A Reduction from Stochastic Parity Games to Stochastic Reachability Games

by
Zihan Zhou

Master Thesis at RWTH Aachen University,
Lehrstuhl für Informatik 2

Submitted to:	Fakultät für Mathematik, Informatik und Naturwissenschaften der RWTH Aachen
Submission Date:	March 6, 2024
First supervisor:	Prof. Dr. Ir. Dr. h. c. Joost-Pieter Katoen
Second supervisor:	apl. Prof. Dr. Thomas Noll
Consultant:	Dr. Raphaël Berthon

Communicated by Joost-Pieter Katoen.

Abstract

Stochastic games (SGs) are used to represent complex decision-making scenarios where the outcomes are determined by strategic choices involving probabilistic events and adversarial behavior. In the landscape of SGs with ω-regular objectives, stochastic parity games (SPGs) are pivotal, yet efficient solutions to SPGs remain elusive. On the contrary, significant progress has been made in developing efficient solutions for simple stochastic games (SSGs), which are SGs with reachability objectives. In particular, value iteration algorithms with guarantees have been introduced for SSGs recently.

Although reductions from SPGs to SSGs are possible via several intermediate game models, a simple and direct reduction has been notably absent, which is the gap this thesis aims to bridge. Given an SPG, we construct an SSG that intuitively simulates the former. Through quantitative analysis of Markov Chains and their reachability equation systems, we establish the correctness of such direct reductions from both qualitative and quantitative SPGs to quantitative SSGs. Under binary encoding, we also show that both reductions are polynomial, corroborating the $\mathbf{NP} \cap \mathbf{coNP}$ complexity of SPGs. This work not only enriches the comprehension of SPGs, but also suggests potential solutions to SPGs by reducing them to SSGs and applying algorithms for SSGs.

Acknowledgements

I would like to express appreciation to my advisor, Raphaël, for offering his invaluable guidance throughout this work and taking his time to read the imperfect early versions.

I am grateful to Prof.Katoen for offering me the opportunity to write my thesis at MOVES, and my thanks also go to Prof.Noll for serving as the second examiner. I would thank all members of MOVES, for the high-quality lectures and seminars that introduced me to this field.

Furthermore, I wish to extend my gratitude to my family and friends. Your belief in me and your constant presence have been sources of comfort and motivation during the challenging moments. To Jin, your optimism, understanding, and love have been the bedrock of my resilience and perseverance.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

1.1.1 Stochastic Games

Games are utilized in various fields to model decision-making scenarios where multiple entities, known as players, interact strategically.

Take, for instance, the case of a market competition, where two firms A and B try to expand their market share. If a firm's market share falls below 10%, it gets acquired and exits the market.

- We represent the varying market shares of firms A and B as the game states. The impact of each firm's business decisions on market share is modeled as transitions between these states. This can be illustrated by a *graph*, enabling clear visualization of the competitive interplay.
- From the standpoint of each firm, the other acts as a direct competitor, and therefore we assume the players are *adversaries* in the game. We also incorporate *randomness* in the game to account for the unpredictable factors of the market, such as policy changes and external economic shocks.
- We encapsulate the strictly conflicting interests of both firms by considering a *zero-sum* game. In this setting, the success of one firm directly correlates with the loss of the other, particularly under the rule that a firm is acquired if its market share falls below a certain threshold.

This leads us to the notion of stochastic games.

Stochastic games (SG) [43] are zero-sum $2\frac{1}{2}$-player games played on graphs, where two players take sequential actions in their respective states. The uncertainty of the environment is represented by the other $\frac{1}{2}$ player, manifested as probability distributions governing the transitions between states. Unlike traditional games, where players make decisions in deterministic settings, stochastic games extend game theory to situations where uncertainty plays a role. SGs find their applications in various fields, including artificial intelligence [34], economics [1], operations research [17], and so on. Moreover, *Markov Decision Processes* (MDP) [41], 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.

1.1.2 Simple Stochastic Games and Solutions

Simple Stochastic Games. *Stochastic reachability games*, commonly addressed as *simple stochastic games* (SSG), were first introduced by Condon [14]. They are SGs where both players try to maximize and minimize the probability of reaching some target states respectively.

SSGs play a central role in stochastic games due to their connections to other objectives in SGs. Parity games, mean payoff games and discounted payoff games can all be reduced to SSGs [26, 46], and this also applies to their stochastic extensions, namely stochastic parity games [7], 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].

Solving Simple Stochastic Games. The complexity of solving SSGs, that is, calculating the winning probabilities or optimal strategies of all states, is in $\mathbf{UP} \cap \mathbf{coUP}$[1] as shown in [7]. A polynomial-time algorithm for solving SSGs has not been discovered.

When it comes to the specific algorithms, the work of Condon [17] is seminal, delineating three primary classes of algorithmic approaches: *value iteration* (VI), *strategy iteration* (SI) and *quadratic programming* (QP). The best-known upper bound of the runtime of all three classes of algorithms remains exponential in the size of the SSG. Comparative analyses, such as those conducted in [30], suggest a preference for VI and SI due to their relative efficiency.

Among these methods, only VI offers an *approximate* solution that converges in exponential steps. Standard VI stops when the difference between the two most recent approximations is small, however, it may lead to arbitrarily imprecise results. More specifically, standard VI approximates the exact values, which is a fixed point, only from below, leaving the distance to the fixed point uncertain. Fortunately, variants of standard VI with guarantees have been proposed recently for both MDPs and SSGs.

For MDPs, recent advancements have been made to apply VI with guarantees. Baier et al. [5] proposed interval VI for MDPs, which converges from both below and above to provide more accurate bounds on reachability probabilities. Sound VI, proposed by Quatmann et al. [42], also approximates the reachability probabilities from both below

[1] As a subset of **NP**, **UP** is defined by the condition that there exists a Turing machine such that for every input, there exists at most one computation path leading to an accepting state.

and above, ensuring a sound approximation. Optimistic VI, introduced by Hartmanns et al. [23], obtains a lower bound of the reachability probabilities via standard VI, then heuristically 'guesses' an upper bound and proves its correctness. This approach is notable for its ease of implementation and the fact that it does not require a priori computation of starting vectors for approximating from above.

For SSGs, similar ideas have been applied to obtain reachability probabilities with guarantees. Eisentraut et al. [19] introduce a VI algorithm for under- and over-approximation sequences, both of which converge to the value of the game, as well as the first practical stopping criterion for VI on SSGs. Azeem et al. [3] adapted the optimistic VI to SSGs and demonstrated its effectiveness in this setting. Phalakarn et al. [40] introduced a novel bounded VI with the concept of widest path, highlighting the potential for more effective solutions in complex decision-making scenarios.

Storm [24] and PRISM [31] are two of the most popular model checkers. Both incorporate different variants of standard VI and SI, and employ VI as the default algorithm for tackling MDPs. PRISM-games [32], an extension of PRISM for stochastic games, restricts the algorithm option exclusively to VI for tackling SSGs.

1.1.3 Stochastic Parity Games and Solutions

In games, ω-regular objectives define winning conditions based on infinite sequences of states. Parity objectives are a typical example, capturing exactly all ω-regular objectives. A parity objective categorizes game states with priorities and requires the lowest priority appearing infinitely often to be even.

Stochastic Parity Games. *Stochastic parity games* (SPG) are a natural extension of parity games to stochastic settings where the player (resp. the adversary) tries to maximize (resp. minimize) the probability of satisfying a parity objective. SPGs are of special interest, because every ω-regular objective can be turned into a parity objective by modifying the game graph [10]. This is achievable through the synchronous product of the game graph with a deterministic parity automaton that accepts the ω-regular objective [38].

Moreover, deterministic, as well as stochastic parity games, are closely connected to modal μ-calculus [29]. Emerson et al. [20] showed that solving parity games, that is, deciding the winner, can be reduced to modal μ-calculus model checking. This connection was enriched by Stirling [44], who showed that modal μ-calculus model checking can be reduced to solving parity games. Wilke [45] extended this connection by showing how modal μ-calculus model checking can be reduced to the acceptance problem for alternating tree automata, which in turn can be reduced to solving parity games. The probabilistic extension of the μ-calculus, qMμ, was introduced in [25, 37, 18]. McIver et al. [36] explored the probabilistic dimension of previous connections, showing that model checking qMμ can be reduced to solving the probabilistic generalization of Stirling's games, which are slightly modified stochastic parity games.

Solving Stochastic Parity Games. The complexity of solving SPGs, that is, calculating the winning probabilities or optimal strategies of all states, falls into $\mathbf{NP} \cap \mathbf{coNP}$ [15, 2]. We can list three primary approaches to solving SPGs as follows.

The first approach is to solve SPGs directly. Chatterjee et al. [8] proposed a strategy improvement algorithm for stochastic graph games with ω-regular conditions specified as parity objectives. This led to the development of a randomized sub-exponential time algorithm to solve SPGs. However, with n game states and d priorities, the expected running time goes to $2^{O(\sqrt{dn \log(n)})}$, which is not practically efficient enough.

The second approach is to first reduce SPGs to *deterministic parity games* (DPG), and then solve the DPG instance, as implemented in the probabilistic game solver GIST [11]. However, this approach only solves SPGs qualitatively, that is, determines the almost-sure winning states of both players. Moreover, solving DPGs efficiently also remains a highly challenging task itself. Several quasi-polynomial algorithms [27, 39, 33] have been proposed for DPGs since the breakthrough made by Calude et al. [6], but Czerwinski et al. [16] proved a quasi-polynomial lower bound on the size of a universal tree, which significantly reduces the likelihood for existing approaches to achieve polynomial running time.

The third approach is to first 'reduce' SPGs to SSGs, and apply model-free reinforcement learning to the SSG instance. Hahn et al. [22] introduced a method to construct a family of SSGs given an SPG, where the values of the SPG equal the values of the SSGs in the limit. This approach allows the use of reinforcement learning algorithms to approximate the values without the game's probabilistic transition structure. However, as the reduction is only correct in the limit, it is likely to produce false results.

To conclude, all existing approaches have theoretical and practical drawbacks that cannot be resolved without significant breakthroughs.

1.1.4 Reductions

Markov Chains (MC) are considered a special case of SGs and MDPs, where both players have no states under control, and the system evolves from one state to another based on probabilities that are solely dependent on the current state. In MCs, parity, or generally ω-regular objectives, can be reduced directly to reachability objectives [4].

However, such a direct reduction is missing in SGs and MDPs. Here we distinguish qualitative solutions, which are almost-sure winning states, from quantitative solutions, which are winning probabilities of all states. To reduce SPGs to quantitative SSGs, some intermediate steps are necessary, as depicted in dashed arrows in Figure 1.1.

- To reduce qualitative SPGs to quantitative SSGs, we first reduce qualitative SPGs to DPGs [12, 13] and then reduce DPGs to quantitative SSGs [7], as shown in the upper half of Figure 1.1.

- To reduce quantitative SPGs to quantitative SSGs, we use quantitative stochastic mean payoff and discounted payoff games as intermediate instances [9, 2], as shown in the lower half of Figure 1.1.

