



A Direct Reduction from Stochastic Parity Games to Simple Stochastic Games

Raphaël Berthon 

RWTH Aachen University, Germany

Joost-Pieter Katoen 

RWTH Aachen University, Germany

Zihan Zhou 

National University of Singapore, Singapore

Abstract

Significant progress has been recently achieved in developing efficient solutions for simple stochastic games (SSGs), focusing on reachability objectives. While reductions from stochastic parity games (SPGs) to SSGs have been presented in the literature through the use of multiple intermediate game models, a direct and simple reduction has been notably absent. This paper introduces a novel and direct polynomial-time reduction from quantitative SPGs to quantitative SSGs. By leveraging a gadget-based transformation that effectively removes the priority function, we construct an SSG that simulates the behavior of a given SPG. We formally establish the correctness of our direct reduction. Furthermore, we demonstrate that under binary encoding this reduction is polynomial, thereby directly corroborating the known $\mathbf{NP} \cap \mathbf{coNP}$ complexity of SPGs and providing new understanding in the relationship between parity and reachability objectives in turn-based stochastic games.

2012 ACM Subject Classification Theory of computation \rightarrow Probabilistic computation

Keywords and phrases stochastic games, parity, reduction

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2025.9

Related Version *Full Version*: <https://arxiv.org/abs/2506.06223> [6]

Funding *Raphaël Berthon*: Funded by DFG Project POMPOM (KA 1462/6-1).

1 Introduction

Stochastic games (SGs) are a broadly used framework for decision-making under uncertainty with a well-defined objective. They aim at modeling two important aspects under which sequential decisions must be made: turn-based interaction with an adversary environment, and dealing with randomness. Stochastic games have been introduced long ago by Shapley [40] as important models, and many variations have been shown to be inter-reducible and in $\mathbf{NP} \cap \mathbf{coNP}$ [18, 2, 9]. SGs find their applications in various fields, including artificial intelligence [33], economics [1], operations research [20], or providing tools to graph theory [41]. Moreover, *Markov Decision Processes* (MDP) [38], a foundational framework for modeling decision-making in stochastic environments, are special cases of SGs, where one of the players has no states under control.

While we make use of simple stochastic games (SSGs) with reachability objectives, we focus on the specific case of stochastic *parity* games (SPGs), which are zero-sum, and where the set of winning runs is ω -regular. Solving such a game consists in finding an optimal strategy and determining its winning probability.

Take, for instance, the case of a market competition, where two firms Alpha and Beta try to expand their market share. From the standpoint of each firm, the other acts as a direct competitor, and therefore we assume the players are *adversarial*. States represent the relative valuation of the firms over a sustained period of time, and business decisions



© Raphaël Berthon, Joost-Pieter Katoen, and Zihan Zhou;
licensed under Creative Commons License CC-BY 4.0

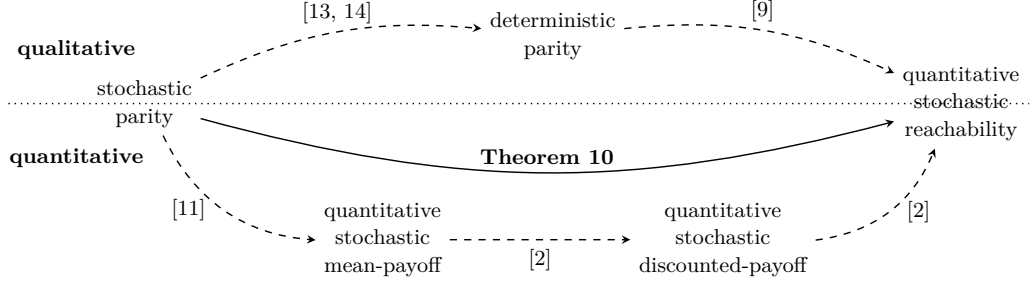
36th International Conference on Concurrency Theory (CONCUR 2025).

Editors: Patricia Bouyer and Jaco van de Pol; Article No. 9; pp. 9:1–9:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Reducing SPGs to SSGs.

(made by either firm Alpha or Beta) and the fluctuation of market shares, policy changes, and external forces (modeled as randomness) lead to transitions between game states. Firm Alpha is interested in keeping its share above a set key threshold, say 50%. We distinguish between three priorities, leading to a *zero-sum* game:

- *Priority 0*: Alpha's market share is significantly above 50%.
- *Priority 1*: Alpha's market share is significantly below 50%.
- *Priority 2*: Alpha's market share fluctuates around 50%. While this is sustainable without major fluctuation, this is not sustainable if the only other fluctuation is Alpha's share regularly dropping below 50% (priority 1).

If the minimum priority visited infinitely often is 0, Alpha can manage in the long term to regularly dominate the market, recovering any loss that occurred in the meantime. If the minimum priority visited infinitely often is 1, despite any temporary success, Alpha's market share will stay near or below 50% in the long run, which is not sustainable for the firm. Finally, if the minimum priority visited infinitely often is 2, Alpha and Beta will eventually find an even balance point.

A variety of algorithms have been considered for (reachability) SSGs [17, 21, 3, 37], which we present in our related work section below. In Markov chains (MCs), ω -regular objectives can be reduced directly to reachability objectives [4]. A similar reduction exists from MDPs with ω -regular objectives to reachability objectives and has been used extensively, for example in [22]. This means that solvers often focus on optimizing specifically the computation of reachability probabilities. Such a direct reduction is lacking for SPGs. To reduce quantitative SPGs to SSGs, some intermediate steps are necessary, via a reduction to stochastic mean-payoff and stochastic discounted-payoff games [11, 2] (see the lower part of Figure 1), making this approach less appealing. For qualitative solutions, a translation via deterministic parity games (i.e. with no random states) exists [13, 14, 9], see the upper part of Figure 1.

Outline and Contribution

In this paper, we propose a *direct* reduction from SPGs to SSGs with a reachability objective (see the solid arrow in Figure 1). To that end, we leverage a gadget whose structure comes from [9], but where we use new probability values, to reduce deterministic parity games to quantitative SSGs. Given an SPG G , where a parity condition is satisfied if the minimum priority seen infinitely often is even, we show in Section 3.1 how to use the gadget to transform G into an SSG \tilde{G} . This introduces two new sink states, one winning and the other losing for the reachability objective. Parity values are removed, and every transition going to a state that used to be even (respectively odd) now has a small chance to go to the winning (resp.

losing) sink. We scale the probabilities, with lower parity values yielding a higher chance to go to a sink. Theorem 10 ensures that any optimal strategy in \tilde{G} is also optimal in G (the reciprocal may not be true). We can then compute an optimal memoryless strategy in \tilde{G} , and compute its value in G .

We show in Theorem 13 that under binary encoding, our reduction is polynomial. We thus reobtain in a direct way the classical result that solving quantitative SPGs is in the complexity class $\mathbf{NP} \cap \mathbf{coNP}$ [2, 9]. While the complexity remains the same as for existing algorithms, and the values used in the reduction make it unlikely to be very efficient in practice, this new approach implies that any efficient SSG solver can be used for SPGs. The direct reduction was already conjectured to exist by Chatterjee and Fijalkow in [9], but as expected, proving its correctness is challenging, and involves the computations of very precise probability bounds. Despite the inspiration drawn from a known gadget, the technical depth of this paper resides in the intricate and novel proofs for the correctness of our reduction. In addition, our direct reduction gives new insights into the relationship between SSGs and SPGs.

In Section 2, we present the necessary background knowledge. In Section 3, we define the gadget and present some related results, which we use in Section 4 to define the reduction properly, and to show its correctness and its complexity. We give some concluding remarks in Section 5.

Related Work

Stochastic parity games, mean-payoff games and discounted payoff games can all be reduced to SSGs [26, 43], and this also applies to their stochastic extensions, namely stochastic parity games [9], stochastic mean payoff games and stochastic discounted payoff games [2]. SSGs also find their applications in the analysis of MDPs, serving as abstractions for large MDPs [28]. The amount of memory required to solve stochastic parity games has been studied in [7].

Various extensions have been considered within this family of inter-reducible stochastic games. Introducing more than two players allows for the analysis of Nash equilibria [15, 42]. Using continuous states can provide tools to represent timed systems [34]. Multi-objective approaches have been employed to synthesize systems that balance average expected outcomes with worst-case guarantees [16]. Parity objectives are significant in many of these scenarios where long-run behavior is relevant, but the classical reduction to SSGs cannot be directly applied.

Common approaches to solving SSGs, as presented in [17], include value iteration (VI), strategy iteration (SI), and quadratic programming, but are all exponential in the size of the SSG. These approaches have been widely studied on MDPs, where recent advancements have been made to apply VI with guarantees, using interval VI [5], sound VI [39], and optimistic VI [24]. Interestingly, optimistic VI does not require an a priori computation of starting vectors to approximate from above. Similar ideas have been lifted to SSGs: Eisentraut et al. [21] introduce a VI algorithm for under- and over-approximation sequences, as well as the first practical stopping criterion for VI on SSGs. Optimistic VI has been adapted to SSGs [3], and a novel bounded VI with the concept of widest path has been introduced in [37]. A comparative analysis [29] suggests VI and SI are more efficient. Storm [25] and PRISM [31] are two popular model checkers incorporating different variants of VI and SI, and both employ VI as the default algorithm for solving MDPs. PRISM-games [30] exploits VI for solving SSGs.

For SPGs, we distinguish three main approaches. Chatterjee et al. [10] use a strategy improvement algorithm requiring randomized sub-exponential time. With n game states and d priorities, the expected running time is in $2^{O(\sqrt{dn \log(n)})}$. The probabilistic game solver GIST [12] reduces qualitative SPGs to deterministic parity games (DPG), and benefits from several quasi-polynomial algorithms for DPGs [27, 36, 32] since the breakthrough made by Calude et al. [8], but this approach is unlikely to achieve polynomial running time [19]. Hahn et al. [23] reduce SPGs to SSGs, allowing the use of reinforcement learning to approximate the values without knowing the game's probabilistic transition structure. Their reduction is only proven correct in the limit.

2 Preliminaries

Our notations on Markov chains and stochastic games on graphs mainly come from [4].

2.1 Discrete-Time Markov Chains

A *discrete distribution* over a countable set \mathcal{A} is a function $\mu : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ with $\sum_{a \in \mathcal{A}} \mu(a) = 1$. The *support* of the discrete distribution μ is defined as $\text{supp}(\mu) \triangleq \{a \in \mathcal{A} \mid \mu(a) > 0\}$. We denote the set of all discrete distributions over \mathcal{A} with $\mathbb{D}(\mathcal{A})$.

