/_2$-player finitary parity and finitary Streett games.

Lemma 5. *Let G be a $2^1\!/_2$-player game graph with the set S of states, and let Φ be a finitary parity or a finitary Streett objective with d pairs. If $\langle\langle 1 \rangle\rangle_{pos}(\Phi) = S$, then the following assertions hold:*

1. *$\langle\langle 1 \rangle\rangle_{almost}(\Phi) = S$;*
2. *if Φ is a finitary parity objective, then memoryless almost-sure winning strategies exist; and if Φ is a finitary Streett objective, then an almost-sure winning strategy with memory $d \cdot 2^d$ exists.*

Proof. The proof proceeds by iteratively removing sure winning sets, and the corresponding attractors from the graphs. Let $G^0 = G$, and $S^0 = S$. For $i \geq 0$, let G^i and S^i be the game graph and the set of states at the i-th iteration. Let

$Z_i = \langle\langle 1 \rangle\rangle_{sure}(\Phi)$ in G^i, and $A_i = Attr_{1,\bigcirc}(Z_i, S^i)$. Let $G^{i+1} = G \upharpoonright (S^i \setminus A_i)$, and $X_i = \bigcup_{j \leq i} A_i$. We continue this process unless for some k we have $X_k = S$. If for some game graph G^i we have $Z_i = \emptyset$ (i.e., $\langle\langle 1 \rangle\rangle_{sure}(\Phi) = \emptyset$ in G^i), then by Lemma 5 we have that $\langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi}) = S^i$, where $\overline{\Phi}$ is the complementary objective to Φ. This would contradict that $\langle\langle 1 \rangle\rangle_{pos}(\Phi) = S$. It follows that for some k we would have $X_k = S$. The almost-sure winning strategy σ^* for player 1 is defined as follows: in Z_i play a sure winning strategy for Φ in G^i, and in $A_i \setminus Z_i$ play a pure memoryless attractor strategy to reach Z_i. The strategy σ^* ensures the following: (a) from Z_i either the game stays in Z_i and satisfies Φ, or reaches X_{i-1} (this follows since a sure winning strategy is followed in G^i, and player 2 may choose to escape only to X_{i-1}); and (b) if A_i is visited infinitely often, then $X_{i-1} \cup Z_i$ is reached with probability 1 (this follows from the attractor properties, i.e., Lemma 4). It follows from the above two facts that with probability 1 the game settles in some Z_i, i.e., for all strategies π and all states s we have $\Pr_s^{\sigma^*,\pi}(\bigcup_{i \leq k} \text{coBuchi}(Z_i)) = 1$. It follows that for all strategies π and all states s we have $\Pr_s^{\sigma^*,\pi}(\Phi) = 1$. By choosing sure winning strategies in Z_i that satisfy the memory requirements (which is possible by Theorem 3 and Theorem 4) we obtain the desired result. ∎

Computation of positive winning set. Given a $2^1/_2$-player game graph G and a finitary parity or a finitary Streett objective Φ, the set $\langle\langle 1 \rangle\rangle_{pos}(\Phi)$ in G can be computed as follows. Let $G^0 = G$, and $S^0 = S$. For $i \geq 0$, let G^i and S^i be the game graph and the set of states at the i-th iteration. Let $Z_i = \langle\langle 1 \rangle\rangle_{sure}(\Phi)$ in G^i, and $A_i = Attr_{1,\bigcirc}(Z_i, S^i)$. Let $G^{i+1} = G \upharpoonright (S^i \setminus A_i)$, and $X_i = \bigcup_{j \leq i} A_i$. If $Z_i = \emptyset$, then $S^i = \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$ and $S \setminus S^i = \langle\langle 1 \rangle\rangle_{pos}(\Phi)$. The correctness follows from Lemma 2.