We also revisit these existing reductions in corresponding sections.

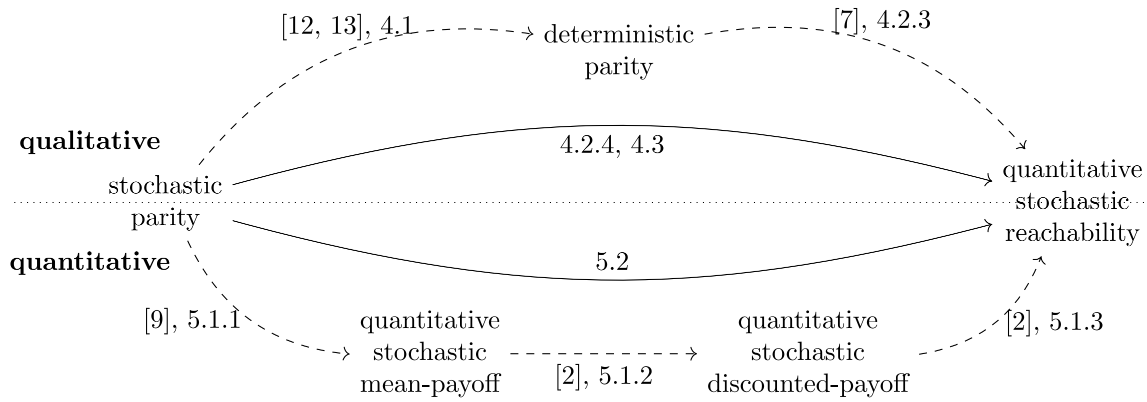

Figure 1.1: Reducing SPGs to SSGs

This thesis is set to look for simple and direct reductions from SPGs to SSGs, to establish a direct connection between two central notions in stochastic games. Furthermore, such a reduction can potentially yield a solution to SPGs by first reducing an SPG to an SSG, and then solving the SSG with existing algorithms.

1.2 Contribution

This thesis gives simple and direct reductions from both qualitative and quantitative SPGs to quantitative SSGs.

For qualitative SPGs, we leverage a gadget from the literature, which reduces DPGs to quantitative SSGs. We extend the idea of this gadget to qualitative SPGs, showing that it yields a direct reduction to quantitative SSGs under certain conditions. We show that these conditions are satisfiable when we arrange a function in the gadget properly, through probabilistic analysis of Markov Chains resulting from fixed strategies.

For quantitative SPGs, this gadget cannot yield a reduction in terms of exact values. However, we show that it reduces strategic SPGs to strategic SSGs under certain conditions, which are again satisfiable, with similar reasoning as in the qualitative case. In the case of SPGs and SSGs, strategic and quantitative solutions are trivially equivalent, and thus we obtain a 'direct' reduction from quantitative SPGs to quantitative SSGs.

In both cases, we show under binary encoding that the reductions are polynomial, and therefore both qualitative and quantitative SPGs are in the complexity class $\mathbf{NP} \cap \mathbf{coNP}$.

1.3 Organization

The organization of the contents is as follows:

In Chapter 2, we present the necessary background knowledge.

In Chapter 3, we clarify some terminologies and notations used later.

In Chapter 4, we first revisit the existing reduction chain from qualitative SPGs to quantitative SSGs. Then we make a direct reduction, and show how to satisfy the conditions ensuring the correctness of this reduction.

Similarly, in Chapter 5, we first revisit the existing reduction chain from quantitative SPGs to quantitative SSGs. Then we make a direct reduction and show how to satisfy the correctness conditions. In the end, we present the complexity results.

In Chapter 6, we conclude this thesis and give an outlook for future works.

Chapter 2

Preliminaries

In this chapter, we discuss the required background for this thesis, which hinges on the principles of Markov Chains and stochastic games on graphs.

2.1 Discrete-Time Markov Chains

A *discrete-time Markov Chain* models a system that undergoes transitions from one state to another, with the next state being dependent only on the current state. The transitions between states are probabilistic, meaning that a transition from one state to one of its possible successors happens with a certain probability. This model is pivotal in analyzing stochastic games.

2.1.1 Basics

Before delving into the specifics, we introduce the notion of *discrete distributions* for the sake of completeness.

Definition 2.1 (Discrete Distributions)
A *discrete distribution* over countable set $\mathcal{A}$ is a function $\mu : \mathcal{A} \to \mathbb{R}_{\geq 0}$ mapping each element of $\mathcal{A}$ to a non-negative weight (probability), such that:

$$\sum_{a \in \mathcal{A}} \mu(a) = 1$$

The *support* of discrete distribution μ, denoted with $supp(\mu)$, is defined as the elements in $\mathcal{A}$ with non-zero probability, formally as follows:

$$supp(\mu) \triangleq \{a \in \mathcal{A} \mid \mu(a) > 0\}$$

A *unit distribution* $unit(a)$ on $a \in \mathcal{A}$ refers to the distribution μ where $\mu(a) = 1$.
We denote the set of all discrete distributions on $\mathcal{A}$ with $\mathbb{D}(\mathcal{A})$. ■

The formal definition of DTMCs goes as follows.

Definition 2.2 (Discrete-Time Markov Chains)
A *discrete-time Markov Chain (DTMC)* $\mathcal{M}$ is a tuple $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ where:

- Σ is the finite set of states,
- $\delta : \Sigma \to \mathbb{D}(\Sigma)$ is a probabilistic transition function, where $\delta(\sigma, \sigma')$ indicates the probability of a transition from σ to σ'. It holds that for all $\sigma \in \Sigma$

$$\sum_{\sigma' \in \Sigma} \delta(\sigma, \sigma') = 1,$$

- $\sigma_I \in \Sigma$ is the initial state.

In the sequel, we refer to discrete-time Markov Chains (DTMCs) as *Markov Chains (MC)* for simplicity, unless stated otherwise explicitly. As Markov Chains with finite state spaces suffice for this thesis, we always assume that $|\Sigma| < \infty$. Given a set of states $S \subseteq \Sigma$, we define for all states $\sigma \in \Sigma$ that $\delta(\sigma, S) = \sum_{s \in S} \delta(\sigma, S)$. ■

A path in a Markov Chain refers to a sequence of states connected by transitions. The formal definition goes as follows.

Definition 2.3 (Paths in Markov Chains)
Let $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ be a Markov Chain.

- An infinite sequence $\pi = \sigma_0\sigma_1 \cdots \in \Sigma^\omega$ is an *infinite path* through $\mathcal{M}$ if for all $i \in \mathbb{N}$, $\delta(\sigma_i, \sigma_{i+1}) > 0$. We denote all the infinite paths that start from a state $\sigma \in \Sigma$ with *Paths*(σ).
- Given an infinite path $\pi = \sigma_0\sigma_1 \cdots \in \Sigma^\omega$, its *prefixes* are $\{\sigma_0 \cdots \sigma_i \mid i \in \mathbb{N}\}$.
- The prefixes of infinite paths are *finite paths*, and we denote all the finite paths that start from a state $\sigma \in \Sigma$ with *Paths*$^*(\sigma)$. ■

Given an infinite path $\pi = \sigma_0\sigma_1 \cdots \in \Sigma^\omega$, we are typically interested in the set of states that appear *infinitely often* in the long run, formally defined as follows.

Definition 2.4 (Infinitely Often Visited Set)
Let $\pi = \sigma_0\sigma_1 \cdots \in \Sigma^\omega$ be an infinite path. We define the set of infinitely often visited states as:

$$\mathit{inf}(\pi) = \{\sigma \in \Sigma \mid \forall n \in \mathbb{N}, \exists k \in \mathbb{N} \ s.t. \ \sigma_{n+k} = \sigma\}.$$ ■

Markov Chains can be visualized and analyzed using directed graphs States of an MC are represented as nodes in a directed graph, and the transitions between states are represented as directed edges between the corresponding nodes labeled with the probabilities.

Example 2.5 (A Markov Chain)
The following is a Markov Chain $\mathcal{M}$ with $\sigma_I = s_0$.

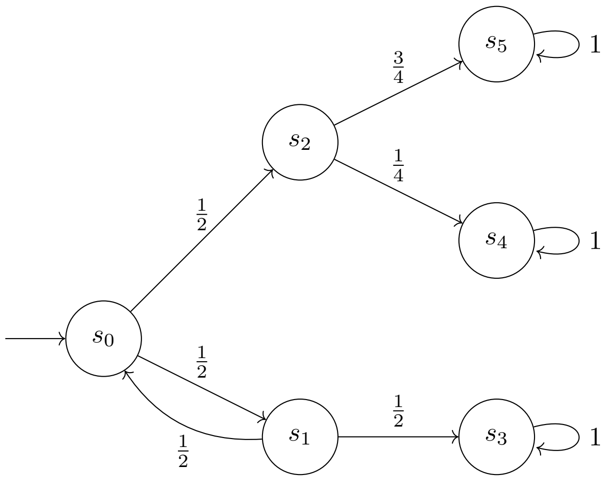

Figure 2.1: A Markov Chain $\mathcal{M}$ with $\sigma_I = s_0$

Applying the previously defined notions, we identify some infinite and finite paths as follows.

- The infinite sequence $\pi = s_0 s_2 s_4^\omega$ is an infinite path, and $\mathit{inf}(\pi) = \{s_4\}$.
- The prefixes of π, such as $s_0 s_2$ and $s_0 s_2 s_4^*$, are finite paths. ■

It comes quite intuitively that the probability of following an infinite path, say $\pi = \sigma_0\sigma_1 \cdots \in \Sigma^\omega$, is $\prod_{i \in \mathbb{N}} \delta(\sigma_i, \sigma_{i+1}) = 0$. However, we can capture the probability measure of a finite path with the notion of *cylinder sets.* Stated in words, the cylinder set of a finite path $\hat{\pi}$ contains all the infinite paths that are extensions of $\hat{\pi}$. The formal definition goes as follows.

Definition 2.6 (Cylinder Sets and Their Probability[4])
Let $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ be a Markov Chain.

- The *cylinder set* of a finite path $\hat{\pi} = \sigma_0\sigma_1 \ldots \sigma_n \in \mathit{Paths}^*(\sigma_0)$ is defined as:

$$\mathit{Cyl}(\hat{\pi}) = \{\pi \in \mathit{Paths}(\sigma_0) \mid \hat{\pi} \textit{ is a prefix of } \pi\}$$

- The probability measure of $\mathit{Cyl}(\hat{\pi})$ is defined as:

$$\Pr(\mathit{Cyl}(\sigma_0\sigma_1 \ldots \sigma_n)) = \prod_{0 \le i < n} \delta(\sigma_i, \sigma_{i+1})$$

  In the sequel, we refer to $\Pr(\mathit{Cyl}(\sigma_0\sigma_1 \ldots \sigma_n))$ as the *probability* of the finite path $\hat{\pi}$ and denote it with $\Pr(\sigma_0\sigma_1 \ldots \sigma_n)$ for simplicity.