A *discrete-time Markov Chain (MC)* \mathcal{M} is a tuple $\mathcal{M} = (V, \delta, v_I)$ where V is a finite set of states, $\delta : V \rightarrow \mathbb{D}(V)$ is a probabilistic transition function, and $v_I \in V$ is the initial state. Given $\delta(v) = \mu$ with $\mu(v') = p$, we write $\delta(v, v') = p$. For $S \subseteq V$ and $v \in V$, let $\delta(v, S) = \sum_{s \in S} \delta(v, s)$.

An infinite sequence $\pi = v_0 v_1 \dots \in V^\omega$ is an *infinite path* through MC \mathcal{M} if $\delta(v_i, v_{i+1}) > 0$ for all $i \in \mathbb{N}$. We denote all infinite paths that start from state $v \in V$ with $\text{Paths}(v)$. *Prefixes* of infinite path $\pi = v_0 v_1 \dots \in V^\omega$ are $\{v_0 \dots v_i \mid i \in \mathbb{N}\}$ and are *finite paths*. We denote all finite paths that start from state $v \in V$ with $\text{Paths}^*(v)$. The set of infinitely often visited states in $\pi = v_0 v_1 \dots \in V^\omega$ is defined as $\text{inf}(\pi) = \{v \in V \mid \forall n \in \mathbb{N}, \exists k \in \mathbb{N} \text{ s.t. } v_{n+k} = v\}$.


The probability Pr of a finite path $\pi = v_0 v_1 \dots v_n \in V^*$ is given by $\prod_{i \in [0, n-1]} \delta(v_i, v_{i+1})$. The set of infinite paths that start with a given finite path is called a *cylinder*, and as in [4], we extend the probability of cylinders in a unique way to all *measurable sets* of V^ω .

Reachability Probabilities

Let $\mathcal{M} = (V, \delta, v_I)$ be an MC. For target states $T \subseteq V$ and starting state $v_0 \in V$, the event of reaching T is defined as $\text{Reach}(T) = \{v_0 v_1 \dots \in V^\omega \mid \exists i \in \mathbb{N}, v_i \in T\}$. The probability to reach T from v_0 is defined as $\text{Pr}^{v_0}(\text{Reach}(T)) = \text{Pr}(\{\hat{\pi} \mid \hat{\pi} \in \text{Paths}^*(v_0) \cap ((V \setminus T)^* T)\})$.

Let variable x_v denote the probability of reaching T from any $v \in V$. Whether T is reachable from a given state v can be determined using standard graph analysis. Let $\text{Pre}^*(T)$ denote the set of states from which T is reachable. If $v \notin \text{Pre}^*(T)$, then $x_v = 0$. If $v \in T$, then $x_v = 1$. Otherwise, $x_v = \sum_{u \in V \setminus T} \delta(v, u) \cdot x_u + \sum_{w \in T} \delta(v, w)$. This is equivalent to a linear equation system, formalized as follows:

► **Theorem 1** (Reachability Probability of Markov Chains [4]). *Given MC $\mathcal{M} = (V, \delta, v_I)$ and target states $T \subseteq V$, let $V_Q = \text{Pre}^*(T) \setminus T$, $\mathbf{A} = (\delta(v, v'))_{v, v' \in V_Q}$ and $\mathbf{b} = (b_v)_{v \in V_Q} = (\delta(v, T))_{v \in V_Q}$. Then, the vector $\mathbf{x} = (x_v)_{v \in V_Q}$ with $x_v = \text{Pr}^v(\text{Reach}(T))$ is the **unique** solution of the linear equation system $\mathbf{x} = \mathbf{A} \cdot \mathbf{x} + \mathbf{b}$.*



■ **Figure 2** An example stochastic arena G .

Limit Behavior

Let MC $\mathcal{M} = (V, \delta, v_I)$. A set $L \subseteq V$ is *strongly connected* if for all pairs of states $v, v' \in L$, v and v' are mutually reachable. Hence a singleton $\{v\}$ is strongly connected if $\delta(v, v) > 0$. Set $L \subseteq V$ is a *strongly connected component (SCC)* if it is maximally strongly connected, i.e., there does not exist another set $L' \subseteq V$ and $L \subsetneq L'$ such that L' is strongly connected. $L \subseteq V$ is a *bottom SCC (BSCC)* if L is a SCC and there is no transition leaving L , i.e., there does not exist $v \in L, v' \in V \setminus L$ such that $\delta(v, v') > 0$. We denote the set of BSCCs in MC \mathcal{M} with $BSCC(\mathcal{M})$.

The limit behavior of an MC regarding the infinitely often visited states is captured by the following theorem.

► **Theorem 2** (Limit behavior of Markov Chains [4]). *For MC $\mathcal{M} = (V, \delta, v_I)$, it holds that $\Pr\{\pi \in Paths(v_I) \mid \inf(\pi) \in BSCC(\mathcal{M})\} = 1$.*

2.2 Stochastic Games


A *stochastic arena* G is a tuple $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$, where (V, E) is a directed graph, with a finite set of vertices V , partitioned as $V_\exists \uplus V_\forall \uplus V_R = V$, and a set of edges $E \subseteq V \times V$. The probabilistic transition function Δ is such that for all $v_r \in V_R$, $\Delta(v_r)$ is a distribution over V , and for $v \in V_\exists \uplus V_\forall, v'$, we have $(v_r, v) \in E$ if and only if $v \in \text{supp}(\Delta(v_r))$. We usually uncurry $\Delta(v_r)(v)$ and write $\Delta(v_r, v)$.

Without loss of generality, we assume each vertex has at least one successor. This property is called *non-blocking*. The finite set V of vertices is partitioned into three sets: V_\exists – vertices where Eve chooses the successor, V_\forall – vertices where Adam chooses the successor, and V_R are the random vertices. A stochastic arena is a Markov Decision Process (MDP) if either $V_\exists = \emptyset$ or $V_\forall = \emptyset$, and an MC if both $V_\exists = \emptyset$ and $V_\forall = \emptyset$.

Figure 2 illustrates a stochastic arena G . Square-shaped vertex v_3 is a vertex in V_\exists where Eve chooses the successor, pentagon-shaped vertex v_4 is in V_\forall where Adam chooses the successor, and the circular vertices $V_R = \{v_0, v_1, v_2, v_5\}$ are random. Edges from random vertices are annotated with probabilities from Δ .

Strategies

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. A *strategy* σ of Eve is a function $\sigma : V^* \cdot V_\exists \rightarrow \mathbb{D}(V)$, such that for all $v_0 v_1 \dots v_n \in V^* \cdot V_\exists$, we have $\sigma(v_0 v_1 \dots v_n, v_{n+1}) > 0$ implies $(v_n, v_{n+1}) \in E$. A strategy γ of Adam is defined analogously. We denote the sets of all strategies of Eve and Adam with Σ_\exists^A and Σ_\forall^A respectively.



■ **Figure 3** Sub-arena $G_{\sigma, \gamma}$ induced by strategies $\sigma = [v_4 \mapsto v_5]$ and $\gamma = [v_3 \mapsto v_5]$.

A strategy σ of Eve is a *pure memoryless strategy*, if for all $w, w' \in V^*$ and $v \in V_\exists$, $\sigma(w \cdot v) = \sigma(w' \cdot v)$ and the support of this distribution is a singleton. A pure memoryless strategy γ of Adam is defined analogously. We denote the sets of pure memoryless strategies of Eve and Adam with Σ_\exists and Σ_\forall respectively.

In a stochastic arena G , when Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma_\exists$ and $\gamma \in \Sigma_\forall$ respectively, the arena $G_{\sigma, \gamma} = ((V, E'), (V_\exists, V_\forall, V_R), \Delta)$ results. Here, the new edge set E' is such that for all $u \in V_\exists$, $(u, v) \in E'$ if and only if $\sigma(u) = v$, and for all $u \in V_\forall$, $(u, v) \in E'$ if and only if $\gamma(u) = v$. We refer to such arenas obtained by fixing pure memoryless strategies as sub-arenas. In fact, given a fixed starting vertex $v_I \in V$, we often view the sub-arena $G_{\sigma, \gamma}$ as an MC $\mathcal{M}_{\sigma, \gamma} = (V, \delta, v_I)$, where the state space is the vertex set V in G , and the transition function δ combines deterministic moves indicated by strategies σ and τ , and the transition function Δ defined on random vertices:

$$\delta(u, v) = \begin{cases} \Delta(u, v) & \text{if } u \in V_R, \text{ or } u \in V_\exists, \sigma(u) = v \text{ or } u \in V_\forall, \gamma(u) = v \\ 0 & \text{otherwise} \end{cases}$$

We continue with the stochastic arena G from Figure 2. Fixing strategy $\sigma = [v_3 \mapsto v_5]$ for Eve and $\gamma = [v_4 \mapsto v_5]$ for Adam induces the sub-arena $G_{\sigma, \gamma}$, as shown in Figure 3.

Winning Objectives

A *play* of G is an infinite sequence of vertices $\pi = v_0 v_1 \dots \in V^\omega$ where for all $i \in \mathbb{N}$, $(v_i, v_{i+1}) \in E$. We denote the set of all plays of G with Π_G , or in short Π when G is clear from the context.

Let G be a stochastic arena. A *winning objective* for Eve is defined as a set of plays $\mathcal{O} \subseteq \Pi$. As we study zero-sum games, the winning objectives of the two players are complementary. The winning objective for Adam is thus $\Pi \setminus \mathcal{O}$. A play π *satisfies* an objective \mathcal{O} if $\pi \in \mathcal{O}$, and is a *winning play* of Eve. A winning play π of Adam satisfies $\Pi \setminus \mathcal{O}$.

A *reachability* objective asserts that the play in G has to reach target vertices $T \subseteq V$, formally given by $RE(T) = \{v_0 v_1 \dots \in \Pi \mid \exists k \in \mathbb{N}, v_k \in T\}$. If $T = \{v\}$ for some vertex v , we simply write $RE(v)$.

Let $p : V \rightarrow \mathbb{N}$ be a *priority function* which assigns a priority $p(v)$ to each vertex $v \in V$. For $T \subseteq V$, let $p(T) = \{p(t) \mid t \in T\}$. A parity objective asserts that the minimum priority visited infinitely often along an infinite path is even: $PA(p) = \{\pi = v_0 v_1 \dots \in \Pi \mid \min(p(\inf(\pi))) \text{ is even}\}$.