Computation of almost-sure winning set. Given a $2^1/_2$-player game graph G and a finitary parity or a finitary Streett objective Φ, the set $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ in G can be computed as follows. Let $G^0 = G$, and $S^0 = S$. For $i \geq 0$, let G^i and S^i be the game graph and the set of states at the i-th iteration. Let $\overline{Z}_i = \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$ in G^i, and $\overline{A}_i = Attr_{2,\bigcirc}(\overline{Z}_i, S^i)$. Let $G^{i+1} = G \upharpoonright (S^i \setminus \overline{A}_i)$, and $X_i = \bigcup_{j \leq i} \overline{A}_i$. In other words, the almost-sure winning set for player 2 and its attractor are iteratively removed from the game graph. If $\overline{Z}_i = \emptyset$, then in G^i we have $\langle\langle 1 \rangle\rangle_{pos}(\Phi) = S^i$, and by Lemma 5 we obtain that $\langle\langle 1 \rangle\rangle_{almost}(\Phi) = S^i$. That is we have $S^i = \langle\langle 1 \rangle\rangle_{almost}(\Phi)$ and $S \setminus S^i = \langle\langle 2 \rangle\rangle_{pos}(\overline{\Phi})$. We have the following theorem summarizing the qualitative complexity of $2^1/_2$-player games with finitary parity and finitary Streett objectives.

Theorem 5. *Given a $2^1/_2$-player game graph $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta))$ with n states and m edges, and given a finitary parity or a finitary Streett objective Φ, the following assertions hold.*

1. $\langle\langle 1 \rangle\rangle_{almost}(\Phi) = S \setminus \langle\langle 2 \rangle\rangle_{pos}(\overline{\Phi})$ *and* $\langle\langle 1 \rangle\rangle_{pos}(\Phi) = S \setminus \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$.
2. *The family of pure memoryless strategies suffices for almost-sure and positive winning with respect to finitary parity objectives on $2^1/_2$-player game graphs. If Φ is a finitary parity objective, then there exist infinite-memory*

almost-sure and positive winning strategies π *for player 2 for the complementary infinitary parity objective* $\overline{\Phi}$ *such that* π *is finite-memory counting with* $\mathrm{nocount}(|\pi|) = 2$. *In general no finite-memory almost-sure and positive winning strategies exist for player 2 for* $\overline{\Phi}$.

3. *If* Φ *is a finitary Streett objective with d pairs, then there exists a finite-memory almost-sure and positive winning strategy* σ *for player 1 such that* $|\sigma| = d \cdot 2^d$. *In general almost-sure and positive winning strategies for player 1 for the objective* Φ *require* $2^{\lfloor \frac{d}{2} \rfloor}$ *memory. There exist infinite-memory almost-sure and positive winning strategies* π *for player 2 for the complementary objective* $\overline{\Phi}$ *such that* π *is finite-memory counting with* $\mathrm{nocount}(|\pi|) = d \cdot 2^d$. *In general no finite-memory almost-sure and positive winning strategies exist for player 2 for* $\overline{\Phi}$.
4. *If* Φ *is a finitary parity objective, then the winning sets* $\langle\langle 1 \rangle\rangle_{pos}(\Phi)$ *and* $\langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$ *can be computed in time* $O(n^3 \cdot m)$, *and the sets* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ *and* $\langle\langle 2 \rangle\rangle_{pos}(\overline{\Phi})$ *can be computed in time* $O(n^4 \cdot m)$.
5. *If* Φ *is a finitary Streett objective with d pairs, then the winning sets* $\langle\langle 1 \rangle\rangle_{pos}(\Phi)$ *and* $\langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi})$ *can be computed in time* $O(n^3 \cdot m \cdot d^2 \cdot 4^d)$, *and the sets* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ *and* $\langle\langle 2 \rangle\rangle_{pos}(\overline{\Phi})$ *can be computed in time* $O(n^4 \cdot m \cdot d^2 \cdot 4^d)$.