- Given a set of finite path $\hat{\Pi} \subseteq \mathit{Paths}^*(\sigma_0)$ with the same starting state $\sigma_0 \in \Sigma$, the probability of $\hat{\Pi}$ is defined as:

$$\Pr(\hat{\Pi}) = \sum_{\hat{\pi} \in \hat{\Pi}} \Pr(\hat{\pi})$$

■

2.1.2 Reachability Probability

Reachability in Markov Chains is a fundamental concept that explores the ability of a system to reach a set of states from a given starting state. Specifically, given a Markov Chain $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ and a set of target states $G \subseteq \Sigma$, we aim to determine the probability of reaching target states G for all states $\sigma \in \Sigma$.

Formally, we define reachability probability as follows.

Definition 2.7 (Reachability Probability[4])
Let $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ be a Markov Chain. Given a set of target states $G \subseteq \Sigma$, and a starting state $\sigma_0 \in \Sigma$,

- The event of reaching G from σ_0 is defined as:

$$\begin{aligned} \mathit{Reach}(G) =& \{\sigma_0\sigma_1 \cdots \in \Sigma^\omega \mid \exists i \in \mathbb{N}, \sigma_i \in G\} \\ =& \bigcup_{\hat{\pi} \in \mathit{Paths}^*(\sigma_0) \cap ((\Sigma \backslash G)^* G)} \mathit{Cyl}(\hat{\pi}) \end{aligned}$$

- The probability of reaching G from σ_0 is defined as:

$$\Pr{}^{\sigma_0}(\mathit{Reach}(G)) = \Pr(\{\hat{\pi} \mid \hat{\pi} \in \mathit{Paths}^*(\sigma_0) \cap ((\Sigma \backslash G)^* G)\})$$

■

Let variable x_σ denote the probability of reaching G from any $\sigma \in \Sigma$. It is trivial to inspect if G is reachable from σ with a graph analysis. We denote the set of states from which G is reachable with $\mathit{Pre}^*(G)$. If $\sigma \notin \mathit{Pre}^*(G)$, then $x_\sigma = 0$. If $\sigma \in G$, then $x_\sigma = 1$. Otherwise,

$$x_\sigma = \sum_{\tau \in \Sigma \backslash G} \delta(\sigma, \tau) \cdot x_\tau + \sum_{\gamma \in G} \delta(\sigma, \gamma).$$

These equations can be written into a linear equation system as follows.

Theorem 2.8 (Reachability Probability of Markov Chains[4])
Given a Markov Chain $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ and a target state set $G \subseteq \Sigma$,

- let $\Sigma_Q = \mathit{Pre}^*(G) \backslash G$, which is the states that can reach G by at least one step,
- $\mathbf{A} = (\delta(\sigma, \tau))_{\sigma, \tau \in \Sigma_Q}$, which is the transition probability in Σ_Q,
- $\mathbf{b} = (b_\sigma)_{\sigma \in \Sigma_Q}$, which is the probability of reaching G in one step.

The vector $\mathbf{x} = (x_\sigma)_{\sigma \in \Sigma_Q}$ where x_σ is the probability of reaching target states G from σ is the unique solution of the linear equation system:

$$\mathbf{x} = \mathbf{A} \cdot \mathbf{x} + \mathbf{b}. \qquad \blacksquare$$

Example 2.9 (A Linear Equation System for Reachability Probability)
We look at the MC $\mathcal{M}$ used in Example 2.5. Let the target state set $G = \{s_4\}$. As a result, we have:

$$\begin{aligned} Pre^*(G) &= \{s_0, s_1, s_2, s_4\} \\ \Sigma_Q &= Pre^*(G) \backslash G = \{s_0, s_1, s_2\} \\ \mathbf{A} &= \begin{pmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \\ \mathbf{b}^T &= \begin{pmatrix} 0 & 0 & 1/4 \end{pmatrix} \end{aligned}$$

Solving the linear equation system, we obtain:

$$\mathbf{x}^T = \begin{pmatrix} \frac{1}{6} & \frac{1}{12} & \frac{1}{4} \end{pmatrix}$$

Thus the probability of reaching s_4 from s_0 is $\frac{1}{6}$. $\blacksquare$

2.1.3 Limit Behavior

The *limit behavior* of a Markov Chain refers to what happens to the chain as it evolves over an infinite number of time steps. We are typically interested in the set of states that are visited infinitely often.

We first introduce a few graph theory-based definitions for Markov Chains.

Definition 2.10 ((Bottom) Strongly Connected Components)
Let $\mathcal{M} = (\Sigma, \delta, \sigma_I)$ be a Markov Chain.

- A set of states $L \subseteq \Sigma$ is *strongly connected* if for all pairs of states $\sigma, \sigma' \in L$, there exists a finite path $\hat{\pi} = \sigma_0\sigma_1 \ldots \sigma_n$ with $n \geq 1$, such that $\sigma_0 = \sigma$, $\sigma_n = \sigma'$ and $\sigma_i \in L$ for all $0 \leq i \leq n$. Hence a single-state set $\{\sigma\}$ is strongly connected if $\delta(\sigma, \sigma) > 0$.
- A set of states $L \subseteq \Sigma$ is referred to as a *strongly connected component (SCC)* if it is maximally strongly connected, i.e., there does not exist another set of states $L' \subseteq \Sigma$ and $L \subsetneq L'$ such that L' is strongly connected.
- A set of states $L \subseteq \Sigma$ is referred to as a *bottom strongly connected component (BSCC)* if L is a SCC and there is no transition leaving L, i.e., there does not exist $\sigma \in L, \sigma' \in \Sigma \backslash L$ such that $\delta(\sigma, \sigma') > 0$.

We denote the set of BSCCs in MC $\mathcal{M}$ with $\mathit{BSCC}(\mathcal{M})$. ■

To give a clear intuition of (bottom) strongly connected components, we revisit the Markov Chain presented in previous examples.

Example 2.11 ((Bottom) Strongly Connected Components)
We recall the MC $\mathcal{M}$ used in Example 2.5. Four SCCs in this Markov Chain are easily recognizable, so $\mathit{BSCC}(\mathcal{M}) = \{\{s_0, s_1\}, \{s_3\}, \{s_4\}, \{s_5\}\}$.

- Since there is a transition from s_0 to s_2 (and a transition from s_1 to s_3), $\{s_0, s_1\}$ is not a BSCC.
- Since s_3 has only a self loop with probability 1, the singleton set $\{s_3\}$ is a BSCC. The same goes for $\{s_4\}$ and $\{s_5\}$. ■

The limit behavior of a Markov Chain regarding the infinitely often visited states is captured by the following theorem.

Theorem 2.12 (Limit behavior of Markov Chains[4])
Given a Markov Chain $\mathcal{M} = (\Sigma, \delta, \sigma_I)$, the following holds:

$$\Pr\{\pi \in \mathit{Paths}(\sigma_I) \mid \mathit{inf}(\pi) \in \mathit{BSCC}(\mathcal{M})\} = 1$$ ■

Recall we have defined the set of states that are visited infinitely often in Definition 2.4. Intuitively speaking, for all states $\sigma \notin \mathit{inf}(\pi)$, there is a point after which σ is not visited any more. In other words, from some point on all the states visited subsequently are $\mathit{inf}(\pi)$, and each of them is visited infinitely often. Therefore Theorem 2.12 can be interpreted as: almost surely any Markov Chain eventually reaches a BSCC and visits all its states infinitely often.

Example 2.13 (Limit Behavior)
Again we look at the MC $\mathcal{M}$ used in previous Example 2.5. We can conclude from Theorem 2.12 that almost surely it eventually reaches one of its BSCCs, namely $\{s_3\}$, $\{s_4\}$ and $\{s_5\}$, and stays there forever. ■

2.2 Stochastic Games on Graphs

In this thesis, we study stochastic games on graphs under a few restrictions, specifically listed as follows.

Perfect-information. By *perfect-information* we refer to the situations where all the players have accurate knowledge about the specifications of the game and the moves made by other players. In contrast, there are games with imperfect information, which leads to a higher level of uncertainty.

$2\frac{1}{2}$-player. There are formally two players in a game, and we address them as Eve and Adam. As we study stochastic games, we refer to stochasticity as the other $\frac{1}{2}$ 'player'. When the game is deterministic, we call it a 2-player game.

Turn-based. The adjective *turn-based* means the formal players, Eve and Adam, take turns to make their moves in the game. Turn-based games are a special case of concurrent games, where all the players make decisions simultaneously and the game proceeds without distinct turns.

Zero-Sum. By *zero-sum*, we refer to situations where the total resources remain constant, and any advantage gained by one player comes at the direct expense of other players. Thus the winning objectives of two players are always complementary.

On Finite Graphs. By *finite graphs* we refer to directed graphs with a finite number of vertices and edges. We always assume a directed graph is finite in the sequel.

In this section, we present in detail the definitions and some relevant theorems regarding this kind of game.

2.2.1 Stochastic Arena

A *stochastic arena* is where a stochastic game is played, which is a finite graph in our case. The states of the game are encoded as vertices, and the possible moves as directed edges. The game is played with players moving a token from vertex to vertex to produce an infinite path. As the game is turn-based, Eve has a set of vertices under control, where she moves the token to a successor vertex, and the same goes for Adam. Sometimes, neither Eve nor Adam has control over the vertices, and the successor vertex is determined stochastically. This leads us to the formal definition of a *stochastic arena*.

Definition 2.14 (Stochastic Arena)
A *stochastic arena* G is a tuple $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ where:

- (V, E) is a directed graph, V is a finite set of vertices and $E \subseteq V \times V$ is a set of edges,

- $(V_\exists, V_\forall, V_R)$ is a vertex partition such that $V_\exists \uplus V_\forall \uplus V_R = V$,
- $\Delta : V_R \to \mathbb{D}(V)$ is a probabilistic transition function such that for all $v_r \in V_R$ and $(v_r, v) \notin E$, $\Delta(v_r, v) = 0$, and
$$\sum_{(v_r,v)\in E} \Delta(v_r, v) = 1.$$

Without loss of generality, we assume each vertex has at least one successor, and we call this property *non-blocking.* ■

The finite set of vertices V is partitioned into three disjoint sets: $V_\exists$ -vertices where Eve chooses the successor, $V_\forall$ -vertices where Adam chooses the successor, and V_R -random vertices. For a random vertex $v \in V_R$, $\Delta(v)$ gives a distribution on v's possible successors. A stochastic arena is a Markov Decision Process (MDP) if $V_\exists = \emptyset$ or $V_\forall = \emptyset$.

The following is an example of a stochastic arena. We use squares to represent vertices in $V_\exists$, which are controlled by Eve, and pentagons for vertices in $V_\forall$, which are controlled by Adam, as 'Eve' is one letter shorter than 'Adam' and squares have one fewer edges than pentagons. Random vertices in V_R are depicted as circles. In the sequel, we always comply with this visual convention.

Example 2.15 (A Stochastic Arena)
The following is a stochastic arena $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$.

