Quantitative Stochastic Parity Games

Krishnendu Chatterjee *Marcin Jurdziński* *Thomas A. Henzinger*

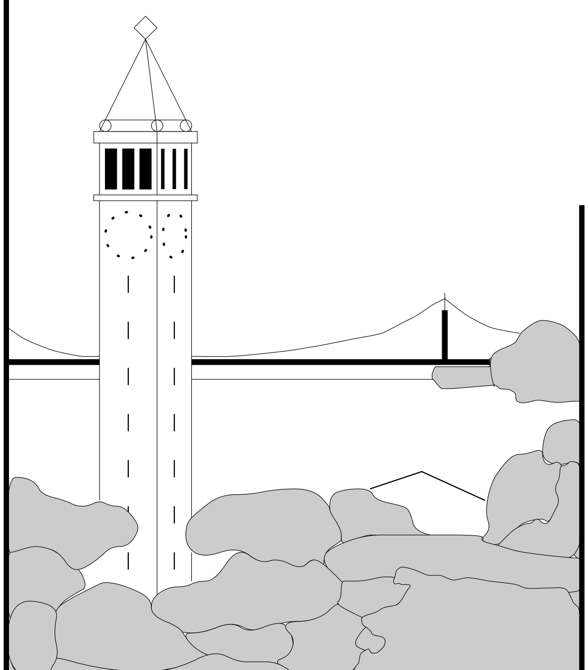

Report No. UCB/CSD-3-1280

October 2003

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

Quantitative Stochastic Parity Games*, **

Krishnendu Chatterjee, Marcin Jurdziński, and Thomas A. Henzinger

Department of Electrical Engineering and Computer Sciences
University of California, Berkeley, USA
{c_krish,mju,tah}@eecs.berkeley.edu

Abstract. We study *perfect-information stochastic parity games*. These are two-player nonterminating games which are played on a graph with turn-based probabilistic transitions. A play results in an infinite path and the conflicting goals of the two players are ω-regular path properties, formalized as parity winning conditions. The *qualitative solution* of such a game amounts to computing the set of vertices from which a player has a strategy to win with probability 1 (or with positive probability). The *quantitative solution* amounts to computing the value of the game in every vertex, i.e., the highest probability with which a player can guarantee satisfaction of his own objective in a play that starts from the vertex.

For the important special case of one-player stochastic parity games (*parity Markov decision processes*) we give polynomial-time algorithms both for the qualitative and the quantitative solution. The running time of the qualitative solution is $O(d \cdot m^{3/2})$ for graphs with m edges and d priorities. The quantitative solution is based on a linear-programming formulation.

For the two-player case, we establish the existence of optimal pure memoryless strategies. This has several important ramifications. First, it implies that the values of the games are rational. This is in contrast to the *concurrent* stochastic parity games of de Alfaro et al.; there, values are in general algebraic numbers, optimal strategies do not exist, and ε-optimal strategies have to be mixed and with infinite memory. Second, the existence of optimal pure memoryless strategies together with the polynomial-time solution for one-player case implies that the quantitative two-player stochastic parity game problem is in NP $\cap$ co-NP. This generalizes a result of Condon for stochastic games with reachability objectives. It also constitutes an exponential improvement over the best previous algorithm, which is based on a doubly exponential procedure of de Alfaro and Majumdar for concurrent stochastic parity games and provides only ε-approximations of the values.

---

* An abbreviated version of this paper will appear in the *ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2004.

** This research was supported in part by the DARPA grant F33615-C-98-3614, the ONR grant N00014-02-1-0671, the NSF grants CCR-9988172 and CCR-0225610, the Polish KBN grant 7-T11C-027-20, and the EU RTN HPRN-CT-2002-00283.

1 Introduction

Perfect-information stochastic games [5] are played on a graph (V, E) by three players —Even, Odd, and Random— who move a token from vertex to vertex so that an infinite path is formed. Given a partition $(V_\Box, V_\Diamond, V_\bigcirc)$ of the set V of vertices, player Even moves if the token is in $V_\Box$, player Odd if the token is in $V_\Diamond$, and player Random if the token is in $V_\bigcirc$. Player Random always moves by choosing a successor vertex uniformly at random. Thus his freedom of choice is limited and we refer to him as a "half player." Henceforth we refer to stochastic games as $2^1\!/_2$-player games, and to stochastic games with $V_\Box = \emptyset$ or $V_\Diamond = \emptyset$ as $1^1\!/_2$-player games, often called *Markov decision processes* [22, 13]. Stochastic games were introduced by Shapley [24] and have been studied extensively in several research communities; the book by Filar and Vrieze [13] provides a unified treatment of the theories of stochastic games and Markov decision processes [22].

The exact computational complexity of stochastic games is a fascinating open problem [23, 5]. In practice, the algorithms known so far are often prohibitively expensive. In their survey on the complexity of solving Markov decision processes, Littman et al. [19] call for investigating subclasses of the general model which admit efficient solution algorithms and are yet sufficiently expressive for modeling correctness or optimality objectives found in practice. Raghavan and Filar [23] list several classes of stochastic games whose structural properties bear a promise for efficient solution algorithms. They single out *perfect-information stochastic games* as one of the most promising, commenting that "despite its obvious appeal, this class has not been studied from an algorithmic point of view." Condon [5] has studied *simple stochastic games* and she has proved that they are log-space complete in the class of logarithmic-space randomized alternating Turing machines [5]. Condon's simple stochastic games can be seen as perfect-information [23, 13] (or "turn-based" [8]) recursive stochastic games [12] with 0-1 payoffs, or alternatively, as perfect-information stochastic games with reachability objectives [8].

We generalize the finitary reachability objectives of Condon's model by considering infinitary *ω-regular* objectives [26], which are popular for the specification and verification of temporal properties of computational systems [14] and subsume reachability objectives as a very special case. On the other hand, our model of perfect-information stochastic parity games is a proper subclass of the general (or "concurrent") stochastic games [13, 20, 7, 9]. For brevity, throughout this paper we will omit the term "perfect information"; whenever we talk about stochastic parity games, we mean perfect-information stochastic parity games. We will explicitly say "concurrent stochastic parity games" when referring to the more general model of de Alfaro, Henzinger, and Majumdar [7, 9].

An objective for a player in an infinite path-forming game is specified by the set of infinite paths that are *winning* for the player. A *stochastic parity game* is given by a game graph as above and a *priority function* $p \colon V \to \mathbb{N}$, which labels every vertex $v \in V$ with a priority $p(v)$. An infinite path is *winning* for the player Even if the smallest priority that occurs infinitely often is even; otherwise the path is winning for the player Odd. This way of specifying a winning objective is called the parity (Rabin chain) condition and it is a normal form for all ω-regular conditions [26]. Two-player parity games have recently attracted considerable attention in the computer-aided verification community, because they can be used to solve model-checking and control problems for the μ-calculus [1] and its numerous modal and temporal sublogics [10, 26, 17, 7, 9, 14, 2]. Parity games are also a powerful theoretical tool for resolving difficult decidability and expressiveness questions in mathematical logic and automata theory [15, 10, 26, 29, 3, 1, 14].

The basic algorithmic problems for stochastic games are that of determining which of the players has a winning strategy, or more accurately, estimating the probability of winning that a player can ensure by following a certain strategy. Several *qualitative* winning criteria have been studied, such as almost-sure winning (i.e., winning with probability 1) or positive-probability winning [7, 4]. The *quantitative* solution [9] amounts to computing the value of a game, i.e., the highest probability of winning that a player can ensure against an arbitrary strategy of the opponent. In Section 2 we define stochastic parity games and the algorithmic problems that we study in this paper. In general, *mixed* strategies have to be considered when solving stochastic games [24], a common phenomenon in non-perfect-information games [28]. A mixed strategy prescribes randomized moves, i.e., the player moves the token at random according to a probability distribution over the successor vertices. By contrast, a strategy is *pure* if it always prescribes deterministic moves, i.e., the player always moves the token deterministically to one of the successor vertices. A strategy is *memoryless* if the moves it prescribes do not depend on the history of the play carried out so far, i.e., if it prescribes the same move at every visit to a vertex.

In Section 3 we give a graph-theoretic characterization of the almost-sure and positive-probability winning sets in $1^1/_2$-player parity games, inspired by the work of Courcoubetis and Yannakakis [6]. From this characterization we derive the existence of *pure memoryless optimal strategies* for the qualitative and quantitative criteria in $1^1/_2$-player parity games, and we present efficient polynomial-time algorithms for solving both qualitative and quantitative $1^1/_2$-player parity games.