4 Quantitative Analysis of Stochastic Finitary Games

In this section we consider the quantitative analysis of $2^1\!/\!_2$-player games with finitary parity and finitary Streett objectives. We start with notion of *value classes.*

Definition 2 (Value classes). *Given a finitary objective* Φ, *for every real* $r \in [0, 1]$ *the* value class *with value* r *is* $\mathrm{VC}(\Phi, r) = \{s \in S \mid \langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = r\}$ *is the set of states with value* r *for player 1. For* $r \in [0, 1]$ *we denote by* $\mathrm{VC}(\Phi, > r) = \bigcup_{q>r} \mathrm{VC}(\Phi, q)$ *the value classes greater than* r *and by* $\mathrm{VC}(\Phi, < r) = \bigcup_{q<r} \mathrm{VC}(\Phi, q)$ *the value classes smaller than* r. ∎

Definition 3 (Boundary probabilistic states). *Given a set* U *of states, a state* $s \in U \cap S_{\bigcirc}$ *is a* boundary probabilistic state *for* U *if* $E(s) \cap (S \setminus U) \neq \emptyset$, *i.e., the probabilistic state has an edge out of the set* U. *We denote by* $Bnd(U)$ *the set of boundary probabilistic states for* U. *For a value class* $\mathrm{VC}(\Phi, r)$ *we denote by* $Bnd(\Phi, r)$ *the set of boundary probabilistic states of value class* r. ∎

Observation. For a state $s \in Bnd(\Phi, r)$ we have $E(s) \cap \mathrm{VC}(\Phi, > r) \neq \emptyset$ and $E(s) \cap \mathrm{VC}(\Phi, < r) \neq \emptyset$, i.e., the boundary probabilistic states have edges to higher and lower value classes.

Reduction of a value class. Given a set U of states, such that U is δ-live, let $Bnd(U)$ be the set boundary probabilistic states for U. We denote by $G_{Bnd(U)}$ the subgame graph $G \upharpoonright U$ where every state in $Bnd(U)$ is converted to an absorbing state (state with a self-loop). Since U is δ-live, we have $G_{Bnd(U)}$ is a subgame graph. We denote by $G_{Bnd(\Phi,r)}$ the subgame graph where every boundary probabilistic state in $Bnd(\Phi, r)$ is converted to an absorbing state. For a tail objective

Φ, we denote by $G_{\Phi,r} = G_{Bnd(\Phi,r)} \upharpoonright \mathrm{VC}(\Phi, r)$: this is a subgame graph since for a tail objective Φ every value class is δ-live, and δ-closed as all states in $Bnd(\Phi, r)$ are converted to absorbing states. We now present a property of tail objectives and we present our results that use the property. Since tail objectives subsume finitary parity and finitary Streett objectives, the desired results would follow for finitary parity and finitary Streett objectives.

Almost-limit property for tail objectives. An objective Φ satisfies the *almost-limit* property if for all $2^1/_2$-player game graphs and for all $F, R \subseteq S$ the following equalities hold:

$$\{s \in S \mid \langle\langle 1\rangle\rangle_{val}(\Phi \cap \mathrm{Safe}(F)) = 1\} = \langle\langle 1\rangle\rangle_{almost}(\Phi \cap \mathrm{Safe}(F));$$

$$\{s \in S \mid \langle\langle 1\rangle\rangle_{val}(\Phi \cup \mathrm{Reach}(R)) = 1\} = \langle\langle 1\rangle\rangle_{almost}(\Phi \cup \mathrm{Reach}(R));$$

$$\{s \in S \mid \langle\langle 2\rangle\rangle_{val}(\overline{\Phi} \cap \mathrm{Safe}(F)) = 1\} = \langle\langle 2\rangle\rangle_{almost}(\overline{\Phi} \cap \mathrm{Safe}(F));$$

$$\{s \in S \mid \langle\langle 2\rangle\rangle_{val}(\overline{\Phi} \cup \mathrm{Reach}(R)) = 1\} = \langle\langle 2\rangle\rangle_{almost}(\overline{\Phi} \cup \mathrm{Reach}(R)).$$