- $V_\exists = \{v_4\}$, $V_\forall = \{v_1, v_2\}$ and $V_R = \{v_0, v_3\}$,
- The values of Δ are labeled on the transitions. ■

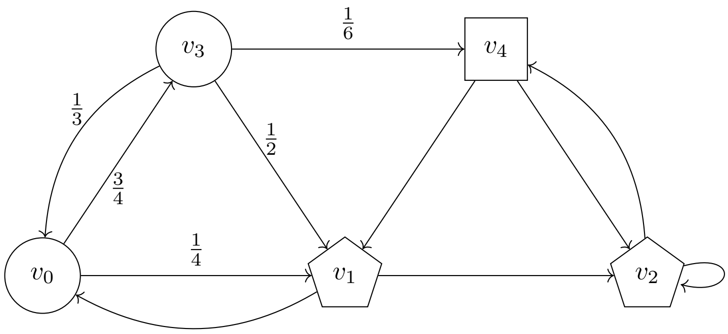

Figure 2.2: A stochastic arena G

2.2.2 Strategies

Players choose their moves at game states according to their *strategies.* Generally speaking, a strategy can be a highly complicated object, depending on long game histories and giving nondeterministic decisions. The general definition of strategies goes as follows.

Definition 2.16 (Strategies)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena.

- A *strategy* σ of Eve is defined as a function $\sigma : V^* \cdot V_\exists \to \mathbb{D}(V)$, such that for all $v_0 v_1 \dots v_n \in V^* \cdot V_\exists$, we have $\sigma(v_0 v_1 \dots v_n, v_{n+1}) > 0$ only if $(v_n, v_{n+1}) \in E$.
- A strategy γ of Adam is defined analogously.
- We denote the sets of all strategies of Eve and Adam with $\mathit{\Sigma}^A$ and $\mathit{\Gamma}^A$ respectively. ■

This general notion of strategies does not pose any restrictions on the size of the game histories. It also allows a strategy to be randomized. Fortunately, game theorists have proved that much simpler strategies are sufficient in many cases. In the following, we introduce *pure memoryless strategies*, which is a simple sub-class of strategies.

Definition 2.17 (Pure Memoryless Strategies)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. A strategy σ of Eve is a *pure memoryless strategy*, if for all $w, w' \in V^*$ and $v \in V_\exists$, $\sigma(w \cdot v) = \sigma(w' \cdot v) = \mathit{unit}(v')$ for some $v' \in V$ and $(v, v') \in E$. A pure memoryless strategy γ of Adam is defined analogously. ■

A pure memoryless strategy is deterministic and independent of the game's history. A player following a pure memoryless strategy always chooses the same unique successor vertex given the current vertex. Note that a pure memoryless strategy of Eve or Adam can be represented as a simpler function $\sigma : V_\exists \to V$ or $\gamma : V_\forall \to V$ respectively, and we always employ this notion in the sequel. We denote the sets of pure memoryless strategies of Eve and Adam with Σ and Γ respectively.

Example 2.18 (Pure Memoryless Strategies)
Recall the stochastic arena G in Example 2.15.

- Let Eve follow pure memoryless strategy σ, where $\sigma(v_4) = v_2$.
- Let Adam follow pure memoryless strategy γ, where $\gamma(v_1) = v_2$ and $\gamma(v_2) = v_4$.

We present Figure 2.3 for some intuition.

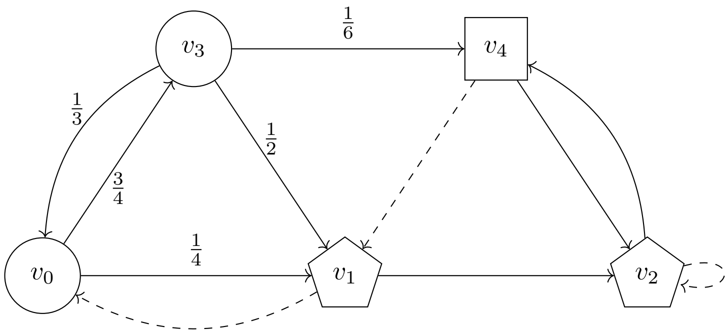

Figure 2.3: Eve and Adam following σ and γ respectively in G

The transitions not taken are depicted with dashed arrows. ■

In the example above, we observe that when Eve and Adam both follow pure memoryless strategies, each vertex $v \in V_\exists \uplus V_\forall$ has a unique successor vertex. The transitions not taken, depicted with dashed arrows, can be ignored, and what is left still fits the definition of an arena. We introduce a special notion of *sub-arena* defined as follows.

Definition 2.19 (Sub-Arenas)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. When Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma$ and $\gamma \in \Gamma$ in G respectively, we call the resulting arena a *sub-arena* denoted with $G_{\sigma,\gamma} = ((V, E'), (V_\exists, V_\forall, V_R), \Delta)$. The new edge set E' is obtained by:

- For all $u \in V_\exists$, $(u, v) \in E'$ if and only if $\sigma(u) = v$.
- For all $u \in V_\forall$, $(u, v) \in E'$ if and only if $\gamma(u) = v$. ■

In this case, the non-determinism in the arena is resolved, and given a fixed starting vertex $v_I \in V$, we can view the sub-arena $G_{\sigma,\gamma}$ as a Markov Chain defined as follows.

Definition 2.20 (Markov Chain View of Sub-Arenas)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. Let Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma$ and $\gamma \in \Gamma$ respectively, and a sub-arena $G_{\sigma,\gamma}$ is obtained. Given a fixed starting vertex $v_I \in V$, we define the Markov Chain view of the sub-arena as a Markov Chain denoted with $\mathcal{M}_{\sigma,\gamma} = (V, \delta, v_I)$, where:

- The state space is the set V of vertices from the arena.
- The transition function combines deterministic moves indicated by strategies

and the original transition function defined on random vertices:

$$\delta(u,v) = \begin{cases} 1 & \text{if } u \in V_\sigma, \sigma(u) = v \text{ or } u \in V_\gamma, \gamma(u) = v, \\ \Delta(u,v) & \text{if } u \in V_R, \\ 0 & \text{otherwise} \end{cases}$$

■

Example 2.21 (Markov Chain View of Sub-Arenas)
We recall Example 2.15 and Example 2.18. Given a starting vertex v_0, the Markov Chain view of the sub-arena $G_{\sigma,\gamma}$ is given in the following Figure 2.4.

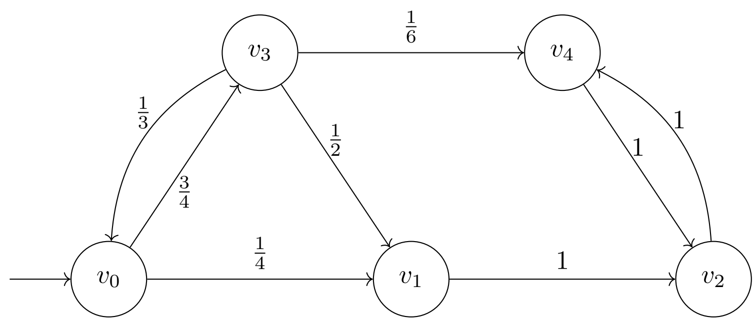

Figure 2.4: The Markov Chain view of $G_{\sigma,\gamma}$

Remark 2.22 (Implicit Reference to the Markov Chain View)
In the sequel, when we apply terminologies from Markov Chains to a sub-arena $G_{\sigma,\gamma} = ((V, E'), (V_\exists, V_\forall, V_R), \Delta)$, we implicitly refer to its Markov Chain view $\mathcal{M}_{\sigma,\gamma} = (V, \delta, v_I)$, where $v_I \in V$ is an arbitrary state. For clarity, we do not explicitly specify an initial state in $\mathcal{M}_{\sigma,\gamma}$, for the following reasons:

- When we reason about BSCCs, the initial state does not matter.
- When we reason about probabilities of finite paths with the notation $\Pr_{\sigma,\gamma}$, the starting vertex v depends on the paths.
- When we reason about reachability probabilities with the notation $\Pr^v_{\sigma,\gamma}$, we explicitly indicate the starting vertex v. ■

2.2.3 Plays and Objectives

When Eve and Adam play in a stochastic arena, we determine the winner based on the infinite sequence of vertices that are visited, which we call a *play*.

Definition 2.23 (Plays)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. We define a *play* as 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 with Π. ■

Both players, Even and Adam, follow certain strategies to win the game, and we now introduce their *winning objectives*. We mainly discuss two classes of objectives, namely *qualitative* and *quantitative* objectives.

Qualitative Winning Objectives

For games played in a stochastic arena, we define a *qualitative winning objective* for Eve as a set of plays $\Phi \subseteq \Pi$. As we study zero-sum games, the winning objectives of the two players are complementary. Therefore the winning objective for Adam is $\Pi \backslash \Phi$. We say a play ϕ *satisfies* an objective Φ if $\phi \in \Phi$, and it is a *winning play* of Eve. Otherwise, ϕ satisfies the objective $\Pi \backslash \Phi$, and it is thus a winning play of Adam.

In the following part, we present several classes of typical qualitative objectives. Note that we discuss winning objectives from the perspective of Eve with a fixed starting vertex $v_0 \in V$, unless indicated otherwise explicitly.

Safety. A *safety* objective asserts that the game never reaches some undesirable states, which are given as a set $T \in V$ of 'target' vertices. The winning plays are those that never visit the 'target' vertices. Formally, a safety objective is defined as:

$$SA(T) = \{v_0 v_1 \dots \in \Pi \mid \forall k \in \mathbb{N}, v_k \notin T\}$$

Reachability. On the contrary to a safety objective, a *reachability* objective asserts that the game has to reach some desirable states given as a set $T \in V$ of target vertices. When Eve's objective is a safety objective, Adam's objective is a reachability objective. Formally, a reachability objective is defined as:

$$RE(T) = \{v_0 v_1 \dots \in \Pi \mid \exists k \in \mathbb{N}, v_k \in T\}$$

Parity. To specify a *parity* objective, we first introduce a *priority function* $p : V \to \mathbb{N}$, which assigns a priority $p(v)$ to each vertex $v \in V$. Given a set of vertices T, we define $p(T) = \{p(t) \mid t \in T\}$. Stated in words, a parity objective asserts that the minimum priority visited infinitely often along an infinite path is even. Formally, given a play $\pi = v_0 v_1 \dots \in \Pi$, we define analogously to Definition 2.4 the infinitely often visited set of vertices as:

$$\mathit{inf}(v_0 v_1 \dots) = \{v \in V \mid \forall n \in \mathbb{N}, \exists k \in \mathbb{N} \ \mathit{s.t.} \ v_{n+k} = v\}$$

Then we can formalize a parity objective as:

$$PA(p) = \{\pi = v_0 v_1 \dots \in \Pi \mid \ \min(p(\mathit{inf}(\pi))) \text{ is even}\}$$

Quantitative Winning Objectives

To specify a quantitative winning objective, we introduce a measurable objective function $f : \Pi \to \mathbb{R}$ mapping a play to a real number as its outcome. Eve tries to maximize the value of f, while Adam tries to minimize it, i.e. to maximize the value of $-f$. In this thesis, we mainly consider payoffs, which are a typical class of quantitative objectives.