Theorem 1. (Complexity of solving $1^1/_2$-player parity games) *Let m be the number of edges in the graph of a parity game, and let d be the number of different values assigned to vertices by the priority function.*

1. *Qualitative $1^1/_2$-player parity games can be solved in time $O(d \cdot m^{3/2})$.*
2. *Quantitative $1^1/_2$-player parity games can be solved in polynomial time by solving a linear program that can be constructed from the qualitative solution.*

Previously, an efficient $O(m^{3/2})$ algorithm was known only for the qualitative solution of the special cases of $1^1/_2$-player reachability and Büchi games [4]. A polynomial-time algorithm for the qualitative solution of $1^1/_2$-player parity games can also be obtained by modifying a reduction [21, 16] to limiting-average (aka. mean-payoff [30]) Markov decision processes (MDP's) [22]. This approach, however, requires the solution of a linear program for limiting-average MDP's [22, 13], whereas we give a direct $O(d \cdot m^{3/2})$ algorithm which uses the solution for Büchi MDP's [4] as a subroutine.

In Section 4 we use the existence of pure memoryless qualitatively optimal strategies in $1^1\!/\!_2$-player parity games to give a simplified proof of existence of *pure memoryless qualitatively optimal strategies* for $2^1\!/\!_2$-player parity games [4]. Then, in Section 5, we apply those results to prove the existence of *pure memoryless optimal strategies* for the quantitative $2^1\!/\!_2$-player games, which generalizes the result of Condon [5] for stochastic reachability games.

Theorem 2. (Quantitative pure memoryless optimal strategies) *Optimal pure memoryless strategies exist in quantitative $2^1\!/\!_2$-player parity games.*

This result is in sharp contrast with the concurrent parity games of de Alfaro et al. [8, 7, 9]. Even for qualitative concurrent games, in general only ε-optimal strategies exist, and mixed strategies with infinite memory are necessary. Theorems 1 and 2 yield an improvement of the computational complexity for solving quantitative $2^1\!/\!_2$-player parity games.

Corollary 1. (Complexity of solving $2^1\!/\!_2$-player parity games) *The decision problems for qualitative and quantitative $2^1\!/\!_2$-player parity games are in NP $\cap$ co-NP, and there is an exponential-time algorithm for computing the exact values. The values of quantitative $2^1\!/\!_2$-player parity games are rational.*

The latter result is again in contrast to concurrent parity games. Their values are in general algebraic numbers, and there are simple examples with irrational values [9]. The previously best complexity for the quantitative solution of $2^1\!/\!_2$-player parity games was based on a doubly exponential algorithm [9] for solving concurrent parity games; so our results yield an exponential improvement. Moreover, the algorithm of de Alfaro and Majumdar [9], which uses a decision procedure for the first-order theory of the reals with addition and multiplication [25], does not compute the exact values of a game but only ε-approximations.

The existence of polynomial-time algorithms for non-stochastic parity games [11, 27], for limiting-average games [23, 30], and for stochastic reachability games [5] are long-standing open questions. We believe that the existence of pure memoryless optimal strategies for stochastic parity games —i.e., Theorem 2— makes it likely that algorithmic techniques that are developed for any of those classes of games can be carried over also to stochastic parity games.

2 Stochastic Parity Games

A *stochastic game graph* (*$2^1\!/\!_2$-player game graph*) $G = (V, E, (V_\Box, V_\Diamond, V_\bigcirc))$ consists of a directed graph (V, E) and a partition $(V_\Box, V_\Diamond, V_\bigcirc)$ of the vertex set V. For technical convenience we assume that every vertex has at least one and at most two out-going edges.

For a set $U \subseteq V$ of vertices, we write $G \restriction U$ for the subgraph of the graph (V, E) induced by the set of vertices U. We say that $G \restriction U$ is a *subgame* of the game G if the for all vertices $u \in U$, we have: if $u \in V_{\bigcirc}$ then $(u, w) \in E$ implies that $w \in U$, and if $u \in V_{\Box} \cup V_{\Diamond}$ then there is an edge $(u, w) \in E$, such that $w \in U$.

An infinite *path* in a game graph G is an infinite sequence $\langle v_0, v_1, v_2, \ldots \rangle$ 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 infinite paths, and for every vertex $s \in V$ we write Ω_s for the set of all infinite paths starting from the vertex s. A nonterminating *stochastic game* G is a $2^1/_2$-player infinite-duration path-forming game played on the graph G. Player Even keeps moving a token along edges of the game graph from vertices in $V_{\Box}$, player Odd keeps moving the token from vertices in $V_{\Diamond}$, and player Random from vertices in $V_{\bigcirc}$. Player Random always passes the token to one of its two successors with probability $1/2$.

Stochastic game graphs can be obviously generalized to non-binary graphs and the transition probabilities for the random vertices to arbitrary rational numbers. All results in this paper easily follow for those generalizations by minor modifications and hence for simplicity we stick to the basic simple model without loss of generality. Moreover, for all algorithmic problems we consider in this paper there is a straightforward translation of game graphs with n vertices and m edges into equivalent binary game graphs with $O(m)$ vertices and $O(m)$ edges. This yields our main results listed in Section 1 from the results for binary graphs in the following sections.

The non-stochastic perfect-information *2-player game graphs* are a special case of the $2^1/_2$-player parity game graphs with $V_{\bigcirc} = \emptyset$. The $1^1/_2$*-player game graphs* (*Markov decision processes*) are a special case of the $2^1/_2$-player game graphs with $V_{\Diamond} = \emptyset$ or $V_{\Box} = \emptyset$.

2.1 Strategies.

For a countable set A, a probability distribution on the set A is a function $\delta : A \to [0, 1]$ such that $\sum_{a \in A} \delta(a) = 1$. We denote the set of probability distributions on the set A by $\mathcal{D}(A)$. A *strategy* for the player Even is a function $\sigma : V^* V_{\Box} \to \mathcal{D}(V)$, assigning a probability distribution to every finite sequence $\boldsymbol{v} \in V^* V_{\Box}$ of vertices, representing the history of the play so far. We say that the player Even *uses* (or *follows*) the strategy σ if in each move, given that the current history of the play is $\boldsymbol{v}$, he chooses the next vertex according to the probability distribution $\sigma(\boldsymbol{v})$. A strategy must prescribe only available moves, i.e., for all $\boldsymbol{w} \in V^*$, $v \in V_{\Box}$, and $u \in V$, if $\sigma(\boldsymbol{w} \cdot v)(u) \neq 0$ then $(v, u) \in E$. A strategy σ is *pure* if for all $\boldsymbol{w} \in V^*$ and $v \in V_{\Box}$, there is a vertex $u \in V$ such that $\sigma(\boldsymbol{w} \cdot v)(u) = 1$. Strategies for player Odd are defined analogously. We write Σ and Π for the sets of all strategies for players Even and Odd, respectively. A *pure memoryless strategy* is a pure strategy which does not depend on the whole history of the play but only on the current vertex. A pure memoryless strategy for player Even can be represented as a function $\sigma : V_{\Box} \to V$ such that $(v, \sigma(v)) \in E$, for all $v \in V_{\Box}$. Analogously we define pure memoryless strategies for player Odd.

Given a strategy $\sigma \in \Sigma$ for the player Even, we write G_σ for the game played as the game G with the constraint that the player Even uses the strategy σ. Note that if $\sigma : V_\Box \to V$ is a pure memoryless strategy then we can think of G_σ as the subgraph of the game graph G obtained from G by removing the edges $(v, w) \in E$, such that $v \in V_\Box$ and $\sigma(v) \neq w$. The definitions for a strategy π for the player Odd are similar.

Once a starting vertex $s \in V$ and strategies $\sigma \in \Sigma$ and $\pi \in \Pi$ for the two players are fixed, the play of the game is a random path $\omega_s^{\sigma,\pi}$ for which the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega_s$ is a measurable set of paths. For a vertex $s \in V$ and an event $\mathcal{A} \subseteq \Omega_s$, we write $\Pr_s^{\sigma,\pi}(\mathcal{A})$ for the probability that a path belongs to $\mathcal{A}$ if the game starts from the vertex s, and the players use the strategies σ and π, respectively.

2.2 Winning objectives.

A general way of specifying *objectives* of players in nonterminating stochastic games is by providing the set of *winning plays* $\mathcal{W} \subseteq \Omega$ for the player. In this paper we study only zero-sum games [23, 13] in which the objectives of the players are strictly competitive. In other words, it is implicit that if the objective of a player is a set $\mathcal{W}$ then the objective of the other player is the set $\Omega \setminus \mathcal{W}$. Given a game graph G and an objective $\mathcal{W} \subseteq \Omega$, we write $G(\mathcal{W})$ for the game played on the graph G with the objective $\mathcal{W}$ for the player Even.