We formally define stochastic games as follows:

► **Definition 3** (Stochastic Games). Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. A **stochastic game (SG)** with winning objective $\Phi \subseteq \Pi$ is defined as (G, Φ) . If Φ is a reachability or parity objective, (G, Φ) is a stochastic reachability game (SRG) or **stochastic**

*parity game (SPG) respectively. SRGs are also referred to as **simple stochastic games (SSG)**. When the winning objective is clear from the context, we refer to G as a stochastic game.*

Solving stochastic games

Let (G, Φ) be an SG, and let Eve and Adam follow strategies $\sigma \in \Sigma_{\exists}^A$ and $\gamma \in \Sigma_{\forall}^A$. Given a starting vertex $v \in V$, the probability for play π to satisfy Φ – the probability for Eve to win – is denoted $\mathbb{P}_{\sigma, \gamma}^v(\Phi)$. The probability for Adam to win is $\mathbb{P}_{\sigma, \gamma}^v(\Pi \setminus \Phi)$.

Let the *value* of a vertex v be the maximal probability of generating a play from v that satisfies Φ , formally defined using a *value function* $\langle E \rangle(\Phi)(v) = \sup_{\sigma \in \Sigma_{\exists}^A} \inf_{\gamma \in \Sigma_{\forall}^A} \mathbb{P}_{\sigma, \gamma}^v(\Phi)$ for Eve, and $\langle A \rangle(\Pi \setminus \Phi)(v) = \sup_{\gamma \in \Sigma_{\forall}^A} \inf_{\sigma \in \Sigma_{\exists}^A} \mathbb{P}_{\sigma, \gamma}^v(\Pi \setminus \Phi)$ for Adam. A strategy σ for Eve is *optimal* from vertex v if $\inf_{\gamma \in \Sigma_{\forall}^A} \mathbb{P}_{\sigma, \gamma}^v(\Phi) = \langle E \rangle(\Phi)(v)$. Optimal strategies for Adam are defined analogously.

We divide solving stochastic games into three distinct tasks. Given an SG, solving the SG *quantitatively* amounts to computing the values of all vertices in the arena. Solving the SG *strategically* amounts to computing an optimal strategy of Eve (or Adam) for the game.

Since for both SSGs and SPGs, solving quantitatively and strategically is polynomially equivalent [2], we just say “solving” in what follows. We mainly consider quantitative solving, but Theorem 10 applies to both quantitative and strategic solving.

Determinacy

Determinacy refers to the property of an SG where both players, Eve and Adam, have optimal strategies, meaning they can guarantee to achieve the values of the game, regardless of the strategies employed by the other player. *Pure memoryless determinacy* means that both players have pure memoryless optimal strategies.

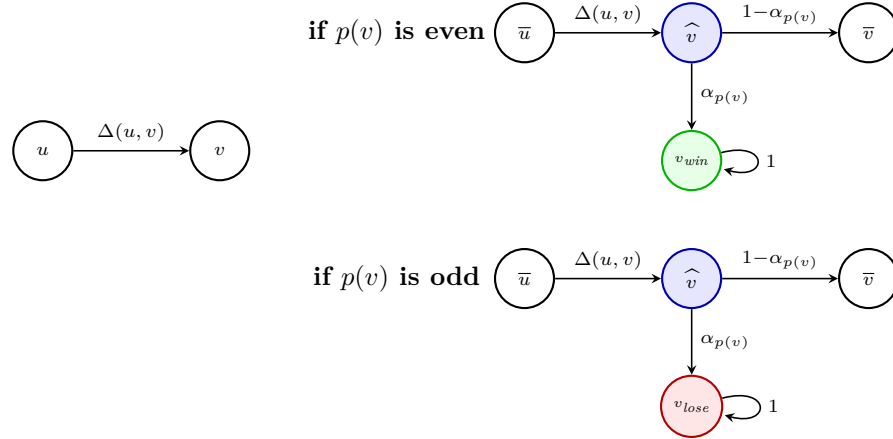
► **Theorem 4** (Pure Memoryless Determinacy [35]). *Let (G, Φ) be an SG, where Φ is a reachability or parity objective. For all $v \in V$, it holds that $\langle E \rangle(\Phi)(v) + \langle A \rangle(\Pi \setminus \Phi)(v) = 1$. Pure memoryless optimal strategies exist for both players from all vertices.*

When Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma_{\exists}$ and $\gamma \in \Sigma_{\forall}$ respectively, we obtain sub-arena $G_{\sigma, \gamma}$, which can be seen as the MC $\mathcal{M}_{\sigma, \gamma}$. We can reduce the winning probabilities $\mathbb{P}_{\sigma, \gamma}^{v_I}$ to reachability probabilities in $G_{\sigma, \gamma}$ as follows. Given a reachability objective $RE(T)$, $\mathbb{P}_{\sigma, \gamma}^{v_I}(RE(T)) = \Pr_{\sigma, \gamma}^{v_I}(Reach(T))$. Given a parity objective $PA(p)$, i.e. $(G, PA(p))$, we call $B \in BSCC(\mathcal{M}_{\sigma, \gamma})$ an *even BSCC* if $\min(p(B))$ is even, meaning intuitively the smallest priority of its vertices is even. *Odd BSCCs* are defined analogously. Then $\mathbb{P}_{\sigma, \gamma}^{v_I}(PA(p)) = \Pr_{\sigma, \gamma}^{v_I}(Reach(B_E))$, where $B_E = \bigcup_{\min(p(B)) \text{ is even}} B \in BSCC(\mathcal{M}_{\sigma, \gamma})$.

► **Corollary 5** (Sufficiency of Pure Memoryless Strategies [11]). *Let $G = ((V, E), (V_{\exists}, V_{\forall}, V_R), \Delta)$ be a stochastic arena, $RE(T)$ a reachability objective, and $PA(p)$ a parity objective. For all vertices $v \in V$, it holds:*

$$\begin{aligned} \text{— } \langle E \rangle(RE(T))(v) &= \sup_{\sigma \in \Sigma_{\exists}^A} \inf_{\gamma \in \Sigma_{\forall}^A} \mathbb{P}_{\sigma, \gamma}^v(RE(T)) = \sup_{\sigma \in \Sigma_{\exists}} \inf_{\gamma \in \Sigma_{\forall}} \Pr_{\sigma, \gamma}^v(Reach(T)) \\ \text{— } \langle E \rangle(PA(p))(v) &= \sup_{\sigma \in \Sigma_{\exists}^A} \inf_{\gamma \in \Sigma_{\forall}^A} \mathbb{P}_{\sigma, \gamma}^v(PA(p)) = \sup_{\sigma \in \Sigma_{\exists}} \inf_{\gamma \in \Sigma_{\forall}} \Pr_{\sigma, \gamma}^v(Reach(B_E)) \\ &\text{where } B_E = \bigcup_{\min(p(B)) \text{ is even}} B \in BSCC(\mathcal{M}_{\sigma, \gamma}). \end{aligned}$$

Therefore we consider only pure memoryless strategies in the sequel, unless stated otherwise.



■ **Figure 4** The gadgets for reducing SPG G (left) to SSG \tilde{G} (right).

3 A Gadget for Transforming SPGs into SSGs

The aim of this paper is to reduce an SPG $(G, PA(p))$ to an SSG $(\tilde{G}, RE(v_{win}))$ such that the probability of reaching target vertex v_{win} in this SSG is related to the probability of winning in the SPG $(G, PA(p))$. As an important step toward this goal, we introduce in this section a gadget that expands each transition of G while removing the priority function.

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, $p : V \rightarrow \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. Section 3.1 presents the gadget enabling the reduction from SPG to SSG $(\tilde{G}, RE(v_{win}))$. We then analyze how probabilistic events in \tilde{G} are related to those in G . Section 3.2 presents a bound on the probability of reaching BSCCs in \tilde{G} . Section 3.3 provides a bound on the winning probability once a BSCC in \tilde{G} is reached, while Section 3.4 gives interval bounds on the winning probabilities in \tilde{G} with regard to those in G .

3.1 Gadget Construction

To reduce the parity objective to a reachability objective, we transform the SPG $(G, PA(p))$ into the SSG $(\tilde{G}, RE(v_{win}))$ by means of a gadget, whose structure was defined by Chatterjee and Fijalkow in [9] to reduce deterministic parity games (DPG) to SSGs.

When both players' strategies are fixed in a DPG, every play forms a lasso, *i.e.* a finite simple path that ends in a simple cycle. Chatterjee and Fijalkow's analysis proceeds by examining lasso plays to choose transition probabilities of the gadget. The specific probabilities they introduce do not extend to a reduction from SPGs to SSGs, because in our stochastic setting, a lasso lifts to a more complicated structure, namely an induced Markov chain: the simple path becomes the transient segment, and each cycle becomes a BSCC. Hence, our approach makes use of smaller probability values, requiring more representation space.

The structure of the gadget remains applicable, and its intuition is as follows: whenever a play visits a vertex with even priority in G , give a small but positive chance to reach a winning sink in \tilde{G} . Vertices with odd priority yield a small chance to reach a losing sink. Finally, to represent that smaller priorities have precedence over larger ones, the probability of reaching a sink from a vertex depends on the priority it is associated to. We introduce a monotonically decreasing function α for this purpose.

We obtain the stochastic arena \tilde{G} by modifying G as indicated in Figure 4. Each vertex v in G is duplicated in \tilde{G} yielding vertices \hat{v} and \bar{v} . A transition $\Delta(u, v)$ in G is replaced by first moving to \hat{v} , which can then either evolve to a sink with probability $\alpha_{p(v)}$, or to the copy \bar{v} with the complementary probability. Depending on p being even or odd, the sink is v_{win} or v_{lose} .

Formally, for $U \subseteq V$, let $\bar{U} = \{\bar{v} \mid v \in U\}$, $\hat{U} = \{\hat{v} \mid v \in U\}$ and $\tilde{U} = \bar{U} \uplus \hat{U}$. We define the arena $\tilde{G} = ((\tilde{V} \uplus \{v_{win}, v_{lose}\}, \tilde{E}), (\bar{V}_\exists, \bar{V}_\forall, \bar{V}_R \uplus \hat{V} \uplus \{v_{win}, v_{lose}\}), \tilde{\Delta})$ where the new edge set \tilde{E} is as follows: $\tilde{E} = \{(\bar{u}, \bar{v}) \mid (u, v) \in E\} \uplus \{(\hat{v}, \bar{v}), (\hat{v}, v_{win}) \mid v \in V, p(v) \text{ is even}\} \uplus \{(\hat{v}, \bar{v}), (\hat{v}, v_{lose}) \mid v \in V, p(v) \text{ is odd}\} \uplus \{(v_{win}, v_{win}), (v_{lose}, v_{lose})\}$.