We associate a reward with each vertex, given by a function $r : V \to \mathbb{R}$. The value of $r(v)$ is rewarded whenever a vertex $v \in V$ is visited. As a game goes on forever, the sum of the rewards can be unbounded. There are typically two ways to obtain a bounded payoff function, namely taking the mean value, or discounting the rewards with a decreasing rate.

Mean Payoff. Let $r : V \to \mathbb{R}$ be a reward function. The *mean payoff* function calculates the limiting average rewards. Specifically, given a play $\pi = v_0 v_1 \cdots \in \Pi$, the mean payoff function is defined as:

$$MP(r)(\pi) = \liminf_{n\to\infty} \frac{1}{n+1} \sum_{i=0}^{n} r(v_i)$$

Discounted Payoff. Let $r : V \to \mathbb{R}$ be a reward function and $\lambda \in (0,1)$ be a discount factor. The *discounted payoff* function discounts the i^{th} reward with the discounting rate $(1-\lambda)\cdot\lambda^i$. Formally, given a play $\pi = v_0 v_1 \cdots \in \Pi$, the discounted payoff function is defined as:

$$DP(r,\lambda)(\pi) = (1-\lambda) \lim_{n\to\infty} \sum_{i=0}^{k} \lambda^i \cdot r(v_i)$$

2.2.4 Define A Stochastic Game, Formally

We now have the necessary ingredients to formally define a *stochastic game (SG)*.

Definition 2.24 (Stochastic Games)
Let $G = ((V,E),(V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. Let Φ be a qualitative objective, and let f be a quantitative objective function.

- A stochastic game with a qualitative winning objective is defined with (G,Φ). If Φ is a reachability or parity objective, we call (G,Φ) a *stochastic reachability game (SRG)* or *stochastic parity game (SPG)* respectively. We also address SRGs as *simple stochastic games* (SSG).
- If there are no random vertices in G, formally $V_R = \emptyset$, we call (G,Φ) a *deterministic reachability game (DRG)* or a *deterministic parity game (DPG)* respectively.

- A stochastic game with a quantitative winning objective is defined with (G, f). If f is a mean payoff or discounted payoff objective function, we call (G, f) a *stochastic mean payoff game (SMPG)* or *stochastic discounted payoff game (SDPG)* respectively. ■

2.2.5 Values and Optimal Strategies

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, and let Eve and Adam follow strategies $\sigma \in \Sigma^A$ and $\gamma \in \Gamma^A$. We specify a starting vertex $v \in V$. Now the game is on, and a play $\pi = vv_1 \cdots \in \Pi$ is generated.

If Eve aims for a qualitative objective Φ, we obtain a stochastic game (G, Φ). The expected outcome of the game is the probability for π to satisfy Φ for Eve. We denote this probability with $\mathbb{P}^v_{\sigma,\gamma}(\Phi)$. On the contrary, the expected outcome of the game is $\mathbb{P}^v_{\sigma,\gamma}(\Pi \backslash \Phi)$ for Adam.

If Eve aims for a quantitative objective represented by function f, we obtain a stochastic game (G, f). The expected outcome of the game is the expected value of $f(\pi)$ for Eve. We denote this expected value with $\mathbb{E}^v_{\sigma,\gamma}(f)$. On the contrary, the expected outcome of the game is $\mathbb{E}^v_{\sigma,\gamma}(-f)$ for Adam.

We introduce the notion of *value* of a vertex v, which represents the best possible expected outcome of a play starting from v, given the opponent player plays optimally. If a player can achieve the value from v following a specific strategy, this strategy is *optimal* from v.

- In the case of a stochastic game with a qualitative objective (G, Φ), we define the *value* of a vertex v as the maximal probabilities of generating a play from v that satisfies Φ, formally defined as follows.

  Definition 2.25 (Values for Qualitative Objectives[7])
  Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena.

  - The *value functions*, denoted with $\langle E\rangle(\Phi) : V \to [0, 1]$ for Eve and $\langle A\rangle(\Pi\backslash\Phi) : V \to [0, 1]$ for Adam, are defined as:

  $$\langle E\rangle(\Phi)(v) = \sup_{\sigma\in\Sigma^A} \inf_{\gamma\in\Gamma^A} \mathbb{P}^v_{\sigma,\gamma}(\Phi)$$
  $$\langle A\rangle(\Pi\backslash\Phi)(v) = \sup_{\gamma\in\Gamma^A} \inf_{\sigma\in\Sigma^A} \mathbb{P}^v_{\sigma,\gamma}(\Pi\backslash\Phi)$$

  - A strategy σ for Eve is *optimal* from vertex v if the following holds:

  $$\inf_{\gamma\in\Gamma^A} \mathbb{P}^v_{\sigma,\gamma}(\Phi) = \langle E\rangle(\Phi)(v)$$

  The optimal strategies for Adam are defined analogously. ■

- In the case of a stochastic game with a quantitative objective (G, f), we define the value of a vertex v as the maximal expected values of f over a play starting from v, formally defined as follows.

  Definition 2.26 (Values for Quantitative Objectives[7])
  Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena.

  - The *value functions*, denoted with $\langle E\rangle(f) : V \to \mathbb{R}$ for Eve and $\langle A\rangle(-f) : V \to \mathbb{R}$ for Adam, are defined as:

  $$\langle E\rangle(f)(v) = \sup_{\sigma\in\Sigma^A} \inf_{\gamma\in\Gamma^A} \mathbb{E}^v_{\sigma,\gamma}(f)$$
  $$\langle A\rangle(-f)(v) = \sup_{\gamma\in\Gamma^A} \inf_{\sigma\in\Sigma^A} \mathbb{E}^v_{\sigma,\gamma}(-f)$$

  - A strategy σ for Eve is *optimal* from vertex v if the following holds:

  $$\inf_{\gamma\in\Gamma^A} \mathbb{E}^v_{\sigma,\gamma}(f) = \langle E\rangle(f)(v)$$

  The optimal strategies for Adam are defined analogously. ■

2.2.6 Determinacy

Determinacy refers to the property of a $2\frac{1}{2}$ game 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. We present the following theorem which captures the pure memoryless determinacy of the scenarios we have introduced above.

Theorem 2.27 (Pure Memoryless Determinacy[35])
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena.

- For a qualitative objective Φ that is safety, or reachability, or parity, for all vertices $v \in V$, the following holds for the game (G, Φ):

  $$\langle E\rangle(\Phi)(v) + \langle A\rangle(\Pi\backslash\Phi)(v) = 1$$

  Pure memoryless optimal strategies exist for both players from all vertices. If the game is deterministic, either $\langle E\rangle(\Phi)(v) = 1$ or $\langle E\rangle(\Phi)(v) = 0$.

- For a quantitative objective f that is either mean payoff or discounted payoff, for all vertices $v \in V$, the following holds for the game (G, f):

  $$\langle E\rangle(f)(v) + \langle A\rangle(-f)(v) = 0$$

  Pure memoryless optimal strategies exist for both players from all vertices. ■

According to Theorem 2.27, in our context, it suffices to consider only pure memoryless strategies. Thus we obtain the following corollaries for the value function.

Corollary 2.28 (Sufficiency of Pure Memoryless Strategies)
In a stochastic game with qualitative objectives, say (G, Φ), when Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma$ and $\gamma \in \Gamma$ respectively, we obtain a sub-arena $G_{\sigma,\gamma}$. We can reduce the expected outcome $\mathbb{P}^{v_I}_{\sigma,\gamma}$ to reachability probabilities in $G_{\sigma,\gamma}$ in the following way:

- Given a reachability objective $RE(T)$, $\mathbb{P}^{v_I}_{\sigma,\gamma}(RE(T)) = \Pr^{v_I}_{\sigma,\gamma}(\mathit{Reach}(T))$
- Given a safety objective $SA(T)$, $\mathbb{P}^{v_I}_{\sigma,\gamma}(SA(T)) = 1 - \Pr^{v_I}_{\sigma,\gamma}(\mathit{Reach}(T))$
- Given a parity objective $PA(p)$, $\mathbb{P}^{v_I}_{\sigma,\gamma}(PA(p)) = \Pr^{v_I}_{\sigma,\gamma}(\mathit{Reach}(B_E))$, where

$$B_E = \bigcup_{min(p(B)) \text{ is even}} B \in BSCC(\mathcal{M}_{\sigma,\gamma})$$

■

Corollary 2.29 (Sufficiency of Pure Memoryless Strategies)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. Let $RE(T)$ be a reachability objective, and $PA(p)$ be a parity objective. For all vertices $v \in V$, the following holds:

- For the reachability objective $RE(T)$,

$$\begin{aligned}\langle E\rangle(RE(T))(v) &= \sup_{\sigma\in\Sigma^A} \inf_{\gamma\in\Gamma^A} \mathbb{P}^{v}_{\sigma,\gamma}(RE(T)) \\ &= \sup_{\sigma\in\Sigma} \inf_{\gamma\in\Gamma} \mathbb{P}^{v}_{\sigma,\gamma}(RE(T)) \\ &= \sup_{\sigma\in\Sigma} \inf_{\gamma\in\Gamma} \Pr^{v}_{\sigma,\gamma}(\mathit{Reach}(T))\end{aligned}$$

- For the parity objective $PA(p)$,

$$\begin{aligned}\langle E\rangle(PA(p))(v) &= \sup_{\sigma\in\Sigma^A} \inf_{\gamma\in\Gamma^A} \mathbb{P}^{v}_{\sigma,\gamma}(PA(p)) \\ &= \sup_{\sigma\in\Sigma} \inf_{\gamma\in\Gamma} \mathbb{P}^{v}_{\sigma,\gamma}(PA(p)) \\ &= \sup_{\sigma\in\Sigma} \inf_{\gamma\in\Gamma} \Pr^{v}_{\sigma,\gamma}(\mathit{Reach}(B_E))\end{aligned}$$

  where

$$B_E = \bigcup_{min(p(B)) \text{ is even}} B \in BSCC(\mathcal{M}_{\sigma,\gamma})$$

■

As a result, we consider only pure memoryless strategies in the sequel, unless stated otherwise explicitly.

Chapter 3

Terminologies and Notations

In this short chapter, we introduce some necessary terminologies and notations for SGs.

Qualitative, Quantitative, and Strategic Solution. We divide solving stochastic games into three distinct tasks. Given an SG, we have:

- Solving the SG qualitatively is the task of computing all the vertices from which Eve (or Adam) wins with probability 1. This is only applicable to qualitative objectives.
- Solving the SG quantitatively is the task of computing the values of all vertices in the arena. Specifically, if the objective is qualitative, we compute the winning probabilities; if the objective is quantitative, we compute the expected values of the objective function.
- Solving the SG strategically is the task of computing the optimal strategy of Eve (or Adam) for the game.

For the sake of clarity,

- When we talk about the objectives of games, we always say 'an SG with a qualitative objective'. The same applies to 'an SG with a quantitative objective'.
- When we say 'a qualitative SG', it means we target the qualitative solution of this SG. The same applies to 'a quantitative SG' and 'a strategic SG'.

(Almost-sure) Winning Vertices and Regions. We now introduce some notions for *winning with probability* 1, which is only applicable to SGs with qualitative objectives. Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, and Φ be a qualitative objective. We specify an SG (G, Φ).