If $n \in \mathbb{N}$ we write $[n]$ for the set $\{0, 1, 2, \ldots, n\}$ and $[n]_+$ for the set $\{1, 2, \ldots, n\}$. In this paper we consider the following classes of winning objectives.

- *Reachability.* For a set T of target vertices, the reachability objective is defined as $\text{Reach}(T) = \{ \langle v_0, v_1, v_2 \ldots \rangle \in \Omega \; : \; v_k \in T, \text{ for some } k \}$.

- *Büchi.* For a set T of target vertices, the Büchi (*repeated reachability*) objective is defined as $\text{Büchi}(T) = \{ \langle v_0, v_1, v_2 \ldots \rangle \in \Omega \; : \; v_k \in T, \text{ for infinitely many } k \}$.

- *Parity.* Let p be a function $p : V \to [d]$ assigning a *priority* $p(v)$ to every vertex $v \in V$, where $d \in \mathbb{N}$. For an infinite path $\omega = \langle v_0, v_1, \ldots \rangle \in \Omega$, we define $\text{Inf}(\omega) = \{ v \in V \; : \; v_k = v \text{ for infinitely many } k \geq 0 \}$. The *even parity objective* is defined as $\text{Even}(p) = \{ \omega \in \Omega \; : \; \min(p(\text{Inf}(\omega))) \text{ is even} \}$, and the *odd parity objective* as $\text{Odd}(p) = \{ \omega \in \Omega \; : \; \min(p(\text{Inf}(\omega))) \text{ is odd} \}$.

For every $s \in V$, we write $\text{Reach}(T)_s$ (resp. $\text{Büchi}(T)_s$, $\text{Even}(p)_s$, and $\text{Odd}(p)_s$) for the set $\text{Reach}(T) \cap \Omega_s$ (resp. $\text{Büchi}(T) \cap \Omega_s$, $\text{Even}(p) \cap \Omega_s$, and $\text{Odd}(p) \cap \Omega_s$). The reachability and Büchi objectives are special cases of the parity objectives. A Büchi objective $\text{Büchi}(T)$ can be encoded as an even parity objective for the priority function $p : V \to [1]$ with only two priorities (0 and 1), such that $p(s) = 0$ if and only if $s \in T$, for all $s \in V$. A reachability objective is a special case of the Büchi objective, where all vertices in the target set T are absorbing, i.e., if $s \in T$ and $(s, v) \in E$ then $s = v$.

A *stochastic parity game* ($2^1\!/_2$*-player parity game*) is a stochastic game G with the set Even(p) as the objective of the player Even, where $p : V \to [d]$ is a priority function. For brevity, we sometimes write $G(p)$ to denote the game $G(\mathrm{Even}(p))$. Similarly, a *parity Markov decision process* ($1^1\!/_2$-player parity game) is a $1^1\!/_2$-player game with the objective of the player derived from a priority function p as above.

2.3 Value functions and optimal strategies.

Consider a parity game $G(p)$. We define the *value* functions $Val_G^{\Box}, Val_G^{\Diamond} : V \to [0, 1]$ for the players Even and Odd, respectively, as follows:

$$Val_G^{\Box}(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Even}(p))(s)$$
$$Val_G^{\Diamond}(s) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Odd}(p))(s)$$

for all vertices $s \in V$. Sometimes, when the objective of a player is not clear from the context, or if it is different from the objective of the original parity game $G(p) = G(\mathrm{Even}(p))$, we make it explicit, e.g., $Val_G^{\Box}(\mathrm{Reach}(T))(s)$ stands for $\sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Reach}(T)_s)$. On the other hand, if the game G is clear from the context, we sometimes write $Val^{\Box}$ and $Val^{\Diamond}$ for the value functions $Val_G^{\Box}$ and $Val_G^{\Diamond}$, respectively.

Proposition 1. (Optimality conditions) *For every $v \in V$, if $(v, w) \in E$ then the following hold.*

1. *If $v \in V_{\Box}$ then $Val_G^{\Box}(v) \geq Val_G^{\Box}(w)$ and $Val^{\Diamond}(v) \leq Val^{\Diamond}(w)$.*
2. *If $v \in V_{\Diamond}$ then $Val_G^{\Box}(v) \leq Val_G^{\Box}(w)$ and $Val^{\Diamond}(v) \geq Val^{\Diamond}(w)$.*
3. *If $v \in V_{\bigcirc}$ then $Val_G^{\Box}(v) = \frac{1}{2}\big(\sum_{(v,w) \in E} Val_G^{\Box}(w)\big)$.*

Similar conditions hold for the value function of the player Odd.

The following is the fundamental determinacy result for stochastic parity games, due to de Alfaro and Majumdar [9]; it can be also derived from the determinacy result of Martin [20] for the zero-sum stochastic games with Borel measurable payoffs.

Theorem 3. (Determinacy [20, 9]) *In every stochastic parity game we have $Val^{\Box}(s) + Val^{\Diamond}(s) = 1$, for all vertices $s \in V$.*

A strategy σ for the player Even is *optimal* if for all vertices $s \in V$, we have:

$$Val_G^{\Box}(s) = Val_{G_\sigma}^{\Box}(s) = \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Even}(p)_s);$$

it is ε-*optimal* if we have:

$$Val_G^{\Box}(s) - \varepsilon \leq Val_{G_\sigma}^{\Box}(s) = \inf_{\pi \in \Pi} \mathrm{Pr}_s^{\sigma,\pi}(\mathrm{Even}(p)_s);$$

and it is *qualitatively optimal* if the following two conditions hold:

1. $Val^{\Box}_{G}(s) = 1$ implies $Val^{\Box}_{G_\sigma}(s) = 1$, and
2. $Val^{\Box}_{G}(s) > 0$ implies $Val^{\Box}_{G_\sigma}(s) > 0$.

Optimal, ε-optimal, and qualitatively optimal strategies for the player Odd are defined similarly.

The *almost-sure winning set* $W^{\Box}_{G}$ for the player Even in a game G is defined by $W^{\Box}_{G} = \{ v \in V : Val^{\Box}(v) = 1 \}$ and the almost-sure winning set for player Odd is $W^{\Diamond}_{G} = \{ v \in V : Val^{\Diamond}(v) = 1 \}$. The algorithmic problem of *solving qualitative stochastic parity games* is, given a game $G(p)$, to compute the sets $W^{\Box}_{G}$ and $W^{\Diamond}_{G}$. The algorithmic problem of *solving quantitative stochastic parity games* is to compute the functions $Val^{\Box}_{G}$ and $Val^{\Diamond}_{G}$. The *quantitative (resp. qualitative) decision problem* for stochastic parity games is, given a vertex $s \in V$ and a rational number $r \in (0, 1]$ as a part of the input, to determine whether $Val^{\Box}_{G}(s) \geq r$ (resp. whether $s \in W^{\Box}_{G}$).

3 Solving $1^1\!/_2$-Player Parity Games

In this section we study $1^1\!/_2$-player parity games, i.e., a subclass of stochastic parity games, such that $V_\Diamond = \emptyset$. We establish Theorem 1: first we develop an $O(d \cdot n^{3/2})$ algorithm for the qualitative solution of $1^1\!/_2$-player parity games and we argue that the quantitative solution then reduces to the quantitative solution of the simpler $1^1\!/_2$-player reachability stochastic games for which linear programming formulations are known [6, 5]. We also argue that pure memoryless optimal strategies exist. Our technique is inspired by the work of Courcoubetis and Yannakakis [6]. The key concept underlying the proofs and our algorithm is that of *controllably win recurrent vertices.*

For a vertex $v \in V$ we define $\mathcal{F}_v = \{ U \subseteq V : v \in U \text{ and } p(v) = \min(p(U)) \}$, i.e., the family of sets of vertices containing vertex v and in which v has the smallest priority. For $s, v \in V$, we define $\Omega^v_s = \{ \omega \in \Omega_s : \text{Inf}(\omega) \in \mathcal{F}_v \}$, and for $U \subseteq V$ we define $\Omega^U_s = \{ \omega \in \Omega_s : \text{Inf}(\omega) = U \}$.

Definition 1. (Controllably win recurrent vertex) *A vertex $v \in V$ in a $1^1\!/_2$-player game (G, p) is* controllably win recurrent *(c.w.r.) if $p(v)$ is even and there is a strongly connected set of vertices $U \subseteq V$, such that $U \in \mathcal{F}_v$ and there are no random edges out of set U in graph G (i.e., if $u \in U \cap V_{\bigcirc}$ and $(u, w) \in E$ imply $w \in U$). The set of vertices U as above is a* witness *for the vertex v. We write V_{cwr} for the set of all c.w.r. vertices.* ∎

Lemma 1. *For every vertex $s \in V$, and for every set of vertices $U \subseteq V$, if there is a random edge out of the set U then for every strategy $\sigma \in \Sigma$, we have* $\Pr^\sigma_s(\Omega^U_s) = 0$.