To define the new transition function $\tilde{\Delta}$, let $\alpha : \mathbb{N} \rightarrow [0, 1]$ where α_i represents the probability of entering the winning (resp. losing) sink before visiting a vertex with even (resp. odd) priority i . We give suitable values for α later, in Lemma 11 on page 14. Now, we define $\tilde{\Delta} : \tilde{V} \times \tilde{V} \rightarrow [0, 1]$ as follows:

$$\tilde{\Delta}(\tilde{u}, \tilde{w}) = \begin{cases} \Delta(u, w) & \text{if } \tilde{u} \in \bar{V}, \tilde{w} \in \hat{V} \\ 1 - \alpha_{p(u)} & \text{if } \tilde{u} \in \hat{V}, \tilde{w} \in \bar{V}, u = w \\ \alpha_{p(u)} & \text{if } \tilde{u} \in \hat{V}, p(u) \text{ is even, } \tilde{w} = v_{win} \\ \alpha_{p(u)} & \text{if } \tilde{u} \in \hat{V}, p(u) \text{ is odd, } \tilde{w} = v_{lose} \\ 1 & \text{if } \tilde{u} = \tilde{w} = v_{win} \text{ or } \tilde{u} = \tilde{w} = v_{lose} \\ 0 & \text{otherwise} \end{cases}$$

When the context is clear, we also address the SPG $(G, PA(p))$ and the SSG $(\tilde{G}, RE(v_{win}))$ with G and \tilde{G} respectively.

Since all new vertices $\hat{V} \uplus \{v_{win}, v_{lose}\}$ are random vertices, a strategy of either player in SPG G is a strategy in SSG \tilde{G} and vice versa. That is, there is a one-to-one relationship between strategies in G and \tilde{G} . Hence, to keep notations simpler, we do not distinguish between strategies in G and \tilde{G} .

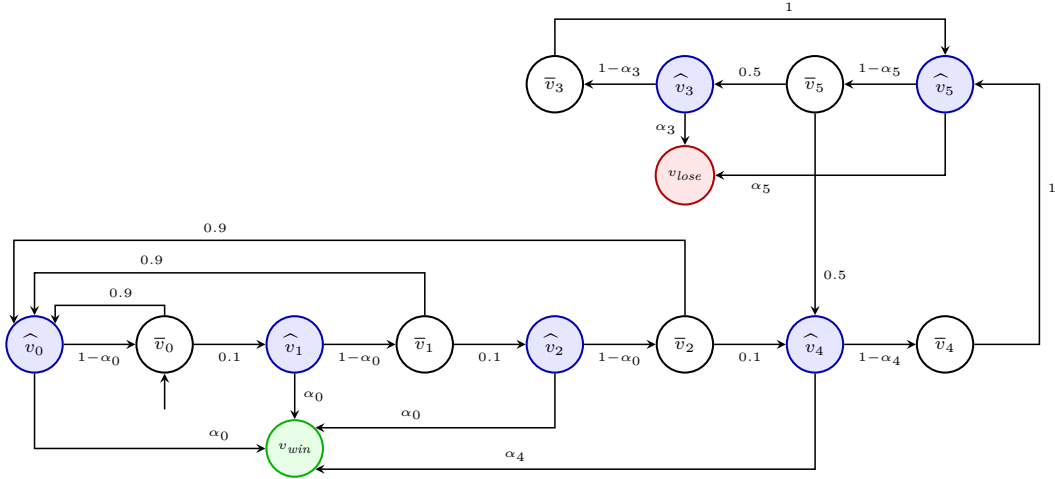
A pair of strategies $\sigma, \gamma \in \Sigma_\exists \times \Sigma_\forall$ for Eve and Adam in G induces the sub-arena $G_{\sigma, \gamma}$. Similarly, we obtain $\tilde{G}_{\sigma, \gamma}$. If U is an even or odd BSCC in SPG $G_{\sigma, \gamma}$, we denote with \tilde{U} what we call the associated *even pBSCC* or *odd pBSCC* in SSG $\tilde{G}_{\sigma, \gamma}$ respectively. While those are not BSCCs, they correspond to the BSCC of the associated parity game, and we never consider the only true BSCCs of $\tilde{G}_{\sigma, \gamma}$, i.e. $\{v_{win}\}$ and $\{v_{lose}\}$.

We continue with the example in Figure 2 and Figure 3. For the vertices v_0, v_1 , and v_2 , we assign priority 0, and for each remaining vertex $v_i, i \in \{3, 4, 5\}$, we assign priority i . An illustration of the corresponding sub-arena $\tilde{G}_{\sigma, \gamma}$, induced by our gadget construction, is provided in Figure 5. Since we do not need to distinguish between different types of vertices, we use circles for all vertices. In $G_{\sigma, \gamma}$, the set $\{v_3, v_4, v_5\}$ forms an odd BSCC. We refer to the set $\{\hat{v}_3, \bar{v}_3, \hat{v}_4, \bar{v}_4, \hat{v}_5, \bar{v}_5\}$ in $\tilde{G}_{\sigma, \gamma}$ as the associated odd pBSCC with \hat{v}_3, \hat{v}_4 , and \hat{v}_5 having outgoing transitions to either v_{win} or v_{lose} .

3.2 Before Entering a pBSCC in SSG \tilde{G}

Recall that when Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma_\exists$ and $\gamma \in \Sigma_\forall$, the resulting sub-arenas $G_{\sigma, \gamma}$ and $\tilde{G}_{\sigma, \gamma}$ can be viewed as finite Markov chains. We first focus on what happens before a play reaches a pBSCC in $\tilde{G}_{\sigma, \gamma}$. Specifically, we give a lower bound on the probability of reaching an entry state of a pBSCC without entering a winning or losing sink. Later, in Lemma 11, we use this bound to determine a suitable value for α_0 .

Intuitively, we show that the probability of entry is minimized in a classical worst-case scenario extending the sub-arena $\{v_0, v_1, v_2\}$ in Figure 3, and $\{\hat{v}_0, \bar{v}_0, \hat{v}_1, \bar{v}_1, \hat{v}_2, \bar{v}_2\}$ in Figure 5. More precisely, we consider an original sub-arena $K_{\sigma, \gamma}$, depicted in Figure 6a with n states



■ **Figure 5** The sub-arena $\tilde{G}_{\sigma, \gamma}$ induced by the gadget.

arranged in a sequence (as $\{v_0, v_1, v_2\}$ in Figure 3) before reaching a BSCC. Each state has the minimal probability δ_{min} of progressing to the next state and a maximal probability $1 - \delta_{min}$ of returning to the initial state. All these states have parity value 0. Upon applying our gadget construction to obtain $\tilde{K}_{\sigma, \gamma}$ in Figure 6b, we introduce random states ($\{\hat{v}_0, \hat{v}_1, \hat{v}_2\}$ in Figure 5) that have the highest probability α_0 of going to the winning sink v_{win} .

In the following, let $(G, PA(p))$ be an SPG, and $(\tilde{G}, RE(v_{win}))$ be its associated SSG. For all strategy pairs $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, let $\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(crossPath)$ denote the probability for a play starting from $\bar{v} \in \bar{V}$ to reach a pBSCC in \tilde{G} . Note that we never consider $\hat{v} \in \hat{V}$ as the starting vertex.

► **Lemma 6.** *For all strategy pairs $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, for all $v \in V$, it holds:*

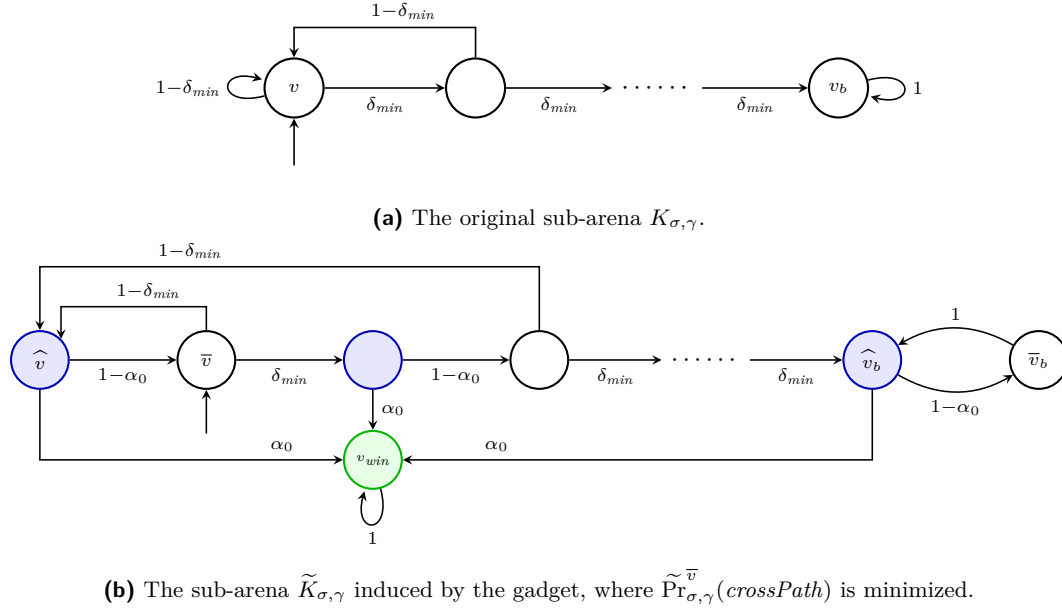
$$\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(crossPath) \geq \frac{(1 - x_0)x_0^n}{(1 - x_0) - (1 - x_0^n)x_1}$$

where $n = |V|$, $x_0 = \delta_{min}(1 - \alpha_0)$, $x_1 = (1 - \delta_{min})(1 - \alpha_0)$, and $\delta_{min} = \min_{u, v \in V} \{\Delta(u, v) \mid \Delta(u, v) > 0\}$.