If Φ is a tail objective, then the objective $\Phi \cap \mathrm{Safe}(F)$ can be interpreted as a tail objective $\Phi \cap \mathrm{coBuchi}(F)$ by transforming every state in $S \setminus F$ as a loosing absorbing state. Similarly, if Φ is a tail objective, then the objective $\Phi \cup \mathrm{Reach}(R)$ can be interpreted as a tail objective $\Phi \cup \mathrm{Buchi}(R)$ by transforming every state in R as winning absorbing state. From the results of [13] (Chapter 3) it follows that for all tail objectives Φ we have

$$\{s \in S \mid \langle\langle 1\rangle\rangle_{val}(\Phi)(s) = 1\} = \langle\langle 1\rangle\rangle_{almost}(\Phi);$$

$$\{s \in S \mid \langle\langle 2\rangle\rangle_{val}(\overline{\Phi})(s) = 1\} = \langle\langle 2\rangle\rangle_{almost}(\overline{\Phi}).$$

Hence it follows that all tail objective satisfy the almost-limit property. We now present a lemma, that extends a property of $2^1/_2$-player games with ω-regular objectives to tail objectives (that subsumes finitary parity and Streett objectives).

Lemma 6 (Almost-sure reduction). *Let G be a $2^1/_2$-player game graph and Φ be a tail objective. For $0 < r < 1$, the following assertions hold.*

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

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

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

that is, q is the maximum value a successor state t of a player 1 state $s \in \mathrm{VC}(\Phi, r)$ such that the successor state t is not in $\mathrm{VC}(\Phi, r)$. We must have $q < r$. Hence if player 1 chooses to escape the value class $\mathrm{VC}(\Phi, r)$, then player 1 gets to see a state with value at most $q < r$. We consider the subgame graph $G_{\Phi,r}$. Let $U = \mathrm{VC}(\Phi, r)$ and $Z = \mathit{Bnd}(\Phi, r)$. Assume towards contradiction, there exists a state $s \in U$ such that $s \notin \langle\langle 1 \rangle\rangle_{almost}(\Phi \cup \mathrm{Reach}(Z))$. Then we have $s \in (U \setminus Z)$; and since Φ is a tail objective satisfying the almost-limit property and $s \notin \langle\langle 1 \rangle\rangle_{almost}(\Phi \cup \mathrm{Reach}(Z))$ we have $\langle\langle 2 \rangle\rangle_{val}(\overline{\Phi} \cap \mathrm{Safe}(U \setminus Z))(s) > 0$. Observe that in $G_{\Phi,r}$ we have all states in Z are absorbing states, and hence the objective $\overline{\Phi} \cap \mathrm{Safe}(U \setminus Z)$ is equivalent to the objective $\overline{\Phi} \cap \mathrm{coBuchi}(U \setminus Z)$, which can be considered as a tail objective. Since $\langle\langle 2 \rangle\rangle_{val}(\overline{\Phi} \cap \mathrm{Safe}(U \setminus Z))(s) > 0$, for some state s, it follows from Theorem 1 of [2] that there exists a state $s_1 \in (U \setminus Z)$ such that $\langle\langle 2 \rangle\rangle_{val}(\overline{\Phi} \cap \mathrm{Safe}(U \setminus Z)) = 1$. Then, since Φ is a tail objective satisfying the almost-limit property, it follows that there exists a strategy $\widehat{\pi}$ for player 2 in $G_{\Phi,r}$ such that for all strategies $\widehat{\sigma}$ for player 1 in $G_{\Phi,r}$ we have $\mathrm{Pr}_{s_1}^{\widehat{\sigma},\widehat{\pi}}(\overline{\Phi} \cap \mathrm{Safe}(U \setminus Z)) = 1$. We will now construct a strategy π^* for player 2 as a combination of the strategy $\widehat{\pi}$ and a strategy in the original game G. By Martin's determinacy result (Theorem 2), for all $\varepsilon > 0$, there exists an ε-optimal strategy π_ε for player 2 in G such that for all $s \in S$ and for all strategies σ for player 1 we have

$$\mathrm{Pr}_s^{\sigma,\pi_\varepsilon}(\overline{\Phi}) \geq \langle\langle 2 \rangle\rangle_{val}(\overline{\Phi})(s) - \varepsilon.$$

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

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

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

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

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

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

Hence for all strategies σ we have

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

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

Lemma 7 (Almost-sure to optimality). *Let G be a $2^1\!/_2$-player game graph and Φ be a tail objective. Let σ be a strategy such that*

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

Then σ is an optimal strategy. Analogous result holds for player 2 strategies.