Proof. Let $(u, w) \in E$ be a random edge out of U, i.e., $u \in U \cap V_{\bigcirc}$ and $w \notin U$. Each time the random path ω^σ_s visits vertex u it then visits vertex w in the next step with probability $1/2$. Therefore, if ω^σ_s visits vertex u infinitely many times then it visits vertex w infinitely many times with probability 1. Hence, $\text{Inf}(\omega^\sigma_s) = U$ with probability 0, i.e., $\Pr^\sigma_s(\Omega^U_s) = 0$. ∎

Lemma 2. *Let $v \in V$ be a c.w.r. vertex and let $U \subseteq V$ be a witness for v. Then there is a strategy $\sigma \in \Sigma$, such that for every vertex $u \in U$, we have* $\mathrm{Pr}_u^\sigma(\Omega_u^v) = 1$.

Proof. For every vertex $u \in U \cap V_\Box$, we set $\sigma(u)$ to be a successor of u with the minimum distance to vertex v in the subgraph induced by the set U; the minimum distance to v is well defined since by definition of a witness the set U is strongly connected. Every play starting from a vertex in set U and following the strategy σ never leaves the set U because by the definition of a witness no random edges leave the set U. Moreover, in every step the distance to the vertex v is decreased with a fixed positive probability and hence from every vertex in the set U there is a fixed positive probability of reaching the vertex v in at most $|U| - 1$ steps. Therefore, the vertex v is visited infinitely many times with probability 1, and since $p(v)$ is even and by definition of a witness we have $p(v) = \min(p(U))$, we obtain $\mathrm{Pr}_u^\sigma(\Omega_u^v) = 1$. ∎

Proposition 2. *If a vertex $v \in V$ is not c.w.r. and $p(v)$ is even then for all strategies $\sigma \in \Sigma$, and for all starting vertices $s \in V$, we have* $\mathrm{Pr}_s^\sigma(\Omega_s^v) = 0$.

Proof. By Lemma 1 we have $\mathrm{Pr}_s^\sigma(\Omega_s^U) = 0$, for every $U \in \mathcal{F}_v$, and for every strategy $\sigma \in \Sigma$, because vertex v is not c.w.r. Observe that $\Omega_s^v = \bigcup_{U \in \mathcal{F}_v} \Omega_s^U$, and hence $\mathrm{Pr}_s^\sigma(\Omega_s^v) = 0$, since the set $\mathcal{F}_v$ is finite. ∎

Lemma 3. (Reach the winning set strategy) *If $W \subseteq W^\Box$ then $Val^\Box \geq Val^\Box(\mathrm{Reach}(W))$.*

Proof. Let $\sigma \in \Sigma$ be a strategy such that $\mathrm{Pr}_w^\sigma(\mathrm{Even}(p)) = 1$, for all vertices $w \in W$, and let $\sigma' \in \Sigma$ be a strategy such that $\mathrm{Pr}_v^{\sigma'}(\mathrm{Reach}(W)) = Val^\Box(\mathrm{Reach}(W))$, for all vertices $v \in V$. We define a strategy $\tau \in \Sigma$ as follows:

$$\tau(\boldsymbol{w}) = \begin{cases} \sigma'(\boldsymbol{w}) & \text{if } \boldsymbol{w} \in (V \setminus W)^+, \\ \sigma(\boldsymbol{u}) & \text{otherwise,} \end{cases}$$

where $\boldsymbol{u} \in V^*$ is the longest suffix of the sequence $\boldsymbol{w}$ starting from a vertex in the set W.

Fix a vertex $s \in V$. Then we have the following:

$$\begin{aligned} \mathrm{Pr}_s^\tau(\mathrm{Even}(p)_s) &\geq \mathrm{Pr}_s^\tau(\mathrm{Even}(p)_s \cap \mathrm{Reach}(W)_s) \\ &= \mathrm{Pr}_s^\tau(\mathrm{Even}(p)_s \mid \mathrm{Reach}(W)_s) \cdot \mathrm{Pr}_s^\tau(\mathrm{Reach}(W)_s) \\ &= \mathrm{Pr}_s^\tau(\mathrm{Reach}(W)_s), \end{aligned}$$

since $\mathrm{Pr}_s^\tau(\mathrm{Even}(p)_s \mid \mathrm{Reach}(W)_s) = \mathrm{Pr}_s^\sigma(\mathrm{Even}(p)_s) = 1$. Hence we get $Val^\Box(s) \geq Val^\Box(\mathrm{Reach}(W))(s)$. ∎

Lemma 4. *Let W be a set of vertices such that $V_{\mathrm{cwr}} \subseteq W \subseteq W^\Box$. Then we have $Val^\Box = Val^\Box(\mathrm{Reach}(W))$.*

Proof. Observe that for every vertex $s \in V$, we have $\text{Even}(p)_s = \bigcup_{v\,:\,p(v) \text{ is even}} \Omega_s^v$. Therefore for every strategy $\sigma \in \Sigma$, we have

$$\Pr_s^\sigma(\text{Even}(p)_s) = \Pr_s^\sigma\big(\bigcup \{ \Omega_s^v \mid v \,:\, p(v) \text{ is even} \} \big) = \Pr_s^\sigma\big(\bigcup \{ \Omega_s^v \mid v \in V_{\text{cwr}} \} \big)$$

where the last equality follows from Proposition 2. Note that for every $v \in V$, we have $\Omega_s^v \subseteq \text{Reach}(v)_s$ and hence $\bigcup_{v \in V_{\text{cwr}}} \Omega_s^v \subseteq \text{Reach}(V_{\text{cwr}})_s$. From the above equalities it then follows that $\Pr_s^\sigma(\text{Even}(p)_s) \leq \Pr_s^\sigma(\text{Reach}(V_{\text{cwr}})_s)$, and hence we have that $Val^\Box(s) \leq Val^\Box(\text{Reach}(V_{\text{cwr}}))(s) \leq Val^\Box(\text{Reach}(W))(s)$.

On the other hand, by Lemma 3 we have that $Val^\Box(s) \geq Val^\Box(\text{Reach}(W))(s)$. ∎

Theorem 4. (Qualitative pure memoryless strategies [4]) *Player Even in $1^1/_2$-player parity games has pure memoryless strategies for almost-sure win and for positive-probability win criteria.*

Proof. For a vertex $v \in V_{\text{cwr}}$, the strategy σ described in Lemma 2 defines a pure memoryless almost-sure winning strategy. For other vertices $v \in V_\Box$, we set $\sigma(v)$ to be a successor of vertex v with the minimum distance to the vertex set V_{cwr} in the graph G. This defines a pure memoryless winning strategy for almost-sure and positive-probability win criteria. ∎

Definition 2. (Attractor) *For a set $T \subseteq V$ of vertices in the $1^1/_2$-player game graph $G = (V, E, (V_\Box, V_\bigcirc))$ we inductively define the* attractor *$Attr_\bigcirc(G, T)$ as follows. Set $R_0 = T$, and for $k \geq 0$, set $R_{k+1} = R_k \cup \{ v \in V_\bigcirc \,:\, (v, u) \in E$ for some $u \in R_k \} \cup \{ v \in V_\Box \,:\, u \in R_k$ for all $(v, u) \in E \}$. Then $Attr_\bigcirc(G, T) = \bigcup_k R_k$.* ∎

It is easy to see that the attractor $Attr_\bigcirc(G, T)$ is the set of vertices from which player $\bigcirc$ has a strategy to reach set T in the sense of a non-stochastic 2-player game [29, 26]. For every $k \in [d]_+$, define $P_{<k} = p^{-1}([k-1])$, i.e., the set of vertices with priorities strictly *smaller* than k; $P_k = p^{-1}(k)$, i.e., the set of vertices with priority *equal* to k; and $P_{>k} = V \setminus (P_{<k} \cup P_k)$, i.e., the set of vertices with priorities strictly *bigger* than k.

The following observations lead to efficient algorithms for solving $1^1/_2$-player parity games, yielding a proof of Theorem 1.