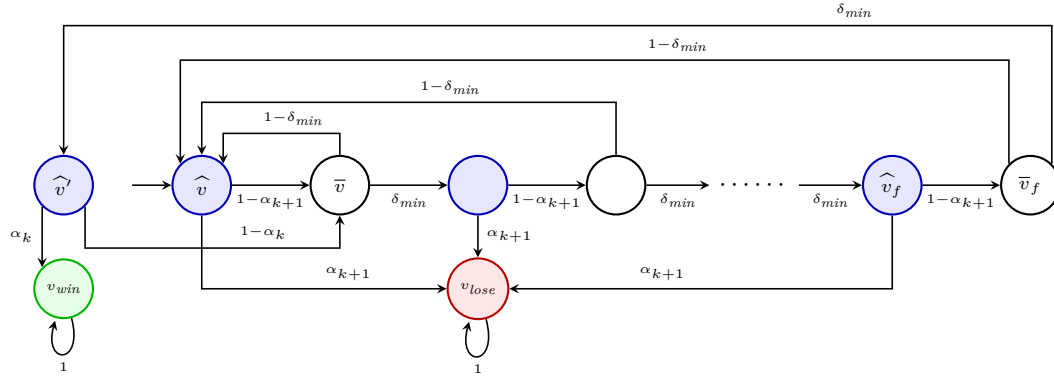
Sketch of Proof. We fix an arbitrary strategy pair $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, and analyze the corresponding MC $\tilde{G}_{\sigma, \gamma}$. We simplify the MC while either preserving or under-approximating the probability of reaching a pBSCC in $\tilde{G}_{\sigma, \gamma}$. These steps merge all pBSCCs into a sink v_b , eliminate auxiliary states to simplify the MC, increase all values of α , and restructure transitions so that only one designated vertex can reach the sink v_b directly. We denote the resulting MC with \tilde{G}_4 . We then derive a lower bound on the probability of reaching v_b in \tilde{G}_4 , which provides a reachability lower bound in a template MC with absorbing sinks and bounded transition probabilities. As the reachability probabilities of v_b in \tilde{G}_4 underapproximate those in $\tilde{G}_{\sigma, \gamma}$, this yields the desired lower bound on $\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(crossPath)$.

3.3 Inside a pBSCC in SSG \tilde{G}

We now focus on what happens after a play reaches a pBSCC in sub-arena $\tilde{G}_{\sigma, \gamma}$. Specifically, we consider MCs and give a lower bound on the probability of reaching the winning sink after reaching an even pBSCC, and dually an upper bound on the probability of reaching the winning sink after reaching an odd pBSCC.



■ **Figure 6** Example sub-arenas where the probability of *crossPath* is minimized.



■ **Figure 7** The MC where the lower bound on minimum *winEven* probability is attained.

The lower bound is attained in the MC shown in Figure 7, where k is an even parity value. There are $2n + 1$ states in a line, and winning and losing sinks. Each white state has maximal probability $1 - \delta_{\min}$ to return to the initial state, and otherwise proceeds to the next blue state. Each blue state, except \hat{v}' , can with probability α_{k+1} go to the losing sink, and otherwise proceeds to the next white state. The special state \hat{v}' goes with probability α_k to v_{win} , and otherwise proceeds to \bar{v} . Unlike the case with $\{\hat{v}_0, \bar{v}_0, \hat{v}_1, \bar{v}_1, \hat{v}_2, \bar{v}_2\}$ in Figure 5, this MC cannot be obtained by applying our gadget on some sub-arena $G_{\sigma, \gamma}$, and hence this bound is not guaranteed to be tight. More precisely, the outgoing transitions of \hat{v} indicate that v has an odd parity value, while \hat{v}' suggests otherwise. The upper bound is obtained by considering the same MC, where k is an odd parity value. Later, in Lemma 11, we use these two bounds to find suitable values for all α_k with $k \in \mathbb{N}$.

Let U be an even BSCC with smallest priority k in $G_{\sigma, \gamma}$ and \tilde{U} its associated pBSCC in $\tilde{G}_{\sigma, \gamma}$. Let $\tilde{\Pr}_{\sigma, \gamma}^{k, \tilde{U}}(\text{winEven})$ denote the minimum probability of reaching the winning sink after reaching \tilde{U} . That is, $\tilde{\Pr}_{\sigma, \gamma}^{k, \tilde{U}}(\text{winEven}) = \min\{\tilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{Reach}(v_{\text{win}})) \mid \bar{v} \in \tilde{U}\}$. We denote

it $\widetilde{\Pr}_{\sigma,\gamma}^{k,\widetilde{U}}(\text{winEven})$ when \widetilde{U} is clear from context. Analogously, given an odd BSCC U with smallest priority k in $G_{\sigma,\gamma}$, we use $\widetilde{\Pr}_{\sigma,\gamma}^k(\text{winOdd})$ to denote the maximum probability of reaching the winning sink after reaching \widetilde{U} .

► **Lemma 7.** *For all strategy pairs $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, for all even k , it holds:*

$$\widetilde{\Pr}_{\sigma,\gamma}^k(\text{winEven}) \geq (1 - \alpha_{k+1}) \cdot \frac{(1 - x_2) \cdot x_2^{n-1} \cdot x_4}{1 - (x_2 + x_3) + x_5 \cdot x_2^n + t \cdot x_2^{n-1} - x_5 \cdot x_2^{n-1}}$$

and for all odd k , it holds:

$$\widetilde{\Pr}_{\sigma,\gamma}^k(\text{winOdd}) \leq 1 - (1 - \alpha_{k+1}) \cdot \frac{(1 - x_2) \cdot x_2^{n-1} \cdot x_4}{1 - (x_2 + x_3) + x_5 \cdot x_2^n + t \cdot x_2^{n-1} - x_5 \cdot x_2^{n-1}}$$

where $n = |V|$, $x_2 = \delta_{\min}(1 - \alpha_{k+1})$, $x_3 = (1 - \delta_{\min})(1 - \alpha_{k+1})$, $x_4 = \delta_{\min}\alpha_k$, $x_5 = \delta_{\min}(1 - \alpha_k) + (1 - \delta_{\min})(1 - \alpha_{k+1})$, and δ_{\min} is as before.

Sketch of Proof. The proof follows the same structure as the one for Lemma 6, applied to a BSCC in $\mathcal{M}_{\sigma,\gamma}$. We apply a similar four-step transformation and obtain a simplified MC where we can directly derive an upper bound. The lower bound comes as the dual.

3.4 Range of Winning Probabilities in the SSG

We now relate the winning probabilities in the constructed SSG \widetilde{G} to the original SPG G . Intuitively, with a fixed strategy pair $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, the value $\widetilde{\mathbb{P}}_{\sigma,\gamma}$ of the SSG \widetilde{G} falls into a range around the value $\mathbb{P}_{\sigma,\gamma}$ of the SPG G , and the range size depends on the probabilities crossPath , winEven and winOdd .

► **Lemma 8.** *Let $x, y \in (0, 1)$ such that $\widetilde{\Pr}_{\sigma,\gamma}^{\overline{v}}(\text{crossPath}) > x$ for all even k , and $\widetilde{\Pr}_{\sigma,\gamma}^k(\text{winEven}) \geq y$, and for all odd k , $\widetilde{\Pr}_{\sigma,\gamma}^k(\text{winOdd}) \leq 1 - y$, then it holds:*

$$y \cdot \mathbb{P}_{\sigma,\gamma}^v - y + x \cdot y \leq \widetilde{\mathbb{P}}_{\sigma,\gamma}^{\overline{v}} \leq \mathbb{P}_{\sigma,\gamma}^v + 1 - x \cdot y$$

4 Reducing SPGs to SSGs

We now present the direct reduction from SPGs to SSGs. Let $G = ((V, E), (V_{\exists}, V_{\forall}, V_R), \Delta)$ be a stochastic arena, $p : V \rightarrow \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct the SSG $(\widetilde{G}, RE(v_{\text{win}}))$ using the gadget presented in Section 3.1. Section 4.1 presents a lower bound on the difference between winning probabilities associated to different strategy pairs in G . Section 4.2 presents the main theorem establishing the reduction, while Section 4.3 gives complexity bounds.

4.1 A Lower Bound on Different Strategies

We consider two strategy pairs $(\sigma, \gamma), (\sigma', \gamma') \in \Sigma_{\exists} \times \Sigma_{\forall}$ and show a general result on all such pair: if they yield different values in G , then there exists a lower bound on the difference between these values.

In the following, we assume for all $u \in V_R, v \in V$ that $\Delta(u, v)$ is a rational number $\frac{a_{u,v}}{b_{u,v}}$, where $a_{u,v} \in \mathbb{N}, b_{u,v} \in \mathbb{N}_+$ and $a_{u,v} \leq b_{u,v}$. Let $M = \max_{(u,v) \in E} \{b_{u,v}\}$ and $n = |V|$.

► **Lemma 9.** *For all $(\sigma, \gamma), (\sigma', \gamma') \in \Sigma_{\exists} \times \Sigma_{\forall}$, for all $v \in V$, the following holds:*

$$\mathbb{P}_{\sigma, \gamma}^v \neq \mathbb{P}_{\sigma', \gamma'}^v \Rightarrow |\mathbb{P}_{\sigma, \gamma}^v - \mathbb{P}_{\sigma', \gamma'}^v| > \frac{1}{(n!)^2 M^{2n^2}} = \epsilon$$

Proof. Let $\Pr_{\sigma, \gamma}^v(\text{enterEven})$ be the probability for a play starting from $v \in V$ to reach an even BSCC. It follows from Corollary 5 that for all $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$ and $v \in V$, we have $\mathbb{P}_{\sigma, \gamma}^v = \Pr_{\sigma, \gamma}^v(\text{enterEven})$. We can obtain $\Pr_{\sigma, \gamma}^v(\text{enterEven})$ by setting all vertices belonging to at least one even BSCC as the target set, and calculating the reachability probability. Calculating $\mathbb{P}_{\sigma, \gamma}^v$ is thus reduced to solving a linear equation system $x = Ax + b$ according to Theorem 1. We omit the details of A and b . Every non-zero entry of A and b is either 1, or $\frac{a_{u,v}}{b_{u,v}}$ for some $u, v \in V_R \times V$.

We use the following notations:

- Let $s = |b|$. It follows that $s < n$ since there is at least one vertex in a BSCC.
- Let $Q = I - A$. For $i \in 1, 2, \dots, n$ we denote the i -th row of Q with $Q[i]$, and the entry of Q at i -th row and j -th column with $Q[i, j]$. It can be written as $Q[i, j] = \frac{c_{i,j}}{d_{i,j}}$, where $|c_{i,j}|$ and $|d_{i,j}|$ are natural numbers bounded by M with $|c_{i,j}| \leq |d_{i,j}|$.
- We denote the i -th entry of b with $b[i]$. It can be written as $b[i] = \frac{c_{i,s+1}}{d_{i,s+1}}$, where $|c_{i,s+1}|$ and $|d_{i,s+1}|$ are natural numbers bounded by M with $|c_{i,s+1}| \leq |d_{i,s+1}|$.