Proof. (Sketch). Consider a strategy σ satisfying the conditions of the lemma, a starting state s, and a counter strategy π. If the play settles in a value-class with $r > 0$, i.e., satisfies $\mathrm{coBuchi}(\mathrm{VC}(\Phi, r))$, for some $r > 0$, then the play satisfies Φ almost-surely. From a value class the play can leave the value class if player 2 chooses to leave to a greater value class, or by reaching the boundary probabilistic states such that average value of the successor states is the value of the value class. Hence it follows that (a) either the event $\bigcup_{r>0} \mathrm{coBuchi}(\mathrm{VC}(\Phi, r))$ holds, and then Φ holds almost-surely; (b) else the event $\mathrm{Reach}(\langle\langle 1 \rangle\rangle_{almost}(\Phi) \cup \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi}))$ holds, and by the conditions on leaving the value class it follows that $\Pr_s^{\sigma,\pi}(\mathrm{Reach}(\langle\langle 1 \rangle\rangle_{almost}(\Phi)) \mid \mathrm{Reach}(\langle\langle 1 \rangle\rangle_{almost}(\Phi) \cup \langle\langle 2 \rangle\rangle_{almost}(\overline{\Phi}))) \geq \langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$. It follows that for all $s \in S$ and all strategies π we have $\Pr_s^{\sigma,\pi}(\Phi) \geq \langle\langle 1 \rangle\rangle_{val}(\Phi)(s)$. The desired result follows. ∎

It follows from Lemma 6 that for tail objectives, strategies satisfying the conditions of Lemma 7 exist. It follows from Lemma 7 that optimal strategies for player 1 for tail objectives (and hence for finitary parity and Streett objectives), and optimal strategies for player 2 for the corresponding complementary objectives is no more complex than the respective almost-sure winning strategies.

Lemma 8. *Let $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$ be a $2^1\!/_2$-player game with a tail objective Φ. Let $\mathcal{P} = (V_0, V_1, \ldots, V_k)$ be a partition of the state space S, and let $r_0 > r_1 > r_2 > \ldots > r_k$ be k-real values such that the following conditions hold:*

1. *$V_0 = \langle\langle 1\rangle\rangle_{almost}(\Phi)$ and $V_k = \langle\langle 2\rangle\rangle_{almost}(\overline{\Phi})$;*
2. *$r_0 = 1$ and $r_k = 0$;*
3. *for all $1 \le i \le k-1$ we have $Bnd(V_i) \neq \emptyset$ and V_i is δ-live;*
4. *for all $1 \le i \le k-1$ and all $s \in S_2 \cap V_i$ we have $E(s) \subseteq \bigcup_{j \le i} V_j$;*
5. *for all $1 \le i \le k-1$ we have $V_i = \langle\langle 1\rangle\rangle_{almost}(\Phi \cup \mathrm{Reach}(Bnd(V_i)))$ in $G_{Bnd(V_i)}$;*
6. *let $x_s = r_i$, for $s \in V_i$, and for all $s \in S_\bigcirc$, let x_s satisfy $x_s = \sum_{t \in E(s)} x_t \cdot \delta(s)(t)$.*

Then we have $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) \ge x_s$ for all $s \in S$. Analogous result holds for player 2.