- For all vertices $v \in V$, we call v an *almost-sure winning vertex* of Eve, if there exists a strategy of Eve such that she wins with probability 1 from v, formally

$\langle E\rangle(\Phi)(v) = 1$.

- We call the set of almost-sure winning vertices of Eve, denoted with $U_\exists$, her *almost-sure winning region*, formally $U_\exists = \{v \in V \mid \langle E\rangle(\Phi)(v) = 1\}$.
- If the arena contains no random vertices, which makes the game deterministic, we call v a *winning vertex* of Eve, and $U_\exists$ her *winning region*.
- The (almost-sure) winning vertices and region of Adam are defined analogously.

Even BSCCs and Cycles. We now introduce some notions for SPGs. Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, and $PA(p)$ be a parity objective. We specify an SPG $(G, PA(p))$.

When Eve and Adam follow pure memoryless strategies $\sigma \in \Sigma$ and $\gamma \in \Gamma$ respectively, we obtain a sub-arena $G_{\sigma,\gamma}$ as Definition 2.19. We recall remark 2.22 for how we implicitly address the Markov chain view $\mathcal{M}_{\sigma,\gamma}$ of the sub-arena $G_{\sigma,\gamma}$.

- 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.
- If the game is deterministic, B is reduced to a cycle, and we call B an *even cycle*.
- An *odd BSCC* and an *odd cycle* are defined analogously.

Projection. Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena. We fix an arbitrary pair of strategies $\sigma, \gamma \in \Sigma \times \Gamma$ to obtain a sub-arena $G_{\sigma,\gamma}$ and its Markov Chain view $\mathcal{M}_{\sigma,\gamma} = (V, \delta, v_I)$. To obtain a part of the sub-arena, which is also a Markov Chain, we introduce the notion of *projection*. Given a set of vertices $T \subseteq V$, if for all $v \in T$, $\delta(v, v') > 0$ implies $v' \in T$, then the projection of $G_{\sigma,\gamma}$ onto T, denoted $G_{\sigma,\gamma} \upharpoonright T$, is well defined. We have:

$$G_{\sigma,\gamma} \upharpoonright T = (T, \delta_T, v_i)$$

where the new transition function $\delta_T : T \to \mathbb{D}(T)$ is defined as: for all $t \in T$, $\delta_T(t) = \delta(t)$. We take an arbitrary vertex $v_i \in T$ as the starting vertex for the reasons stated in Remark 2.22.

A Lemma for Almost-Sure Winning Regions. Finally, we present a lemma concerning the properties of almost-sure winning regions.

Lemma 3.1 ((Almost-Sure) Winning Regions)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, and $PA(p)$ be a parity objective. We specify an SPG $(G, PA(p))$. Let $U_\exists$ be the almost-sure winning region of Eve, and $U_\forall$ be the almost-sure winning region of Adam in $(G, PA(p))$. The following statements hold:

1. For all Adam's vertices $v \in V_\forall \cap U_\exists$, for all $v' \in V$, $(v, v') \in E$ implies $v' \in U_\exists$.

2. For all random vertices $v \in V_R \cap U_\exists$, for all $v' \in V$, $\Delta(v, v') > 0$ implies $v' \in U_\exists$.
3. There exists an optimal strategy $\sigma^* \in \Sigma$ of Eve, such that for all strategies $\gamma \in \Gamma$ of Adam, we have:
   - $G_{\sigma^*,\gamma} \upharpoonright U_\exists$ is well defined.
   - A play starting from $v \in U_\exists$ does not leave $G_{\sigma^*,\gamma} \upharpoonright U_\exists$.
   - All BSCCs are even in $G_{\sigma^*,\gamma} \upharpoonright U_\exists$.
4. For all optimal strategies $\sigma^* \in \Sigma$ of Eve, for all strategies $\gamma \in \Gamma$ of Adam, if a BSCC C is odd in $G_{\sigma^*,\gamma}$, then $C \subseteq U_\forall$. ■

Proof. We prove the statements one by one.

1. We assume there exists $v \in V_\forall \cap U_\exists$ and $v' \in V$, such that $(v, v') \in E$ and $v' \notin U_\exists$. As long as Adam chooses v' as the successor of v and plays optimally from v, Eve does not win with probability 1 from v, which contradicts the assumptions that $U_\exists$ is the almost-sure winning region of Eve.
2. Similarly, if random vertex $v \in V_R \cap U_\exists$ has a successor outside $U_\exists$, then Eve reaches a non-almost-sure winning vertex with a positive probability from v. This implies that v is not almost-sure winning for Eve, which contradicts the assumptions that $U_\exists$ is the almost-sure winning region of Eve.
3. It follows from 1 and 2 that $G_{\sigma^*,\gamma} \upharpoonright U_\exists$ is well defined, and a play starting from $v \in U_\exists$ does not leave $G_{\sigma^*,\gamma} \upharpoonright U_\exists$. If there exists an odd BSCC C in $G_{\sigma^*,\gamma} \upharpoonright U_\exists$, then from all vertices $v \in C \subseteq U_\exists$, Eve wins with probability 0. This contradicts the assumption that $U_\exists$ is the almost-sure winning region of Eve.
4. When Eve follows optimal strategy σ^* and Adam follows strategy γ, Adam wins with probability 1 from all vertices in C, following from C being odd in $G_{\sigma^*,\gamma}$. It indicates that Eve wins with probability 0 from C even if she plays optimally. Thus Adam wins with probability 1 from C as long as he follows strategy γ. □

Chapter 4

Reducing Qualitative Parity Games to Quantitative SSGs

The contents of this chapter are given in Figure 4.1:

- We first revisit the reductions from qualitative SPGs to DPGs, then to quantitative SSGs.
- We then proposed a direct reduction from qualitative SPGs to quantitative SSGs.

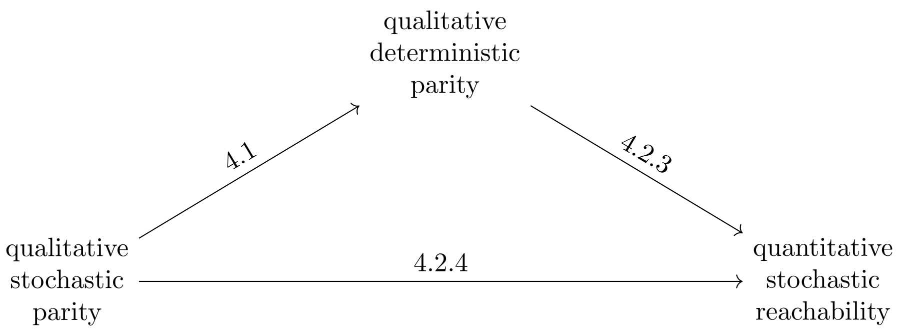

Figure 4.1: Reductions from qualitative parity to quantitative reachability

4.1 Reducing Qualitative SPGs to DPGs

Solving an SPG qualitatively, that is, determining the set of almost-sure winning vertices of either player, can be reduced to solving a DPG [13]. In this section, we revisit this reduction.

4.1.1 Gadget Construction

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, $p : V \to \mathbb{N}$ be a priority function and $(G, PA(p))$ be an SPG. We construct a DPG $(\overline{G}, PA(\overline{p}))$, where $\overline{G}$ is an arena without random vertices, and $\overline{p}$ is a slightly changed priority function. They are specified as follows.

For all vertices $v \in V_\exists \cup V_\forall$ that are not random, we make a copy and rename it as $\overline{v}$ in $\overline{G}$, and the priority of $\overline{v}$ remains the same, formally $\overline{p}(\overline{v}) = p(v)$. For all edges $(u, v) \in E$ where $u, v \in V_\exists \cup V_\forall$, we also make a copy, which is denoted with $(\overline{u}, \overline{v})$ in $\overline{G}$.

For all random vertices $v \in V_R$, we substitute it with a tree-shaped gadget as below in Figure 4.2.

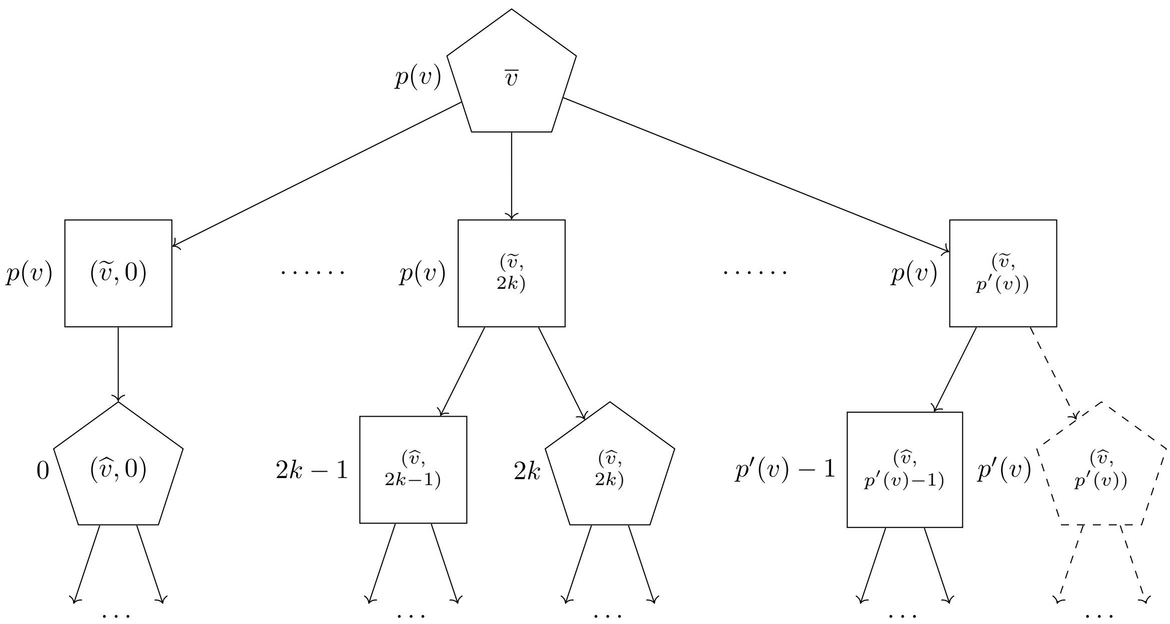

Figure 4.2: The gadget for reducing qualitative SPGs to DPGs

Specifically, we first introduce a function p'. For all random vertices $v \in V_R$, we let $p'(v) = p(v)$ if $p(v)$ is even, and $p'(v) = p(v) + 1$ if $p(v)$ is odd. We construct the tree-shaped gadget substituting v as follows:

- It starts with a root vertex $\overline{v}$ which belongs to Adam, and $\overline{p}(\overline{v}) = p(v)$. For all edges $(u, v) \in E$ where $u \in V_\exists \cup V_\forall$, there is an edge $(\overline{u}, \overline{v})$ in the gadget. In other words, we first 'copy' all edges in G that go from non-random vertices to v.