The equation system can be written as:

$$Qx = b$$

We take an arbitrary row i , and write the i -th equation $Q[i] \cdot x = b[i]$ as follows:

$$\begin{bmatrix} \frac{c_{i,1}}{d_{i,1}} & \frac{c_{i,2}}{d_{i,2}} & \dots & \frac{c_{i,s}}{d_{i,s}} \end{bmatrix} \cdot x = \frac{c_{i,s+1}}{d_{i,s+1}} \quad (1)$$

We multiply equation (1) with $\prod_{t=1}^{s+1} d_{i,t}$ to obtain:

- For all $j = 1, \dots, s$, $Q[i, j]$ equals $(\prod_{t=1}^{s+1} d_{i,t}) \frac{c_{i,j}}{d_{i,j}}$, an integer with absolute value bounded by M^{s+1} .
- For all $i = 1, \dots, s$, $b[i]$ equals $(\prod_{t=1}^s d_{i,t}) c_{i,s+1}$, an integer with absolute value bounded by M^{s+1} .

We apply this transformation to each row of the equation system, and write the new equation system as:

$$Q'x = b'$$

By Cramer's rule, for all $i = 1, 2, \dots, s$, we obtain:

$$x[i] = \frac{\det(Q'_i)}{\det(Q')}$$

where Q'_i is the matrix obtained by replacing the i -th column of Q' with the column vector b' . It follows that all entries of Q'_i are also integers with absolute values bounded by M^{s+1} .

Since $x[i]$ is a reachability probability, we have $x[i] \leq 1$. Following from the calculation of determinants, we obtain the following:

$$|\det(Q'_i)| \leq |\det(Q')| \leq s!(M^{s+1})^s < n!M^{n^2}$$

Therefore, if the equation system resulting from σ', γ' yields $x'[i] > x[i]$, we have:

$$x'[i] - x[i] > \frac{1}{(n!M^{n^2})^2} = \frac{1}{(n!)^2 M^{2n^2}} \quad \blacktriangleleft$$

4.2 Direct Reduction

We now establish the direct reduction from SPGs to SSGs.

► **Theorem 10** (Reducing SPGs to SSGs). *If for all $(\sigma, \gamma) \in \Sigma_{\exists} \times \Sigma_{\forall}$, and $v \in V$, the following conditions hold:*

1. $\widetilde{\Pr}_{\sigma, \gamma}^v(\text{crossPath}) > \frac{4-\epsilon}{4}$
 2. $\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven}) \geq \frac{4}{4+\epsilon}$ for all even k , and $\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winOdd}) \leq 1 - \frac{4}{4+\epsilon}$ for all odd k
- where $\epsilon = \frac{1}{(n!)^2 M^{2n^2}}$, then every optimal strategy $\sigma \in \Sigma_{\exists}$ of Eve in the SSG $(\tilde{G}, RE(v_{\text{win}}))$ is also optimal in the SPG $(G, PA(p))$. The same holds for Adam.

Proof. We assume conditions 1 and 2 hold. We show that every optimal strategy $\sigma \in \Sigma_{\exists}$ in SSG \tilde{G} is also optimal in SPG G . We prove this by contraposition.

We take $\Sigma_{\exists}^* \subseteq \Sigma_{\exists}$ and $\Sigma_{\forall}^* \subseteq \Sigma_{\forall}$ as the sets of optimal strategies of Eve and Adam in \tilde{G} . We obtain by Lemma 8 that for all $v \in V$ and all $\sigma, \gamma \in \Sigma_{\exists}^* \times \Sigma_{\forall}^*$, the following holds:

$$y \cdot \mathbb{P}_{\sigma, \gamma}^v - y + x \cdot y \leq \widetilde{\mathbb{P}}_{\sigma, \gamma}^v \leq \mathbb{P}_{\sigma, \gamma}^v + 1 - x \cdot y$$

Since conditions 1 and 2 hold, we substitute x and y to obtain:

$$\widetilde{\mathbb{P}}_{\sigma, \gamma}^v \leq \mathbb{P}_{\sigma, \gamma}^v + \frac{2\epsilon}{4+\epsilon} \quad (2)$$

If σ is not optimal in G , then there exists another strategy $\sigma' \in \Sigma_{\exists}$ and a vertex $v \in V$ such that $\mathbb{P}_{\sigma', \gamma}^v > \mathbb{P}_{\sigma, \gamma}^v$. It follows again from Lemma 8 that:

$$y \mathbb{P}_{\sigma', \gamma}^v - y + x \cdot y \leq \widetilde{\mathbb{P}}_{\sigma', \gamma}^v \leq \mathbb{P}_{\sigma', \gamma}^v + 1 - x \cdot y \quad (3)$$

Furthermore, Lemma 9 yields:

$$\mathbb{P}_{\sigma', \gamma}^v > \mathbb{P}_{\sigma, \gamma}^v + \epsilon \quad (4)$$

As a result, we obtain the following:

$$\begin{aligned} \widetilde{\mathbb{P}}_{\sigma', \gamma}^v &\geq y \cdot \mathbb{P}_{\sigma', \gamma}^v - y + x \cdot y && \text{by (3)} \\ &> y \cdot (\mathbb{P}_{\sigma, \gamma}^v + \epsilon) - y + x \cdot y && \text{by (4)} \\ &= \frac{4}{4+\epsilon} (\mathbb{P}_{\sigma, \gamma}^v + \epsilon) - \frac{4}{4+\epsilon} \cdot \frac{\epsilon}{4} \\ &= \mathbb{P}_{\sigma, \gamma}^v - \frac{\epsilon}{4+\epsilon} \mathbb{P}_{\sigma, \gamma}^v + \frac{3\epsilon}{4+\epsilon} \\ &\geq \mathbb{P}_{\sigma, \gamma}^v + \frac{2\epsilon}{4+\epsilon} \\ &\geq \widetilde{\mathbb{P}}_{\sigma, \gamma}^v && \text{by (2)} \end{aligned}$$

It indicates that $\widetilde{\mathbb{P}}_{\sigma', \gamma}^v > \widetilde{\mathbb{P}}_{\sigma, \gamma}^v$, which contradicts the assumption that $\sigma \in \Sigma_{\exists}^*$. ◀

Until now, we have used the function α in Theorem 10, obtaining inequalities relating parity values in SPG G to transition probabilities in SSG \tilde{G} . We now give requirements for α that satisfy all these inequalities.

► **Lemma 11.** *When the values of α are arranged as follows, the conditions in Theorem 10 are satisfied:*

1. If $\alpha_0 \leq \frac{\delta_{\min}^n}{8(n!)^2 M^{2n^2}}$, then condition 1 is satisfied.
2. If for all $k \in \mathbb{N}$, the following holds, then condition 2 is satisfied:

$$\frac{\alpha_{k+1}}{\alpha_k} \leq \frac{\delta_{\min}^n (1 - \delta_{\min})}{8(n!)^2 M^{2n^2} + 1}.$$

Sketch of Proof. Both cases follow a similar structure. For α_0 (respectively α_{k+1}/α_k), we derive from the bound given by Lemma 6 (resp. Lemma 7) a corollary giving a bound that explicitly makes use of α . We then directly obtain the two cases of Lemma 11 from these two bounds. The full proof of this lemma, detailing how to compute function α can be found in Appendix A.1.

4.3 Complexity Considerations

To introduce complexity results, we first define size of a stochastic game G as $|G| = |V| + |E| + |\Delta|$ where $|\Delta|$ is the space needed to store the transition function (which may be stored in unary or binary). A now longstanding result shows that most stochastic game settings are polynomially reducible one to the other. In particular:

► **Theorem 12** (From Theorem 1 in [2]). *Solving stochastic parity games and solving simple stochastic games is polynomial-time equivalent. Either can be using unary or binary encoding.*

We show that the reduction we have introduced in this paper is polynomial with binary encoding. We recall that $M = \max_{(u,v) \in E} \{b_{u,v}\}$ and $n = |V|$.

► **Theorem 13.** *Given an SPG G , there exist polynomial values for function α that satisfy Theorem 10, such that the SSG \tilde{G} is of size $\mathcal{O}(n^5 \log M)$ in binary.*

Proof. Since $\delta_{\min} \geq \frac{1}{M}$ and $1 - \delta_{\min} \geq \frac{1}{2}$, the following is a valid instance of α , polynomial in G (and polynomial in the transition probabilities appearing in G) under binary encoding:

$$\forall k \in \mathbb{N}, \alpha_k = \left(\frac{1}{16(n!)^2 M^{2n^2+n} + 1} \right)^{k+1}$$

Then the size of the SSG \tilde{G} is:

$$\begin{aligned} |\tilde{G}| &= |\tilde{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}| + |\tilde{E}| + |\tilde{\Delta}| \\ &= \mathcal{O}(n) + \mathcal{O}(n^2) + \mathcal{O}(n^2) \cdot \mathcal{O}(n \cdot (n \log n + n^2 \log M)) \\ &= \mathcal{O}(n^5 \log M) \end{aligned}$$

◀

According to [2], quantitative SSGs under unary and binary encoding are in the same complexity class, and so in $\mathbf{NP} \cap \mathbf{coNP}$ [18]. We thus obtain that our reduction yields an $\mathbf{NP} \cap \mathbf{coNP}$ algorithm for solving SPGs.

5 Epilogue

We have given a polynomial reduction from quantitative SPGs to quantitative SSGs, taking inspiration from a gadget used in [9] to obtain a reduction from deterministic PGs to quantitative SSGs. After fixing a pair of strategies, the values of both the SPG and the SSG are determined, but the construction of the SSG makes it difficult to establish coinciding values. Using these fixed strategies, we showed that the value of the SSG falls into a range around the value of the SPG, where this range depends on the probability to reach a pBSCC of the SSG and the minimum probability to reach a winning sink in pBSCCs of the SSG. When considering all possible strategy pairs, we obtained a lower bound ϵ on their value differences in the SPG, by restricting transition probabilities to rational numbers and analyzing reachability equation systems of Markov chains. We then showed that by arranging transition probabilities of the SSG properly in terms of the size of the SPG, its smallest probability, and ϵ , the value ranges of different strategy pairs can be narrowed so that they do not overlap. In this case, a reduction from SPGs to SSGs is achieved.

Although under unary encoding, exponential numbers can be introduced into the probability function α of the newly constructed SSGs, both reductions are polynomial. Hence, our construction yields an $\mathbf{NP} \cap \mathbf{coNP}$ algorithm in both qualitative and quantitative SSGs under unary and binary encoding, substantiating the complexity results from previous works [18, 2, 9].