Proof. (Sketch). We fix a strategy σ such that σ is almost-sure winning from $\langle\langle 1\rangle\rangle_{almost}(\Phi)$, and in every V_i, for $1 \le i \le k-1$, it is almost-sure winning for the objective $\Phi \cup \mathrm{Reach}(Bnd(V_i))$. Arguments similar to Lemma 7 shows that for $s \in S$ and for all π we have $\mathrm{Pr}_s^{\sigma,\pi}(\Phi)(s) \ge x_s$. ∎

Algorithm for quantitative analysis. We now present an algorithm for quantitative analysis for $2^1\!/_2$-player games with tail objectives. The algorithm is a NP algorithm with an oracle access to the qualitative algorithms. The algorithm is based on Lemma 8. Given a $2^1\!/_2$-player game $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$ with a finitary parity or a finitary Streett objective Φ, a state s and a rational number r, the following assertion hold: if $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) \ge r$, then there exists a partition $\mathcal{P} = (V_0, V_1, V_2, \ldots, V_k)$ of S and rational values $r_0 > r_1 > r_2 > \ldots > r_k$, such that $r_i = \frac{p_i}{q_i}$ with $p_i, q_i \le \delta_u^{4\cdot|E|}$, where $\delta_u = \max\{q \mid \delta(s)(t) = \frac{p}{q}$ for $p, q \in \mathbb{N},\ s \in S_\bigcirc$ and $\delta(s)(t) > 0\}$, such that conditions of Lemma 8 are satisfied, and $s \in V_i$ with $r_i \ge r$. The witness $\mathcal{P}$ is the value class partition and the rational values represent the values of the value classes, and the precision of the values can also be proved (we omit details due to lack of space). From the above observation we obtain the algorithm for quantitative analysis as follows: given a $2^1\!/_2$-player game graph $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$ with a finitary parity or a finitary Streett objective Φ, a state s and a rational r, to verify that $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) \ge r$, the algorithm guesses a partition $\mathcal{P} = (V_0, V_1, V_2, \ldots, V_k)$ of S and rational values $r_0 > r_1 > r_2 > \ldots > r_k$, such that $r_i = \frac{p_i}{q_i}$ with $p_i, q_i \le \delta_u^{4\cdot|E|}$, and then verifies that all the conditions of Lemma 8 are satisfied, and $s \in V_i$ with $r_i \ge r$. Observe that since the guesses of the rational values can be made with $O(|G| \cdot |S| \cdot |E|)$ bits, the guess is polynomial in size of the game. The condition 1 and the condition 5 of Lemma 8 can be verified by any qualitative algorithms, and all the other conditions can be checked in polynomial time. We now summarize the results on quantitative analysis of $2^1\!/_2$-player games with tail objectives, and then present the results for finitary parity and finitary Streett objectives.

Theorem 6. *Given a $2^1\!/_2$-player game graph and a tail objective Φ, the following assertions hold.*

1. *If a family $\Sigma^{\mathcal{C}}$ of strategies suffices for almost-sure winning for Φ, then the family $\Sigma^{\mathcal{C}}$ of strategies also suffices for optimality for Φ.*
2. *Given a rational number r and a state s, whether $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) \geq r$ can be decided in $NP^{\mathcal{A}}$, where $\mathcal{A}$ is an oracle for the qualitative analysis of Φ on $2^1\!/_2$-player game graphs.*

Theorem 7. *Given a $2^1\!/_2$-player game graph and a finitary parity or a finitary Streett objective Φ, the following assertions hold.*