- There are $\frac{p'(v)}{2} + 1$ vertices in the second layer. For all integers $k \in [0, \frac{p'(v)}{2}]$, $\overline{v}$ has a successor $(\widetilde{v}, 2k)$ that belongs to Eve. The priority of $(\widetilde{v}, 2k)$ is the same as v, formally $\overline{p}((\widetilde{v}, 2k)) = p(v)$.

- In the third layer we introduce $|p(v)|+1$ vertices, namely $(\widehat{v}, 0), (\widehat{v}, 1), \cdots, (\widehat{v}, p(v))$, and for all integers $i \in [0, p(v)]$, $\overline{p}((\widehat{v}, i)) = i$.

  - The vertex $(\widetilde{v}, 0)$ has one successor $(\widehat{v}, 0)$ that belongs to Adam.

  - For all vertices $(\widetilde{v}, 2k)$ where $0 < k < \frac{p'(v)}{2}$, $(\widetilde{v}, 2k)$ has two successors: $(\widehat{v}, 2k-1)$ that belongs to Eve, $(\widehat{v}, 2k)$ that belongs to Adam.

  - The vertex $(\widetilde{v}, p'(v))$ has one successor $(\widehat{v}, p'(v) - 1)$ that belongs to Eve. If $p'(v) = p(v)$, it has another successor $(\widehat{v}, p(v))$ that belongs to Adam. Either way, the largest priority in this gadget is still $p(v)$.

- For all vertices $(\widehat{v}, i)$ where integer $i \in [0, p(v)]$, this vertex has a successor $\overline{w}$ if there exists $w \in V$ such that $(v, w) \in E$. In other words, we 'copy' all edges in G that go from v to other vertices. Note that all edges between random vertices in G are also 'copied' into $\overline{G}$ in this way.

We give the intuition of this gadget in the following remark.

Remark 4.1 (Either Priority or Right of Choice)
Note that whenever there is a random vertex $v \in V_R$ in G, correspondingly Adam chooses a successor of the root vertex $\overline{v}$ in $\overline{G}$, which is $(\widetilde{v}, 2k)$ for some integer $k \in [0, \frac{p'(v)}{2}]$. Then Eve has to pick at each vertex $(\widetilde{v}, 2k)$ one of its (at most) two successors, namely $(\widehat{v}, 2k-1)$ and $(\widehat{v}, 2k)$. If Eve chooses $(\widehat{v}, 2k-1)$ when $k > 0$, then an odd priority $2k-1$ might be introduced to the path. If Eve chooses $(\widehat{v}, 2k)$ given it is available, Adam gets to choose the successor at his vertex $(\widehat{v}, 2k)$. ■

4.1.2 Correctness

In the following we prove the correctness of the reduction, that is, for all $v \in V$, v is an almost-sure winning vertex of Eve in $(G, PA(p))$ if and only if $\overline{v}$ is a winning vertex of Eve in $(\overline{G}, PA(\overline{p}))$. For all sets of vertices $D \subseteq V$ in G, we denote its counterpart in $\overline{G}$ with $\overline{D}$. Formally,

$$\begin{aligned}\overline{D} = & \{\overline{v} \mid v \in D\} \cup \\ & \bigcup_{v \in D \cap V_R} (\{(\widetilde{v}, 2k) \mid k \in \mathbb{N} \cap [0, \frac{p'(v)}{2}]\} \\ & \cup \{(\widehat{v}, i) \mid i \in \mathbb{N} \cap [0, p(v)]\})\end{aligned}$$

1. We first show that if $\overline{v}$ is a winning vertex of Eve in the DPG, then v is an almost-sure winning vertex of Eve in the SPG.

   We take a set of states $U_\exists$ from G such that $\overline{U}_\exists$ is the winning region of Eve in $\overline{G}$. It follows from Lemma 3.1 that there exists an optimal strategy $\overline{\sigma}$ of Eve, such that for all strategies $\overline{\gamma}$ of Adam, all cycles in $\overline{G}_{\overline{\sigma},\overline{\gamma}} \upharpoonright \overline{U}_\exists$ are even. Given Eve's optimal strategy $\overline{\sigma}$ in $\overline{G}$, we define a strategy σ in G, where for all $v \in V_\exists$, $\sigma(v) = u$ if $\overline{\sigma}(\overline{v}) = \overline{u}$. Let Eve follow strategy σ in G.

   We prove the following claims:

   (1) For all random vertices $v \in V_R \cap U_\exists$, v has only successors in $U_\exists$.

   (2) For all Adams's vertices $v \in V_\forall \cap U_\exists$, v has only successors in $U_\exists$.

   (3) For all strategies γ of Adam, all BSCCs in $G_{\sigma,\gamma} \upharpoonright U_\exists$ are even.

   which are sufficient to prove that Eve wins with probability 1 from all vertices in $U_\exists$.

   (1) We also make the proof by contradiction.

   In G, we assume there exists a random vertex $v \in V_R \cap U_\exists$, such that v has a successor w outside $U_\exists$. Formally, there exists $w \in V$ such that $w \notin U_\exists$ and $(v, w) \in E$. As a result, Adam has a strategy in $\overline{G}$ that chooses $(\widetilde{v}, 0)$ at $\overline{v}$ and chooses $\overline{w} \notin \overline{U}_\exists$ at $(\widehat{v}, 0)$, which yields a path from $\overline{v}$ to $\overline{w}$. Note that $\overline{w} \in \overline{V} \backslash \overline{U}_\exists = \overline{U}_\forall$ according to Theorem 2.27, indicating that Adam can win from vertex $\overline{w}$ and thus $\overline{v}$. This contradicts the assumptions that $\overline{U}_\exists$ is the winning region of Eve, and that $\overline{\sigma}$ is a winning strategy of Eve.

   (2) The claim follows trivially with the similar idea as (1).

   (3) It follows from (1) that for all $\gamma \in \Gamma$, $G_{\sigma,\gamma} \upharpoonright U_\exists$ is well defined. We make the proof by contradiction, and give a running example in the meantime. Assume there exists a strategy γ' of Adam such that $G_{\sigma,\gamma'} \upharpoonright U_\exists$ contains an odd BSCC, denoted with C. Let the lowest priority of C be $2r - 1$ for some $r \in \mathbb{N}_{>0}$, and let $v_m \in C$ be an arbitrary vertex with priority $2r - 1$. Below is an example of C where $r = 1$, and we leave out the probabilities of transitions (v, u) and (v, w), which are both positive.

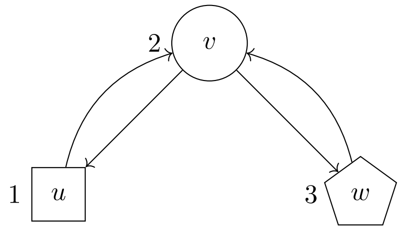

Figure 4.3: BSCC C of the running example

Now we look at $\overline{G}$. We construct a strategy $\overline{\gamma}'$ of Adam defined on $\overline{C}$, which goes as follows.

For all $v \in C \cap V_\forall$, we let $\overline{\gamma}'(\overline{v}) = \overline{u}$ if $\gamma'(v) = u$. Intuitively, Adam makes the 'same' move at $\overline{v}$ if v is his vertex in G. Note that $\overline{\sigma}$ is already fixed in $\overline{G}$, and thus the choices on Adam's vertices in newly introduced gadgets are left open. We denote such a directed graph with $\overline{C}'$.

We present $\overline{C}'$ of the running example in the following figure. The open choices of Adam are represented with dotted arrows, and we make a case analysis of the fixed choices of Eve, represented with dashed arrows.

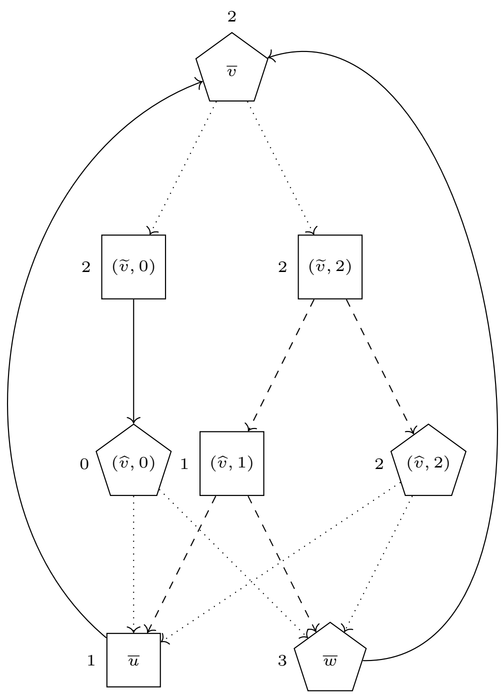

Figure 4.4: Directed graph $\overline{C}'$ of the running example

We introduce a function $d : \overline{C} \to \mathbb{N}$ such that for all $\overline{v} \in \overline{C}$, $d(\overline{v})$ represents the distance between $\overline{v}$ and $\overline{v}_m$ in $\overline{C}'$. We recall Remark 4.1 that Eve can either control the priority or choose the successor. Below is how $\overline{\gamma}'$ behaves in the gadgets:

- If at some vertex $(\widetilde{v}, 2k)$ where integer $k \in [0, r]$, Eve chooses $(\widehat{v}, 2k-1)$, then at $\overline{v}$ Adam chooses $(\widetilde{v}, 2k)$. As a result, the priority $2k-1 \leq 2r-1$ is introduced, although Adam cannot choose the successor of $(\widehat{v}, 2k-1)$.
- Otherwise it implies that Eve chooses $(\widehat{v}, 2r)$ at $(\widetilde{v}, 2r)$. Then at $\overline{v}$ Adam chooses $(\widetilde{v}, 2r)$, and at $(\widehat{v}, 2r)$ Adam choose a successor $\overline{v}'$ such that $d(\overline{v}') < d((\widehat{v}, 2r))$. Note that such $\overline{v}'$ always exists, since $(\widehat{v}, 2r)$ can always reach $\overline{v}_m$ in $\overline{C}'$.

The following figure depicts the running example corresponding to the two cases, where the transitions in dashed arrows represent 'not care'.