Proposition 3. *If a set $U \subseteq V$ is a witness for a c.w.r. vertex $v \in U$ then the set U is disjoint from the set $Attr_\bigcirc(G, P_{<p(v)})$.*

Proof. By an easy induction on the stages R_k in the definition of the attractor $Attr_\bigcirc(G, T)$ one can prove that if there are no random edges leaving a set $U \subseteq V$ then $u \in U \cap Attr_\bigcirc(G, T)$ implies that $U \cap T \neq \emptyset$. Hence if the set U is a witness for the c.w.r. vertex $v \in V$ then it has to be disjoint from the set $Attr_\bigcirc(G, P_{<p(v)})$, since otherwise we have $U \cap P_{<p(v)} \neq \emptyset$ contradicting the $p(v) = \min(p(U))$ condition of a witness. ∎

Algorithm 1 **An** $O(d \cdot n^{3/2})$ **algorithm for the qualitative** $1^1\!/\!_2$**-player parity games.**

Input : $1^1\!/\!_2$-player parity game (G, p).
Output: The almost-sure winning set $W^\square$ for player Even

1. $W := \emptyset$; $G_0 := G$
2. **repeat**
   - $k :=$ smallest priority of a vertex in G
   - **if** k is odd **then** $G := G \setminus Attr_\bigcirc(G, P_k)$
   - **if** k is even **then**
     - $B :=$ Büchi-almost-sure-win(G, P_k)
     - $W := W \cup B$
     - $A := Attr_\bigcirc(G \setminus B, P_k \setminus B)$
     - $G := G \setminus (B \cup A)$

   until $G = \emptyset$
3. $W^\square :=$ Reachability-almost-sure-win(G_0, W)

Lemma 5. (If the smallest priority is odd) *Let the smallest priority k in a $1^1\!/\!_2$-player parity game $G(p)$ be odd. Then the set of c.w.r. vertices in $G(p)$ is equal to the set of c.w.r. vertices in the subgame $(G \setminus Attr_\bigcirc(G, P_k))(p)$.*

Proof. Set $A = Attr_\bigcirc(G, P_k)$. No vertex in P_k is c.w.r. because k is odd. For every vertex $v \in P_{>k}$, (from $p(v) > k$) it follows that $P_k \subseteq P_{<p(v)}$, which implies $A \subseteq Attr_\bigcirc(G, P_{<p(v)})$, and hence by Proposition 3 every c.w.r. vertex in the game $G(p)$ is in the set $V \setminus A$, and moreover, all witnesses of c.w.r. vertices are disjoint from the set A. It follows that if a vertex is c.w.r. in the game $G(p)$ then it is also c.w.r. in the subgame $(G \setminus A)(p)$. Conversely, if a vertex is c.w.r. in the subgame $(G \setminus A)(p)$ then it is also c.w.r. in the game $G(p)$ because by the definition of an attractor there are no random edges entering the set A, so a witness in the subgame is also a witness in the game $G(p)$. ∎

Lemma 6. (If the smallest priority is even) *Let the smallest priority k in a $1^1\!/\!_2$-player parity game $G(p)$ be even, and let B be the set of almost-sure winning vertices in the $1^1\!/\!_2$-player Büchi game $G(P_k)$. Then all vertices in set B are almost-sure winning in the parity game $G(p)$ and all c.w.r. vertices of priority k in the game $G(p)$ are in the set B. Let $A = Attr_\bigcirc(G \setminus B, P_k \setminus B)$. No vertex in the set A is c.w.r. in the game $G(p)$. If a vertex in the set $V \setminus (B \cup A)$ is c.w.r. in the game $G(p)$ then it is c.w.r. in the subgame $(G \setminus (B \cup A))(p)$.*

Proof. For the set B of almost-sure winning vertices in the Büchi game $G(P_k)$ there is a strategy to visit vertices of priority k infinitely often, and hence it is an almost-sure winning strategy for player Even in the game $G(p)$.

If a vertex of priority k is c.w.r. then by Lemma 2 there is an almost-sure winning strategy from the vertex, i.e., it belongs to the set B. Hence no vertex in the set $P_k \setminus B$ is c.w.r. Note that there is no edge entering the almost-sure winning set B from the set $V_\Box \setminus B$ since otherwise its source would also be almost-sure winning and hence in the set B. Therefore, we have $A \subseteq Attr_\bigcirc(G, P_k \setminus B)$, which by Proposition 3 implies that no vertex in the set $A \setminus P_k$ is c.w.r. This concludes the proof that no vertex in the set A is c.w.r. in the game $G(p)$.

Let a vertex $v \in V \setminus (B \cup A)$ be c.w.r. in the game $G(p)$. We have $P_k \setminus B \subseteq P_{<p(v)}$ because $p(v) > k$, and hence $A \subseteq Attr_\bigcirc(G, P_k \setminus B) \subseteq Attr_\bigcirc(G, P_{<p(v)})$. Let a set $U \subseteq V$ be a witness for the vertex v in the game $G(p)$. Then Proposition 3 implies that the set U is disjoint from the set $Attr_\bigcirc(G, P_{<p(v)})$, and hence the set U is disjoint from the set A. The set U is also disjoint from the set B, since otherwise Lemma 2 implies that $v \in B$. It follows that $U \subseteq V \setminus (B \cup A)$, and hence the vertex v is c.w.r. in the game $(G \setminus (B \cup A))(p)$ with the set U as a witness. ∎

Lemma 7. *The Algorithm 1 computes the almost-sure winning set $W^\Box$ for player Even in $O(d \cdot n^{3/2})$ time for binary $1\frac{1}{2}$-player parity game graphs with n vertices and d different priorities. Quantitative solution, i.e., the value function can be computed in polynomial time.*

Proof. It follows from Lemmas 5 and 6 that the set of vertices W computed in Step 2. satisfies that $V_{\text{cwr}} \subseteq W \subseteq W^\Box$. Correctness of the algorithm then follows from Lemma 4. The work done in every iteration is either $O(n)$ if the smallest priority is odd (attractor computation can be easily implemented in $O(n)$ time) or $O(n^{3/2})$ if the smallest priority is even, using the $O(n^{3/2})$ time algorithm for solving Büchi $1\frac{1}{2}$-player games of Chatterjee et al. [4]. As almost-sure reachability is a special case of the Büchi almost-sure reachability (with *absorbing* target vertices) Step 3. can also be computed in $O(n^{3/2})$ time. There are at most d iterations of the algorithm hence the total complexity is $O(d \cdot n^{3/2})$.

It follows from Lemma 4 that for every vertex $s \in V$, we have $Val^\Box(s) = Val^\Box(\text{Reach}(W^\Box))(s)$, i.e., the value of vertex s in the $1\frac{1}{2}$-player parity game $G(p)$ is its value in the reachability game $G(\text{Reach}(W^\Box))$ with the almost-sure winning set as the target set. The quantitative solution then amounts to computing the qualitative solution in time $O(d \cdot n^{3/2})$ and then quantitatively solving the reachability game $G(\text{Reach}(W^\Box))$ for which linear programming formulations are known [6, 5]; see also Section 5 for more details, including the linear program. ∎

Lemma 8. *Given a parity MDP $(G = (V, E, (V_\Box, V_\bigcirc)), p)$ let s be a vertex with even priority. Then s is controllably win recurrent if and only if the player Even has an almost sure winning strategy from vertex s in the parity MDP (G, p'), where the priority function p' is defined by:*

$$p'(s') = \begin{cases} 1 & \text{if } s \neq s' \text{ and } p(s') < p(s), \\ 2 & \text{if } s = s', \\ 3 & \text{if } s \neq s' \text{ and } p(s') \geq p(s). \end{cases}$$

Proof. If s is controllably win recurrent then there is a strategy $\sigma \in \Sigma$, such that $\Pr_s^\sigma(\Omega_s^s) = 1$ (by Lemma 2) and hence s belongs to the almost-sure winning vertex set.

If s is an almost-sure winning vertex then by Lemma 4 the set of controllably win recurrent vertices must be reachable from s. Since vertex s is the only vertex with even priority in the game (G, p'), it must be that s is controllably win recurrent in (G, p'), and hence by definition of p' it is also controllably win recurrent in the game (G, p). ∎

We use the notation $[3]_+$ MDP to denote a parity MDP with the priority function $p : V \to [3]_+$.

Lemma 9. *Given a vertex s whether $s \in W^{\Box}$ can be determined in time $O(n^{3/2} \cdot \log d)$.*

Proof. Given a vertex s whether $s \in W^{\Box}$ can be computed with a few modifications to the *procedure solve* described Section 3 of [18] which solves non-emptiness of parity automata. The key idea is to apply *binary search technique* along with *hierarchical clustering* of graphs which is described in details in Section 3 of [18]. We point out the difference below:

- Replace the procedure of *solving* $[3]_+$ *MDP's* with almost sure winning criteria for the procedure of *finding maximal strongly connected component* in the *procedure solve* (Section 3) of [18].
- Replacing the procedure of *solving MDP reachability* with almost sure winning criteria for the procedure of *reachability* in the *procedure solve* (Section 3) of [18].