1. *The family of pure memoryless strategies suffices for optimality with respect to finitary parity objectives on $2^1\!/_2$-player game graphs. If Φ is a finitary parity objective, then there exist infinite-memory optimal strategies π for player 2 for the complementary infinitary parity objective $\overline{\Phi}$ such that π is finite-memory counting with* $\mathrm{nocount}(|\pi|) = 2$. *In general no finite-memory optimal strategies exist for player 2 for $\overline{\Phi}$.*
2. *If Φ is a finitary Streett objective with d pairs, then there exists a finite-memory optimal strategy σ for player 1 such that $|\sigma| = d \cdot 2^d$. In general optimal strategies for player 1 for the objective Φ require $2^{\lfloor \frac{d}{2} \rfloor}$ memory. There exist infinite-memory optimal strategies π for player 2 for the complementary objective $\overline{\Phi}$ such that π is finite-memory counting with* $\mathrm{nocount}(|\pi|) = d \cdot 2^d$. *In general no finite-memory optimal strategy exists for player 2 for $\overline{\Phi}$.*
3. *If Φ is a finitary parity objective, then given a rational r and a state s, whether $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) \geq r$ can be decided in NP $\cap$ coNP.*
4. *If Φ is a finitary Streett objective, then given a rational r and a state s, whether $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) \geq r$ can be decided in EXPTIME.*

Remark 2. For $2^1\!/_2$-player games with finitary objectives, the qualitative analysis can be achieved in polynomial time, however, we only prove a NP $\cap$ coNP bound for the quantitative analysis. It may be noted that for $2^1\!/_2$-player game graphs, the quantitative analysis for finitary and nonfinitary Büchi objectives coincide. The best known bound for quantitative analysis of $2^1\!/_2$-player games with Büchi objectives is NP $\cap$ coNP, and obtaining a polynomial time algorithm is a major open problem. Hence obtaining a polynomial time algorithm for quantitative analysis of $2^1\!/_2$-player games with finitary parity objectives would require the solution of a major open problem. ∎

References

1. Alur, R., Henzinger, T.A.: Finitary fairness. In: LICS 1994, pp. 52–61. IEEE, Los Alamitos (1994)
2. Chatterjee, K.: Concurrent games with tail objectives. Theoretical Computer Science 388, 181–198 (2007)
3. Chatterjee, K., de Alfaro, L., Henzinger, T.A.: The complexity of stochastic Rabin and Streett games. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 878–890. Springer, Heidelberg (2005)

4. Chatterjee, K., Henzinger, T.A.: Finitary winning in ω-regular games. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 257–271. Springer, Heidelberg (2006)
5. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in ω-regular games. Technical Report: UCB/EECS-2007-120 (2007)
6. Chatterjee, K., Jurdziński, M., Henzinger, T.A.: Quantitative stochastic parity games. In: SODA 2004, pp. 121–130. SIAM, Philadelphia (2004)
7. de Alfaro, L., Henzinger, T.A.: Concurrent omega-regular games. In: LICS 2000, pp. 141–154. IEEE, Los Alamitos (2000)
8. Dziembowski, S., Jurdzinski, M., Walukiewicz, I.: How much memory is needed to win infinite games? In: LICS 1997, pp. 99–110. IEEE, Los Alamitos (1997)
9. Emerson, E.A., Jutla, C.: The complexity of tree automata and logics of programs. In: FOCS 1988, pp. 328–337. IEEE, Los Alamitos (1988)
10. Gurevich, Y., Harrington, L.: Trees, automata, and games. In: STOC 1982, pp. 60–65. ACM Press, New York (1982)
11. Horn, F.: Dicing on the streett. IPL 104, 1–9 (2007)
12. Horn, F.: Faster algorithms for finitary games. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 472–484. Springer, Heidelberg (2007)
13. Horn, F.: Random Games. PhD thesis, Université Denis-Diderot and RWTH, Aachen (2008)
14. Martin, D.A.: Borel determinacy. Annals of Mathematics 102(2), 363–371 (1975)
15. Martin, D.A.: The determinacy of Blackwell games. The Journal of Symbolic Logic 63(4), 1565–1581 (1998)
16. Piterman, N., Pnueli, A.: Faster solution of Rabin and Streett games. In: LICS 2006, pp. 275–284. IEEE, Los Alamitos (2006)
17. Zielonka, W.: Perfect-information stochastic parity games. In: Walukiewicz, I. (ed.) FOSSACS 2004. LNCS, vol. 2987, pp. 499–513. Springer, Heidelberg (2004)