Our result enables solving SPGs by first reducing them to SSGs and then applying algorithms for SSGs. However, its implementability is in question, due to the possibly huge representation of α . Our reduction also captures the transformation from an MDP with a parity objective into an SSG. As we assume the minimum transition probability to be $\delta_{\min} \in (0, \frac{1}{2}]$ in the original SPG, we cannot capture the subcase of reducing DPGs to quantitative SSGs. Although our reduction is unlikely to be leveraged to effectively solve SPGs in practice, some improvements are possible. First, we have not formally examined the optimal arrangement of α . It is possible to find the weakest requirements on α so that the reductions are correct, thus optimizing possible implementations. Second, some specific cases lead to very small values of α . These cases are similar to the ones that can challenge classical MDP solvers using VI, and so we can benefit from any family of arena structure where these cases are avoided, leading to implementable valuations of α .

References

- 1 Rabah Amir. Stochastic games in economics and related fields: an overview. *Stochastic Games and Applications*, pages 455–470, 2003.
- 2 Daniel Andersson and Peter Bro Miltersen. The complexity of solving stochastic games on graphs. In *International Symposium on Algorithms and Computation*, pages 112–121. Springer, 2009. doi:10.1007/978-3-642-10631-6_13.
- 3 Muqsit Azeem, Alexandros Evangelidis, Jan Křetínský, Alexander Slivinskiy, and Maximilian Weininger. Optimistic and topological value iteration for simple stochastic games. In *International Symposium on Automated Technology for Verification and Analysis*, pages 285–302. Springer, 2022.
- 4 Christel Baier and Joost-Pieter Katoen. *Principles of Model Checking*. MIT press, 2008.
- 5 Christel Baier, Joachim Klein, Linda Leuschner, David Parker, and Sascha Wunderlich. Ensuring the reliability of your model checker: Interval iteration for Markov decision processes. In *International Conference on Computer Aided Verification*, volume 10426, pages 160–180. Springer, 2017. doi:10.1007/978-3-319-63387-9_8.
- 6 Raphaël Berthon, Joost-Pieter Katoen, and Zihan Zhou. A direct reduction from stochastic parity games to simple stochastic games, 2025. arXiv:2506.06223.
- 7 Patricia Bouyer, Youssef Oualhadj, Mickael Randour, and Pierre Vandenhover. Arena-independent finite-memory determinacy in stochastic games. *Log. Methods Comput. Sci.*, 19(4), 2023. doi:10.46298/LMCS-19(4:18)2023.
- 8 Cristian S Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 252–263, 2017. doi:10.1145/3055399.3055409.
- 9 Krishnendu Chatterjee and Nathanaël Fijalkow. A reduction from parity games to simple stochastic games. *Electronic Proceedings in Theoretical Computer Science*, 54:74–86, 2011. doi:10.4204/eptcs.54.6.
- 10 Krishnendu Chatterjee and Thomas A Henzinger. Strategy improvement and randomized subexponential algorithms for stochastic parity games. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 512–523. Springer, 2006. doi:10.1007/11672142_42.
- 11 Krishnendu Chatterjee and Thomas A Henzinger. Reduction of stochastic parity to stochastic mean-payoff games. *Information Processing Letters*, 106(1):1–7, 2008. doi:10.1016/J.IPL.2007.08.035.

- 12 Krishnendu Chatterjee, Thomas A Henzinger, Barbara Jobstmann, and Arjun Radhakrishna. Gist: A solver for probabilistic games. In *Computer Aided Verification: 22nd International Conference*, volume 6174, pages 665–669. Springer, 2010. doi:10.1007/978-3-642-14295-6_57.
- 13 Krishnendu Chatterjee, Marcin Jurdziński, and Thomas A. Henzinger. Simple stochastic parity games. In Matthias Baaz and Johann A. Makowsky, editors, *Computer Science Logic*, volume 2803, pages 100–113, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. doi:10.1007/978-3-540-45220-1_11.
- 14 Krishnendu Chatterjee, Marcin Jurdzinski, and Thomas A Henzinger. Quantitative stochastic parity games. In *SODA*, volume 4, pages 121–130, 2004. URL: <http://dl.acm.org/citation.cfm?id=982792.982808>.
- 15 Krishnendu Chatterjee, Rupak Majumdar, and Marcin Jurdzinski. On Nash equilibria in stochastic games. In *Computer Science Logic*, volume 3210 of *Lecture Notes in Computer Science*, pages 26–40. Springer, 2004. doi:10.1007/978-3-540-30124-0_6.
- 16 Krishnendu Chatterjee and Nir Piterman. Combinations of qualitative winning for stochastic parity games. In *30th International Conference on Concurrency Theory*, volume 140 of *LIPICs*, pages 6:1–6:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.CONCUR.2019.6.
- 17 Anne Condon. On algorithms for simple stochastic games. *Advances in Computational Complexity Theory*, 13:51–72, 1990. doi:10.1090/DIMACS/013/04.
- 18 Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992. doi:10.1016/0890-5401(92)90048-K.
- 19 Wojciech Czerwiński, Laure Daviaud, Nathanaël Fijalkow, Marcin Jurdziński, Ranko Lazić, and Paweł Parys. Universal trees grow inside separating automata: Quasi-polynomial lower bounds for parity games. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2333–2349. SIAM, 2019.
- 20 Matthew Darlington, Kevin D. Glazebrook, David S. Leslie, Rob Shone, and Roberto Szechtman. A stochastic game framework for patrolling a border. *Eur. J. Oper. Res.*, 311(3):1146–1158, 2023. doi:10.1016/J.EJOR.2023.06.011.
- 21 Julia Eisentraut, Edon Kelmendi, Jan Křetínský, and Maximilian Weininger. Value iteration for simple stochastic games: Stopping criterion and learning algorithm. *Information and Computation*, 285:104886, 2022. doi:10.1016/J.IC.2022.104886.
- 22 Kousha Etessami, Marta Z. Kwiatkowska, Moshe Y. Vardi, and Mihalis Yannakakis. Multi-objective model checking of Markov decision processes. *Log. Methods Comput. Sci.*, 4(4), 2008. doi:10.2168/LMCS-4(4:8)2008.
- 23 Ernst Moritz Hahn, Mateo Perez, Sven Schewe, Fabio Somenzi, Ashutosh Trivedi, and Dominik Wojtczak. Model-free reinforcement learning for stochastic parity games. In *31st International Conference on Concurrency Theory*, volume 171 of *LIPICs*, pages 21:1–21:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.CONCUR.2020.21.
- 24 Arnd Hartmanns and Benjamin Lucien Kaminski. Optimistic value iteration. In *Computer Aided Verification - 32nd International Conference*, volume 12225 of *Lecture Notes in Computer Science*, pages 488–511. Springer, 2020. doi:10.1007/978-3-030-53291-8_26.
- 25 Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, Tim Quatmann, and Matthias Volk. The probabilistic model checker storm. *International Journal on Software Tools for Technology Transfer*, pages 1–22, 2021.
- 26 Marcin Jurdziński. Deciding the winner in parity games is in $\mathbf{UP} \cap \mathbf{co-UP}$. *Information Processing Letters*, 68(3):119–124, 1998.
- 27 Marcin Jurdziński and Ranko Lazić. Succinct progress measures for solving parity games. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–9. IEEE, 2017.

- 28 Mark Kattenbelt, Marta Kwiatkowska, Gethin Norman, and David Parker. A game-based abstraction-refinement framework for Markov decision processes. *Formal Methods in System Design*, 36:246–280, 2010. doi:10.1007/S10703-010-0097-6.
- 29 Jan Křetínský, Emanuel Ramneantu, Alexander Slivinskiy, and Maximilian Weininger. Comparison of algorithms for simple stochastic games. *Information and Computation*, 289:104885, 2022. doi:10.1016/J.IC.2022.104885.
- 30 Marta Kwiatkowska, Gethin Norman, David Parker, and Gabriel Santos. PRISM-games 3.0: Stochastic game verification with concurrency, equilibria and time. In *Computer Aided Verification - 32nd International Conference*, volume 12225 of *Lecture Notes in Computer Science*, pages 475–487. Springer, 2020. doi:10.1007/978-3-030-53291-8_25.
- 31 Marta Z. Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: Verification of probabilistic real-time systems. In *Computer Aided Verification - 23rd International Conference*, volume 6806 of *Lecture Notes in Computer Science*, pages 585–591. Springer, 2011. doi:10.1007/978-3-642-22110-1_47.
- 32 Karoliina Lehtinen. A modal μ perspective on solving parity games in quasi-polynomial time. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 639–648, 2018. doi:10.1145/3209108.3209115.
- 33 David S Leslie, Steven Perkins, and Zibo Xu. Best-response dynamics in zero-sum stochastic games. *Journal of Economic Theory*, 189:105095, 2020. doi:10.1016/J.JET.2020.105095.
- 34 Rupak Majumdar, Kaushik Mallik, Anne-Kathrin Schmuck, and Sadegh Soudjani. Symbolic control for stochastic systems via finite parity games. *Nonlinear Analysis: Hybrid Systems*, 51:101430, 2024. doi:10.1016/j.nahs.2023.101430.
- 35 Donald A Martin. The determinacy of Blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998. doi:10.2307/2586667.
- 36 Pawel Parys. Parity games: Zielonka’s algorithm in quasi-polynomial time. In *44th International Symposium on Mathematical Foundations of Computer Science*, volume 138 of *LIPICs*, pages 10:1–10:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPICs.MFCS.2019.10.
- 37 Kittiphon Phalakarn, Toru Takisaka, Thomas Haas, and Ichiro Hasuo. Widest paths and global propagation in bounded value iteration for stochastic games. In *Computer Aided Verification: 32nd International Conference*, pages 349–371. Springer, 2020. doi:10.1007/978-3-030-53291-8_19.
- 38 Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- 39 Tim Quatmann and Joost-Pieter Katoen. Sound value iteration. In *Computer Aided Verification - 30th International Conference*, volume 10981 of *Lecture Notes in Computer Science*, pages 643–661. Springer, 2018. doi:10.1007/978-3-319-96145-3_37.
- 40 Lloyd S Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953.
- 41 Frédéric Simard, Josée Desharnais, and François Laviolette. General cops and robbers games with randomness. *Theor. Comput. Sci.*, 887:30–50, 2021. doi:10.1016/J.TCS.2021.06.043.
- 42 Michael Ummels and Dominik Wojtczak. The complexity of Nash equilibria in stochastic multiplayer games. *Log. Methods Comput. Sci.*, 7(3), 2011. doi:10.2168/LMCS-7(3:20)2011.
- 43 Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996. doi:10.1016/0304-3975(95)00188-3.