The correctness follows from Lemma 8 and the correctness of *procedure solve*. The time complexity is $O(n^{3/2} \cdot \log d)$ follows from the work analysis of *procedure solve* and the fact that $[3]_+$ MDP's with almost-sure winning criteria and MDP reachability can be solved in time $O(n^{3/2})$ (follows from Lemma 7). ∎

4 Qualitative 2½-Player Parity Games

In this section we give an alternative, more elementary proof of our earlier result on solving qualitative stochastic parity games by a reduction to non-stochastic parity games [4]. The key consequence of this solution is the existence of pure memoryless qualitatively optimal strategies that we use in the next section to prove existence of pure memoryless optimal strategies.

Theorem 5. (Qualitatively optimal pure memoryless strategies [4]) *In 2½-player stochastic parity games both players have qualitatively optimal pure memoryless strategies.*

First we recall the reduction [4]. Given a 2½-player parity game $(G = (V, E, (V_\Box, V_\Diamond, V_\bigcirc)), p : V \to [d])$, we construct a 2-player parity game $\overline{G}$ with the same set $[d]$ of priorities. For every vertex $v \in V_\Box \cup V_\Diamond$, there is a vertex $\overline{v} \in \overline{V}$ with "the same" outgoing edges, i.e., $(v, u) \in E$ if and only if $(\overline{v}, \overline{u}) \in \overline{E}$. Each random vertex $v \in V_\bigcirc$ is substituted by the gadget presented in Figure 1.

More formally, the players play the following 3-step game in $\overline{G}$ from vertex $\overline{v}$ of priority $p(v)$. First, in vertex $\overline{v}$ player Odd chooses a successor $(\widetilde{v}, k)$, each of priority $p(v)$, where $k \in [p(v)+1]$ is even. Then in vertex $(\widetilde{v}, k)$ player Even chooses from at most two successors: vertex $(\widehat{v}, k-1)$ of priority $k-1$ if $k > 0$, or vertex $(\widehat{v}, k)$ of priority k if $k \leq p(v)$. Finally, in a vertex $(\widehat{v}, k)$ the choice is between all vertices $\overline{u}$ such that $(v, u) \in E$, and it belongs to player Even if k is odd, and to player Odd if k is even.

Let $\overline{U^\Box}$ and $\overline{U^\Diamond}$ be the winning sets for players Even and Odd, respectively, in the 2-player parity game $\overline{G}$. Define sets $U^\Box$ and $U^\Diamond$ of vertices in the $2^1/_2$-player parity game G by $U^\Box = \{\, v \in V \;:\; \overline{v} \in \overline{U^\Box} \,\}$, and $U^\Diamond = \{\, v \in V \;:\; \overline{v} \in \overline{U^\Diamond} \,\}$. By the determinacy of 2-player parity games [10] we have that $\overline{U^\Box} \cup \overline{U^\Diamond} = \overline{V}$, and hence $U^\Box \cup U^\Diamond = V$. Therefore, the following lemma yields the Theorem 5.

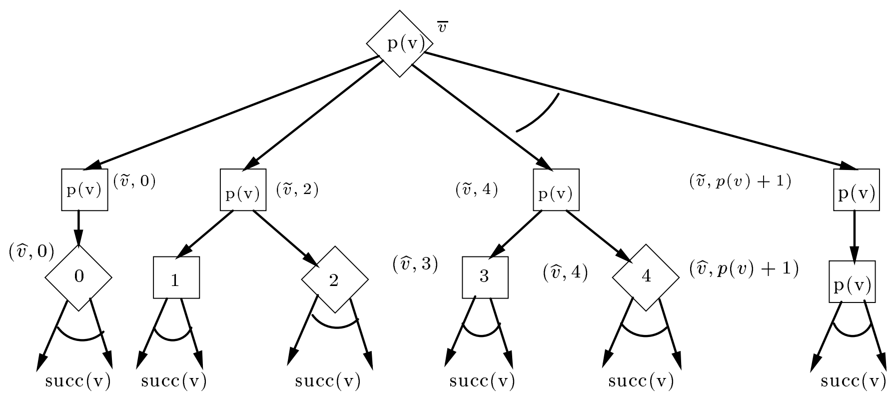

Fig. 1. The gadget for the reduction of a $2^1/_2$-player parity game to a 2-player parity game.

Lemma 10. *There are pure memoryless strategies σ and π in the stochastic game G for the player Even and Odd, respectively, such that $Val^\Box_{G_\sigma}(v) = 1$, for all $v \in U^\Box$, and $Val^\Diamond_{G_\pi}(v) > 0$, for all $v \in U^\Diamond$.*

Proof. For brevity, we say that a cycle (resp. strongly connected component, or *s.c.c.*) in a (stochastic) parity game graph is even (resp. odd) if the minimum priority of a vertex on the cycle (resp. in the s.c.c.) is even (resp. odd). By the pure memoryless determinacy of 2-player parity games [10] there are pure memoryless strategies $\overline{\sigma}$ and $\overline{\pi}$ in the game $\overline{G}$ for the player Even and Odd, respectively, such that all cycles in the subgraph $\overline{G}_{\overline{\sigma}} \restriction \overline{U^\Box}$ are even, and all cycles in the subgraph $\overline{G}_{\overline{\pi}} \restriction \overline{U^\Diamond}$ are odd [17].

We define a pure memoryless strategy σ for the player Even in the game G as follows: for all $v \in V_\Box$, if $\overline{\sigma}(\overline{v}) = \overline{u}$ then we set $\sigma(v) = u$. In a similar fashion we get a pure memoryless strategy π for the player Odd in the game G from the pure memoryless strategy $\overline{\pi}$.

Our goal is to establish that the player Even wins with probability 1 from the set $U^{\Box}$ in the game G using strategy σ, and that the player Odd wins with positive probability from the set $U^{\Diamond}$ in the game G using strategy π. By Theorem 4 it suffices to prove it for strategies σ and π only against memoryless strategies. In other words, it suffices to establish the following.

1. For all pure memoryless strategies π', every terminal s.c.c. in the graph $G_{\sigma,\pi'} \restriction U^{\Box}$ is even. For every random vertex $v \in V_{\bigcirc} \cap U^{\Box}$, all of its successors are in the set $U^{\Box}$. Hence for every strategy π', in the Markov chain $G_{\sigma,\pi'}$, if the play starts at a vertex $v \in U^{\Box}$ then it reaches some terminal strongly connected component in $U^{\Box}$ with probability 1.
2. For all pure memoryless strategies σ', every terminal s.c.c. in the graph $G_{\sigma',\pi} \restriction U^{\Diamond}$ is odd. For every random vertex $v \in V_{\bigcirc} \cap U^{\Diamond}$, at least one of its successors is in the set $U^{\Diamond}$. Hence for every strategy σ', in the Markov chain $G_{\sigma',\pi}$, from every vertex $v \in U^{\Diamond}$ there is a path from $v \in U^{\Diamond}$ to a terminal strongly connected component in $U^{\Diamond}$. Therefore for any strategy σ', in the Markov chain $G_{\sigma',\pi}$, if the play starts at any vertex $v \in U^{\Diamond}$ there is a positive probability that it reaches a terminal strongly connected component in $U^{\Diamond}$.

1. We prove that every terminal strongly connected component in the strategy subgraph $G_{\sigma,\pi'} \restriction U^{\Box}$ is even. We argue that if there is an odd terminal strongly connected component in $G_{\sigma,\pi'} \restriction U^{\Box}$ then we can construct an odd cycle in the subgraph $\overline{G}_{\overline{\sigma}} \restriction \overline{U^{\Box}}$, which is impossible because $\overline{\sigma}$ is a winning strategy for the player Even from the set $\overline{U^{\Box}}$ in the 2-player parity game $\overline{G}$. Let C be an odd terminal strongly connected component in $G_{\sigma,\pi'} \restriction U^{\Box}$, and let its minimum priority be $2r-1$, for some $r \in \mathbb{N}$. We fix a strategy $\overline{\pi}$ for the player Odd in the game $\overline{G}_{\overline{\sigma}}$ as follows. For any vertex $v \in C \cap V_{\Diamond}$ we have $\overline{\pi}(v) = \pi'(v)$. For a vertex $v \in C \cap V_{\bigcirc}$ we define the strategy as follows: if the vertex $(\widetilde{v}, 2\ell-1)$ for $\ell < r$ is in $\overline{G}_{\overline{\sigma}}$ then the strategy $\overline{\pi}$ chooses at vertex $\overline{v}$ the successor leading to $(\widetilde{v}, 2\ell-1)$. Otherwise the strategy $\overline{\pi}$ chooses at vertex $\overline{v}$ the successor leading to $(\widetilde{v}, 2r)$, and at vertex $(\widetilde{v}, 2r)$ it choses a successor shortening the distance to the fixed vertex of priority $2r-1$.
   Consider an arbitrary cycle in the subgraph $\overline{G}_{\overline{\sigma}\,\overline{\pi}} \restriction \overline{C}$, where $\overline{C}$ is the set of vertices in the gadgets of vertices in C. There are two cases.
   - If there is at least one vertex $(\widetilde{v}, 2\ell-1)$ with $\ell < r$ on the cycle then the minimum priority on the cycle is odd.
   - Otherwise, in all vertices choices shortening the distance to the vertex with priority $2r-1$ are taken and hence priority $2r-1$ is visited and all other priorities on the cycle are $\geq 2r-1$, so $2r-1$ is the minimum priority on the cycle.