A **Appendix**

We present how we compute values for function α . We include the details, hoping these techniques may later be modified to yield better values on specific classes of games.

A.1 Proof of Lemma 11

► **Lemma 11.** *When the values of α are arranged as follows, the conditions in Theorem 10 are satisfied:*

1. *If $\alpha_0 \leq \frac{\delta_{\min}^n}{8(n!)^2 M^{2n^2}}$, then condition 1 is satisfied.*
2. *If for all $k \in \mathbb{N}$, the following holds, then condition 2 is satisfied:*

$$\frac{\alpha_{k+1}}{\alpha_k} \leq \frac{\delta_{\min}^n(1 - \delta_{\min})}{8(n!)^2 M^{2n^2} + 1}.$$

A.1.1 Arranging α_0

We start by giving a bound on $\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath})$ that involves α_0 . To do so, we make use of Lemma 6.

► **Corollary 14** (Another Lower Bound of *crossPath* Probability). *For all strategy pairs $\sigma, \gamma \in \Sigma_{\exists} \times \Sigma_{\forall}$, for all $\bar{v} \in \bar{V}$, the following holds:*

$$\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath}) > \frac{\delta_{\min}^n(1 - \alpha_0)^{n+1}}{2\alpha_0 + \delta_{\min}^n(1 - \alpha_0)^{n+1}}$$

Proof. We scale down $\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath})$ as follows:

$$\begin{aligned} & \widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath}) \\ & \geq \frac{(1 - x_0)x_0^n}{(1 - x_0) - (1 - x_0^n)x_1} && \text{Lemma 6} \\ & = \frac{(1 - \delta_{\min}(1 - \alpha_0))\delta_{\min}^n(1 - \alpha_0)^n}{\alpha_0 + (1 - \delta_{\min})\delta_{\min}^n(1 - \alpha_0)^{n+1}} \\ & > \frac{(1 - \delta_{\min})(1 - \alpha_0)\delta_{\min}^n(1 - \alpha_0)^n}{\alpha_0 + (1 - \delta_{\min})\delta_{\min}^n(1 - \alpha_0)^{n+1}} && \text{since } 1 - \delta_{\min}(1 - \alpha_0) > (1 - \delta_{\min})(1 - \alpha_0) \\ & = \frac{\delta_{\min}^n(1 - \alpha_0)^{n+1}}{\frac{\alpha_0}{1 - \delta_{\min}} + \delta_{\min}^n(1 - \alpha_0)^{n+1}} \\ & \geq \frac{\delta_{\min}^n(1 - \alpha_0)^{n+1}}{2\alpha_0 + \delta_{\min}^n(1 - \alpha_0)^{n+1}} && \text{since } 1 - \delta_{\min} \geq \frac{1}{2} \end{aligned}$$

◀

We can now arrange α_0 . It follows from Corollary 14 that:

$$\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath}) > \frac{\delta_{\min}^n(1 - \alpha_0)^{n+1}}{2\alpha_0 + \delta_{\min}^n(1 - \alpha_0)^{n+1}}$$

Therefore to show:

$$\widetilde{\Pr}_{\sigma, \gamma}^{\bar{v}}(\text{crossPath}) > \frac{4 - \epsilon}{4} \tag{5}$$

where $\epsilon = \frac{1}{(n!)^2 M^{2n^2}}$, it suffices to show that:

$$\frac{\delta_{\min}^n(1 - \alpha_0)^{n+1}}{2\alpha_0 + \delta_{\min}^n(1 - \alpha_0)^{n+1}} \geq \frac{4 - \epsilon}{4} \tag{6}$$

which can be further simplified as:

$$\epsilon \geq \frac{8\alpha_0}{2\alpha_0 + \delta_{\min}^n(1 - \alpha_0)^{n+1}} \tag{7}$$

We show that when $\alpha_0 \leq \frac{\delta_{min}^n}{8(n!)^2 M^{2n^2}}$, inequality (7) holds. We start with the right side:

$$\begin{aligned}
& \frac{8\alpha_0}{2\alpha_0 + \delta_{min}^n(1 - \alpha_0)^{n+1}} \\
& \leq \frac{8\alpha_0}{2\alpha_0 + \delta_{min}^n(1 - (n+1)\alpha_0)} \quad \text{follows from Bernoulli's inequality} \\
& = \frac{8\alpha_0}{\alpha_0 + \delta_{min}^n + \alpha_0(1 - (n+1)\delta_{min}^n)} \\
& \leq \frac{8\alpha_0}{\alpha_0 + \delta_{min}^n} \quad \text{since } 1 - (n+1)\delta_{min}^n > 0 \\
& \leq \frac{\frac{\delta_{min}^n}{(n!)^2 M^{2n^2}}}{\frac{\delta_{min}^n}{8(n!)^2 M^{2n^2}} + \delta_{min}^n} \\
& = \frac{1}{\frac{1}{8} + (n!)^2 M^{2n^2}} \\
& \leq \frac{1}{(n!)^2 M^{2n^2}} = \epsilon
\end{aligned}$$

Therefore we obtain that when $\alpha_0 \leq \frac{\delta_{min}^n}{8(n!)^2 M^{2n^2}}$, inequality (5) holds, and thus condition (1) is satisfied.

A.1.2 Arranging α_{k+1}/α_k

We start by getting a lower bound on $\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven})$ for even k 's, which makes use of α_{k+1} and α_k . The reasoning for odd k 's is symmetric. To do so, we use Lemma 7.

► **Corollary 15** (Another Lower Bound of *winEven* Probability). *For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all even k , the following holds:*

$$\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven}) > \frac{\delta_{min}^n(1 - \delta_{min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{min}^n(1 - \delta_{min}) + \frac{\alpha_{k+1}}{\alpha_k}}$$

Proof. We scale down the right side as follows:

$$\begin{aligned}
& \widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven}) \\
& \geq (1 - \alpha_{k+1}) \cdot \frac{(1 - x_2) \cdot x_2^{n-1} \cdot x_4}{1 - (x_2 + x_3) + x_5 \cdot x_2^n + t \cdot x_2^{n-1} - x_5 \cdot x_2^{n-1}} \quad \text{Lemma 7} \\
& = \frac{(1 - \delta_{min} + \delta_{min}\alpha_{k+1}) \cdot \delta_{min}^{n-1}(1 - \alpha_{k+1})^n \cdot \delta_{min}\alpha_k}{\alpha_{k+1} + \delta_{min}^{n-1}(1 - \alpha_{k+1})^{n-1}(\delta_{min}^2(-\alpha_k + \alpha_{k+1} + \alpha_k\alpha_{k+1} - \alpha_{k+1}^2) + \delta_{min}(\alpha_k - 2\alpha_{k+1} + \alpha_{k+1}^2))} \\
& > \frac{(1 - \delta_{min}) \cdot \delta_{min}^n \cdot (1 - \alpha_{k+1})^{n+1} \cdot \alpha_k}{\alpha_{k+1} + \delta_{min}^n \cdot (\delta_{min}(-\alpha_k + \alpha_{k+1} + \alpha_k\alpha_{k+1} - \alpha_{k+1}^2) + (\alpha_k - 2\alpha_{k+1} + \alpha_{k+1}^2))} \\
& \quad \text{since } 1 - \delta_{min} + \delta_{min}\alpha_{k+1} > (1 - \delta_{min})(1 - \alpha_{k+1}) \text{ and } (1 - \alpha_{k+1})^{n-1} < 1 \\
& \geq \frac{(1 - \delta_{min})\delta_{min}^n\alpha_k(1 - (n+1)\alpha_{k+1})}{\alpha_{k+1} + \delta_{min}^n(\alpha_k(1 - \delta_{min}) - \alpha_{k+1}(2 - \alpha_{k+1} - \delta_{min}(1 + \alpha_k - \alpha_{k+1})))} \\
& \quad \text{follows from Bernoulli inequality} \\
& > \frac{(1 - \delta_{min})\delta_{min}^n\alpha_k - (n+1)\delta_{min}^n(1 - \delta_{min})\alpha_k\alpha_{k+1}}{\alpha_{k+1} + (1 - \delta_{min})\delta_{min}^n\alpha_k} \\
& \quad \text{since } \alpha_{k+1}(2 - \alpha_{k+1} - \delta_{min}(1 + \alpha_k - \alpha_{k+1})) > 0 \\
& > \frac{\delta_{min}^n\alpha_k(1 - \delta_{min}) - \alpha_{k+1}}{\delta_{min}^n\alpha_k(1 - \delta_{min}) + \alpha_{k+1}} \quad \text{since } (n+1)\delta_{min}^n(1 - \delta_{min})\alpha_k < 1 \\
& = \frac{\delta_{min}^n(1 - \delta_{min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{min}^n(1 - \delta_{min}) + \frac{\alpha_{k+1}}{\alpha_k}}
\end{aligned}$$

◀

We now arrange α_{k+1}/α_k . It follows from Corollary 15 that for all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all even k , the following holds:

$$\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven}) > \frac{\delta_{min}^n(1 - \delta_{min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{min}^n(1 - \delta_{min}) + \frac{\alpha_{k+1}}{\alpha_k}} \quad (8)$$

Therefore to show:

$$\widetilde{\Pr}_{\sigma, \gamma}^k(\text{winEven}) \geq \frac{4}{4 + \epsilon} \quad (9)$$

it suffices to show that for all $k \in \mathbb{N}$:

$$\frac{\delta_{min}^n(1 - \delta_{min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{min}^n(1 - \delta_{min}) + \frac{\alpha_{k+1}}{\alpha_k}} \geq \frac{4}{4 + \epsilon} \quad (10)$$

We denote $\frac{\alpha_{k+1}}{\alpha_k}$ with r in the following calculation. Inequality 10 can be further simplified to:

$$\frac{\epsilon}{4 + \epsilon} \geq \frac{2r}{(1 - \delta_{min})\delta_{min}^n + r} \quad (11)$$

and can be finally simplified to:

$$r \leq \frac{(1 - \delta_{min})\delta_{min}^n}{8(n!)^2 M^{2n^2} + 1} \quad (12)$$

Therefore we obtain that if for all $k \in \mathbb{N}$, the following holds:

$$\frac{\alpha_{k+1}}{\alpha_k} \leq \frac{(1 - \delta_{min})\delta_{min}^n}{8(n!)^2 M^{2n^2} + 1}$$

then inequality (9) holds, and thus condition (2) is satisfied.