   This will construct an odd cycle in the subgraph $\overline{G}_{\overline{\pi}} \restriction \overline{U^{\Box}}$ which is a contradiction.
   Now we argue that for every random vertex $v \in V_{\bigcirc} \cap U^{\Box}$, all of its successors are in $U^{\Box}$. Otherwise, the player Odd in the vertex $\overline{v}$ of the game $\overline{G}$ can choose the successor $(\widetilde{v}, 0)$ and then a successor to its winning set $\overline{U^{\Diamond}}$, which contradicts the assumption that the strategy $\overline{\sigma}$ is a winning strategy for the player Even in the game $\overline{G}$.

2. The proofs for the player Odd are similar. We only briefly sketch the key steps. Let C be an even terminal strongly connected component in $G_{\sigma',\pi} \upharpoonright U^{\diamond}$, and let its minimum priority be $2r$, for some $r \in \mathbb{N}$. We fix a strategy $\overline{\sigma}$ for the player Even in the game $\overline{G_{\overline{\pi}}}$ as follows. For vertices $v \in C \cap V_{\Box}$ we have $\overline{\sigma}(v) = \sigma'(v)$. For vertices $v \in C \cap V_{\bigcirc}$ we define the strategy as follows: if at a vertex $\overline{v}$ the strategy $\overline{\pi}$ chooses a vertex $(\widetilde{v}, 2\ell - 2)$, such that $\ell \leq r$, then the strategy $\overline{\sigma}$ chooses the successor $(\hat{v}, 2\ell - 2)$. Otherwise, $\ell > r$ and then the strategy $\overline{\sigma}$ chooses a successor which shortens the distance to a fixed vertex with priority $2r$.
   Consider an arbitrary cycle in the subgraph $\overline{G_{\overline{\sigma}\,\overline{\pi}}} \upharpoonright \overline{C}$, where $\overline{C}$ is the set of vertices in the gadgets of vertices in C. There are two cases.
   - If there is at least one vertex $(\widetilde{v}, 2\ell)$ with $\ell < r$ on the cycle then the minimum priority on the cycle is even.
   - Otherwise, in all vertices choices shortening the distance to the vertex with priority $2r$ are taken and hence priority $2r$ is visited and all other priorities on the cycle are $\geq 2r$, so $2r$ is the minimum priority on the cycle.

   This will construct an even cycle in the subgraph $\overline{G_{\overline{\pi}}} \upharpoonright \overline{U^{\diamond}}$ which is a contradiction.
   Finally, for every random vertex $v \in V_{\bigcirc} \cap U^{\diamond}$, at least one successor must be in the set $U^{\diamond}$. Otherwise if both the successors of v are in $U^{\Box}$ it follows from the construction of the gadget that $\overline{v} \in Attr_{\Box}(\overline{G}, \overline{U^{\Box}})$. In other words there is a strategy for player Even in the 2-player game to reach the set $\overline{U^{\Box}}$ from $\overline{v}$. Hence $\overline{v} \in \overline{U^{\Box}}$ which is a contradiction. ∎

5 Quantitative 2½-Player Parity Games

In this section we prove that *pure memoryless optimal strategies* exist for quantitative stochastic parity games, i.e., we establish Theorem 2.

Lemma 11. (Subgame strategy) *Let $G' = G \upharpoonright U$ be a subgame of the game G, and let $r < 1$. If $Val^{\Box}_{G'}(u) = 1$, for all $u \in U$, and if $Val^{\Box}_{G}(w) \geq r$, for all edges $(u, w) \in E$ going out of the set U, such that $u \in V_{\diamond}$, then $Val^{\Box}_{G}(u) \geq r$, for all $u \in U$.*

Proof. By the assumption that for all vertices $u \in U$, we have $Val^{\Box}_{G'}(u) = 1$, so by Theorem 5, there is a pure memoryless strategy σ' for the player Even in the game G', such that $\inf_{\pi \in \Pi} \mathrm{Pr}_{u}^{\sigma',\pi}(\mathrm{Even}(p)) = 1$. Let $\overline{\sigma}$ be an ε-optimal strategy for the player Even in the game G.

We define a strategy σ for the player Even in the game G as follows:

$$\sigma(\boldsymbol{w} \cdot v) = \begin{cases} \sigma'(v) & \text{if } \boldsymbol{w} \cdot v \in U^{+}, \\ \overline{\sigma}(\boldsymbol{u} \cdot v) & \text{otherwise,} \end{cases}$$

where $\boldsymbol{u} \in V^{*}$ is the longest suffix of the sequence $\boldsymbol{w}$ starting from a vertex which is not in the set U.

Fix a vertex $u \in U$. Define $\Omega_U = u \cdot U^\omega$, the set of infinite sequences of vertices in the set U, starting from the vertex u. Set $\overline{\Omega_U} = \Omega_u \setminus \Omega_U$. Let $\pi \in \Pi$ be a strategy for the player Odd in the game G. By the assumption that $\inf_{\pi \in \Pi} \Pr_u^{\sigma',\pi}(\mathrm{Even}(p)) = 1$, we have that $\Pr_u^{\sigma',\pi}(\mathrm{Even}(p) \mid \Omega_U) = 1$. For every play $\omega = \langle u = v_0, v_1, v_2, \ldots \rangle \in \overline{\Omega_U}$, there is the smallest $i \in \mathbb{N}$, such that $v_i \notin U$. Note that if in the play ω the player Even uses the strategy σ then by the ε-optimality of the strategy $\overline{\sigma}$, we have $\Pr_u^{\sigma,\pi}(\mathrm{Even}(p) \mid \overline{\Omega_U}) \geq r - \varepsilon$. Therefore, for every strategy $\pi \in \Pi$ of the player Odd, we have:

$$\begin{aligned}\Pr_u^{\sigma,\pi}(\mathrm{Even}(p)) &= \Pr_u^{\sigma,\pi}(\mathrm{Even}(p) \mid \Omega_U) \cdot \Pr_u^{\sigma,\pi}(\Omega_U) \\ &+ \Pr_u^{\sigma,\pi}(\mathrm{Even}(p) \mid \overline{\Omega_U}) \cdot \Pr_u^{\sigma,\pi}(\overline{\Omega_U}) \\ &\geq \Pr_u^{\sigma,\pi}(\Omega_U) + (r - \varepsilon) \cdot \Pr_u^{\sigma,\pi}(\overline{\Omega_U})\end{aligned}$$

Note that $\Pr_u^{\sigma,\pi}(\Omega_U) + \Pr_u^{\sigma,\pi}(\overline{\Omega_U}) = 1$, so we get $\Pr_u^{\sigma,\pi}(\mathrm{Even}(p)) \geq r - \varepsilon$, for all strategies π for the player Odd. This implies that $Val_G^{\Box}(u) \geq r$. ∎

Definition 3. (Locally optimal strategy) *We say that a pure memoryless strategy τ for a player $P \in \{\Box, \Diamond\}$ (i.e., for player Even and Odd, respectively) is* locally optimal *in the game G if for every vertex $v \in V_P$, we have $Val_G^P(\tau(v)) = Val_G^P(v)$.* ∎

We establish the following strengthening of Theorem 5.

Lemma 12. (Locally and qualitatively optimal strategies) *Both players have pure memoryless qualitatively optimal strategies that are also locally optimal.*

Proof. We prove this and the next lemma for the player Odd; the proofs for the player Even are similar.

Let G' be the game obtained from the game G by removing all the edges $(v, w) \in E$, such that $v \in V_\Diamond$, and $Val_G^\Diamond(v) > Val_G^\Diamond(w)$. Observe that the game G' differs from the game G only in that in the game G' the player Odd can only use the locally optimal strategies in the game G, while the player Even can use all his strategies in the game G. Therefore, it suffices to prove for every vertex $v \in V$, that if $Val_G^\Diamond(v) = 1$ then $Val_{G'}^\Diamond(v) = 1$, and that if $Val_G^\Diamond(v) > 0$ then $Val_{G'}^\Diamond(v) > 0$.

By Theorem 5 there is a pure memoryless qualitatively optimal strategy π for the player Odd in the game G. Note that there are no edges going out of the set $W_G^\Diamond$ in the strategy subgraph G_π. Therefore, for every edge $(v, w) \in E_\pi$, such that $v \in W_G^\Diamond$, we have $Val_G^\Diamond(v) = Val_G^\Diamond(w) = 1$, and hence the strategy π is also a valid strategy in the game G', and for all vertices $v \in W_G^\Diamond$, we have $Val_{G'}^\Diamond(v) = Val_{G'_\pi}^\Diamond(v) = 1$.

We now prove the other claim which can be restated as follows: there is no vertex $v \in V$, such that $Val_G^\Box(v) < 1$ and $Val_{G'}^\Box(v) = 1$. Define the number $r = \min\{\, Val_G^\Box(v) \;:\; Val_{G'}^\Box(v) = 1 \,\}$. For the sake of contradiction assume that $r < 1$.

By Theorem 5 there is a pure memoryless qualitatively optimal strategy σ for the player Even in the game G'. Note that there are no edges going out of the set $W^{\Box}_{G'}$ in the strategy subgraph G'_{σ}. It follows that $G \restriction W^{\Box}_{G'}$ is a subgame of the game G. By definition of the set $W^{\Box}_{G'}$, we have $Val^{\Box}_{G'}(v) = 1$, for all vertices $v \in W^{\Box}_{G'}$. Moreover, by the optimality condition for a vertex $v \in V_{\Diamond}$ in the game G we have that for every edge $(v, u) \in E$ going out of the set $W^{\Box}_{G'}$ (and hence not present in the game graph G') we must have $Val^{\Diamond}_{G}(v) > Val^{\Diamond}_{G}(u)$, i.e., $Val^{\Box}_{G}(u) > Val^{\Box}_{G}(v) \geq r$. Applying Lemma 11 to the subgame $G \restriction W^{\Box}_{G'}$ we get that $Val^{\Box}_{G}(v) > r$, for every vertex $v \in W^{\Box}_{G'}$. This, however is a contradiction with the assumption that there is a vertex $v \in W^{\Box}_{G'}$, such that $Val^{\Box}_{G}(v) = r$. ∎

Lemma 13. (Locally and qualitatively optimal strategies are optimal) *If a pure memoryless strategy is locally optimal and qualitatively optimal then it is optimal.*

Proof. Let π be a pure memoryless strategy for player Odd that is locally optimal and qualitatively optimal. The strategy subgraph G_{π} of the strategy π is a Markov decision process so by Lemma 4 we have $Val^{\Box}_{G_{\pi}}(v) = Val^{\Box}_{G_{\pi}}(\mathrm{Reach}(W^{\Box}_{G_{\pi}}))(v)$, and those values are the unique solution to the following linear program [6, 5]: minimize $\left(\sum_{v \in V} x_v\right)$, subject to:

$$
\begin{array}{ll}
x_v = \frac{1}{2}\left(\sum_{(v,w)\in E} x_w\right) & v \in V_{\bigcirc} \\
x_v = \pi(x_v) & v \in V_{\Diamond} \\
x_v \geq x_w & v \in V_{\Box}, (v, w) \in E \\
x_v \geq 0 & v \in V \\
x_v = 1 & v \in W^{\Box}_{G_{\pi}}
\end{array}
$$

It follows from the qualitative optimality of the strategy π that $W^{\Box}_{G_{\pi}} = W^{\Box}_{G}$. The local optimality of π implies that for every vertex $v \in V_{\Diamond}$, we have $Val^{\Box}_{G}(v) = Val^{\Box}_{G}(\pi(v))$. Therefore, the valuation $x_v := Val^{\Box}_{G}(v)$, for all $v \in V$, is a feasible solution of the linear program, so we have $Val^{\Box}_{G_{\pi}} \leq Val^{\Box}_{G}$. This implies that π is an optimal strategy for player Odd in the game G. ∎

Theorem 2 follows from Lemmas 12 and 13.

Proof (of Corollary 1). By Theorem 2 both players have optimal pure memoryless strategies. There are only exponentially many pure memoryless strategies and once a strategy for a player is fixed we have a $1\frac{1}{2}$-player game whose values can be computed in polynomial time by Theorem 1. Hence we get the NP $\cap$ co-NP upper bound and an exponential time algorithm to compute the exact values.

By Theorem 2 the values are a solution of a system of linear equations with rational coefficients [5] and hence they are rational. ∎

The previous best algorithm known for this problem is due to de Alfaro and Majumdar (cf. the full version of their article [9]). It works for concurrent stochastic parity games, and uses a decision procedure for the first-order theory of the reals with addition and multiplication [25]. The algorithm computes

ε-approximations of the values and it has a doubly exponential running time bound. Our algorithm computes in exponential time the *exact* values for the perfect information stochastic parity games.

Acknowledgments.

We thank Rupak Majumdar for clarifying the complexity of the algorithm for solving concurrent stochastic parity games [9], and an anonymous referee for pointing out a polynomial-time solution of qualitative $1^1/_2$-player parity games by a reduction to limiting-average MDP's. We thank Orna Kupferman for explaining the details of the procedure for checking nonemptiness of parity automata [18].

References

1. A. Arnold and D. Niwiński. *Rudiments of μ-Calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 2001.
2. A. Arnold, A. Vincent, and I. Walukiewicz. Games for synthesis of controllers with partial observation. *Theoretical Computer Science*, 303(1):7–34, 2003.
3. J. Bradfield. The modal μ-calculus alternation hierarchy is strict. *Theoretical Computer Science*, 195(2):133–153, 1998.
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. A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
6. C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *ICALP'90*, volume 443 of *LNCS*, pages 336–349. Springer, 1990.
7. L. de Alfaro and T. A. Henzinger. Concurrent ω-regular games. In *LICS'00*, pages 141–154. IEEE Computer Society Press, 2000.
8. L. de Alfaro, T. A. Henzinger, and O. Kupferman. Concurrent reachability games. In *FOCS'98*, pages 564–575. IEEE Computer Society Press, 1998.
9. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. In *STOC'01*, pages 675–683. ACM Press, 2001. Full version to appear in the *Journal of Computer and Systems Sciences*. Elsevier.
10. E. A. Emerson and C. Jutla. Tree automata, μ-calculus and determinacy. In *FOCS'91*, pages 368–377. IEEE Computer Society Press, 1991.
11. E. A. Emerson, C. S. Jutla, and A. P. Sistla. On model-checking for fragments of μ-calculus. In *CAV'93*, volume 697 of *LNCS*, pages 385–396. Springer, 1993.
12. 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.
13. J. Filar and K. Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
14. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games. A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002.
15. Y. Gurevich and L. Harrington. Trees, automata, and games. In *STOC'82*, pages 60–65. ACM Press, 1982.
16. M. Jurdziński. Deciding the winner in parity games is in UP $\cap$ co-UP. *Information Processing Letters*, 68(3):119–124, 1998.

17. M. Jurdziński. Small progress measures for solving parity games. In *STACS'00*, volume 1770 of *LNCS*, pages 290–301. Springer, 2000.
18. V. King, O. Kupferman, and M. Y. Vardi. On complexity of parity word automata. In *FoSSaCS'01*, volume 2030 of *LNCS*, pages 276–286. Springer, 2001.
19. M. L. Littman, T. Dean, and L. P. Kaelbling. On the complexity of solving Markov decision problems. In *UAI'95*, pages 394–402. Morgan Kaufmann, 1995.
20. D. A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, 1998.
21. A. Puri. *Theory of Hybrid Systems and Discrete Event Systems*. PhD thesis, University of California, Berkeley, 1995.
22. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.
23. 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.
24. L. S. Shapley. Stochastic games. *Proceedings Nat. Acad. of Science USA*, 39:1095–1100, 1953.
25. A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley and Los Angeles, 1951.
26. W. Thomas. Languages, automata, and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
27. J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In *CAV'00*, volume 1855 of *LNCS*, pages 202–215. Springer, 2000.
28. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1953.
29. W. Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200:135–183, 1998.
30. U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.