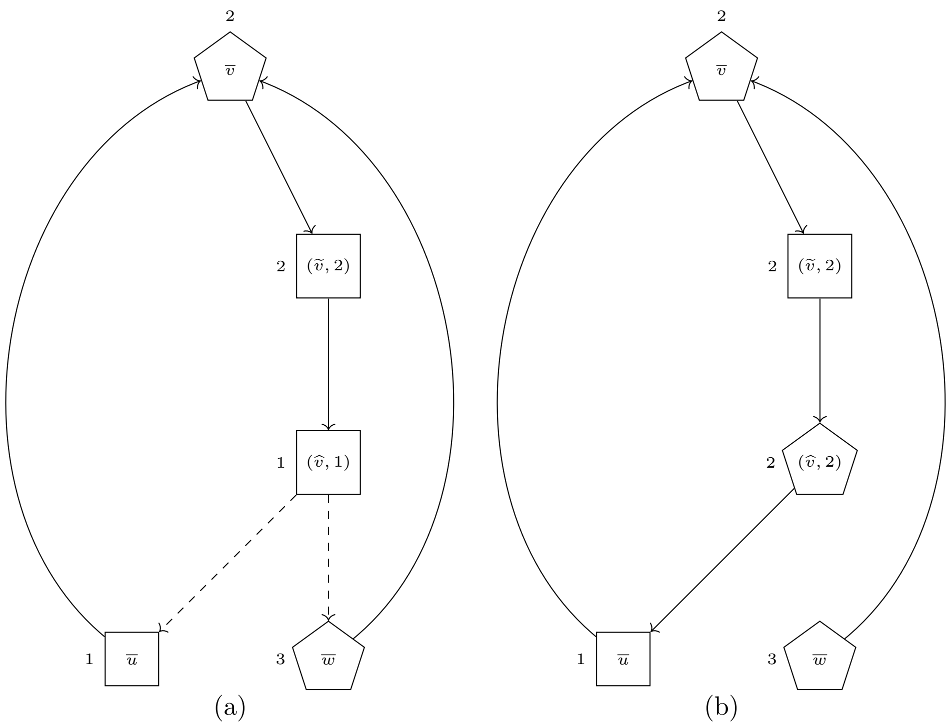

Figure 4.5: Case analysis for strategy inside gadgets of the running example

The result of Adam's strategy $\overline{\gamma}'$ is: whenever in G there is a random vertex $v \in C \cap V_R$, correspondingly in $\overline{G}$ either an odd priority no larger than $2r-1$ is introduced, or Adam can shorten the distance to $\overline{v}_m$, whose priority is

$2r-1$. We now look at the resulting $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{C}$.

If C contains no random vertices, then $G_{\sigma,\gamma'} \restriction C$ is a cycle and it indicates that $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{C}$ is the same cycle up to vertices renaming. Therefore the lowest priority on $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{C}$ is $2r-1$.

Otherwise, C contains at least one random vertex. We look at a cycle in $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{C}$ containing $t \geq 1$ vertices from $\overline{C} \cap \{\overline{v} \mid v \in V_R\}$, which are the root vertices of the gadgets. We arbitrarily fix one as $\overline{v}_0$, proceed from $\overline{v}_0$ along the cycle, and denote the others with $\overline{v}_1, \overline{v}_2, \cdots, \overline{v}_{t-1}$. We distinguish between two cases:

- If an odd priority no larger than $2r-1$ is introduced, the cycle is odd.
- Otherwise for all integers $i \in [0, t-1]$, at $(\widehat{v}_i, 2r)$ Adam has chosen a successor $\overline{v}'_i$ such that $d(\overline{v}'_i) < d((\widehat{v}_i, 2r))$. Thus $\overline{v}_m$ is guaranteed to be on the cycle. If this is not the case, we can find a contradiction as follows:
  - For all integers $i \in [0, t-1]$, it holds that $d(\overline{v}_i) > d((\widehat{v}_i, 2r))$ because of the tree-shaped structure of the gadgets, and thus $d(\overline{v}_i) > d(\overline{v}'_i)$.
  - For all integers $i \in [0, t-1]$, it holds that $d(\overline{v}'_i) > d(\overline{v}_{i+1})$, which follows from the fact that there is a unique path in $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{C}$ from $\overline{v}'_i$ to $\overline{v}_{i+1}$. Note that $d(\overline{v}'_{t-1}) > d(\overline{v}_0)$ also holds since it is a cycle.

  Thus we obtain $d(\overline{v}_i) > d(\overline{v}_{i+1})$ for all integers $i \in [0, t-1]$, and $d(\overline{v}_{t-1}) > d(\overline{v}_0)$. Combining them, we can obtain $d(\overline{v}_1) > d(\overline{v}_1)$ and the contradiction arises.

As a result, we can conclude that $\overline{G}_{\overline{\sigma},\overline{\gamma}'} \restriction \overline{U}_\exists$ contains an odd cycle, which contradicts the assumptions that $\overline{U}_\exists$ is the winning region of Eve, and that $\overline{\sigma}$ is a winning strategy of Eve.

2. For the other direction, we show that if $\overline{v}$ is not a winning vertex of Eve in the DPG, which indicates that $\overline{v}$ is a winning vertex of Adam, then v is not an almost-sure winning vertex of Eve in the SPG.

   We take a set of states $U_\forall$ from G such that $\overline{U}_\forall$ is the winning region of Adam in $\overline{G}$. It follows from Lemma 3.1 that there exists an optimal strategy $\overline{\gamma}$ of Adam, such that for all strategies $\overline{\sigma}$ of Eve, all cycles in $\overline{G}_{\overline{\sigma},\overline{\gamma}} \restriction \overline{U}_\forall$ are odd. Given Adam's optimal strategy $\overline{\gamma}$ in $\overline{G}$, we define a strategy γ in G, where for all $v \in V_\forall$, $\gamma(v) = u$ if $\overline{\gamma}(\overline{v}) = \overline{u}$. Let Adam follow strategy γ in G.

   We claim that by playing strategy γ, Adam wins with a positive probability from all vertices in $U_\forall$. Due to the duality with the previous proof, we refer to [13] for further details.

4.2 Reducing Qualitative SPGs Directly to Quantitative SSGs

In this section, we present:

- A reduction from solving DPGs to solving SSGs quantitatively, which was proposed in [7].
- Our extension of this reduction, which yields a direct reduction from solving SPGs qualitatively to solving SSGs quantitatively.

4.2.1 Gadget Construction

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\text{win}}))$, and the stochastic arena $\overline{G}$ is given as follows in Figure 4.6. Note that when there is only one target vertex, we write $RE(\{v\})$ simply as $RE(v)$.

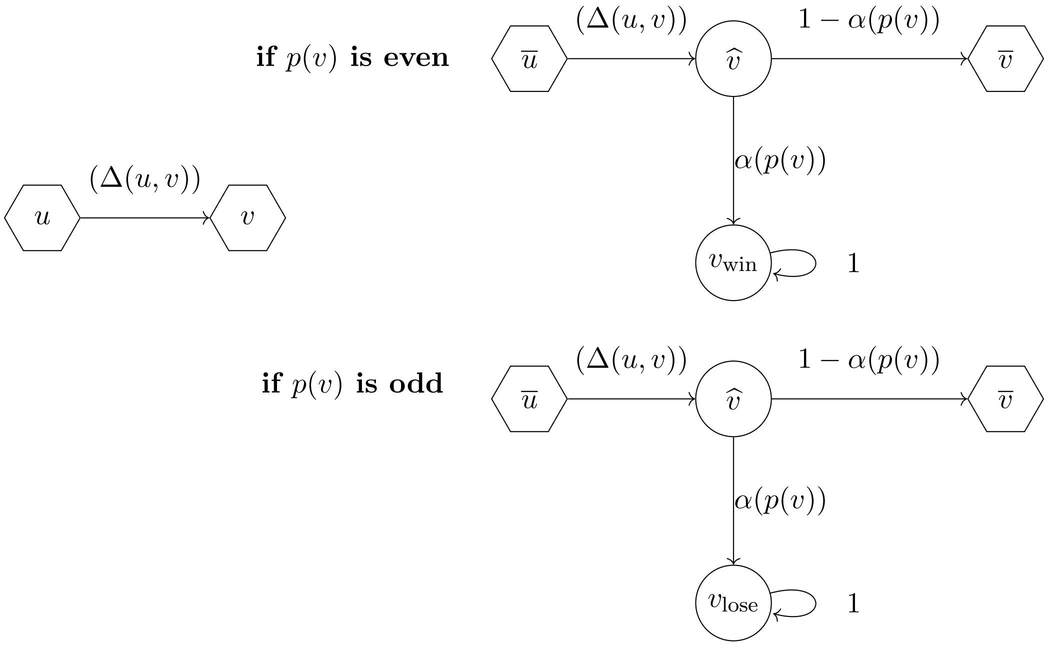

Figure 4.6: The gadget for reducing SPGs to SSGs

We first introduce a few notations for the newly constructed arena $\overline{G}$. Given a set of vertices $U \subseteq V$, we define $\overline{U}$, $\widehat{U}$ and $\widetilde{U}$ as follows:

$$\begin{aligned}\overline{U} &= \{\overline{v} \mid v \in U\}\\ \widehat{U} &= \{\widehat{v} \mid v \in U\}\\ \widetilde{U} &= \overline{U} \uplus \widehat{U}\end{aligned}$$

The newly constructed arena $\overline{G}$ is formally written as:

$$\overline{G} = ((\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{E}), (\overline{V}_{\exists}, \overline{V}_{\forall}, \overline{V}_R \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}), \overline{\Delta})$$

where the new edge set $\overline{E}$ is written as:

$$\begin{aligned}\overline{E} = \; & \{(\overline{u}, \widehat{v}) \mid (u,v) \in E\} \uplus \\ & \{(\widehat{v}, \overline{v}), (\widehat{v}, v_{\text{win}}) \mid v \in V,\ p(v) \text{ is even}\} \uplus \\ & \{(\widehat{v}, \overline{v}), (\widehat{v}, v_{\text{lose}}) \mid v \in V,\ p(v) \text{ is odd}\} \uplus \\ & \{(v_{\text{win}}, v_{\text{win}}), (v_{\text{lose}}, v_{\text{lose}})\}\end{aligned}$$

To define the new transition function $\overline{\Delta}$, we first introduce a function $\alpha : \mathbb{N} \to [0,1]$. It is the probability of entering the winning or losing sink before visiting a vertex with even or odd priority respectively. Then we define the new transition function $\overline{\Delta} : \overline{V}_R \to \mathbb{D}(\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\})$ as follows:

- For all vertices $\overline{u} \in \overline{V}_R$, for all edges $(\overline{u}, \widehat{v}) \in E$, we have $\overline{\Delta}(\overline{u}, \widehat{v}) = \Delta(u,v)$.
- For all vertices $\widehat{v} \in \widehat{V}$ where $p(v)$ is even, we have $\overline{\Delta}(\widehat{v}, \overline{v}) = 1 - \alpha(p(v))$ and $\overline{\Delta}(\widehat{v}, v_{\text{win}}) = \alpha(p(v))$.
- For all vertices $\widehat{v} \in \widehat{V}$ where $p(v)$ is odd, we have $\overline{\Delta}(\widehat{v}, \overline{v}) = 1 - \alpha(p(v))$ and $\overline{\Delta}(\widehat{v}, v_{\text{lose}}) = \alpha(p(v))$.
- The sinks have only self-loops, namely $\overline{\Delta}(v_{\text{win}}, v_{\text{win}}) = 1$ and $\overline{\Delta}(v_{\text{lose}}, v_{\text{lose}}) = 1$.

Remark 4.2 (Intuition of the Gadget)
The intuition of this gadget is that we reward Eve with a small but positive chance to win the game in $\overline{G}$ whenever she visits a vertex with even priority. Naturally, vertices with odd priority are favorable for Adam. Also, small priorities are more rewarding than larger ones, so we use a monotonically decreasing function α. ■

4.2.2 Notations for the New Arena

Given a stochastic arena $G = ((V,E), (V_{\exists}, V_{\forall}, V_R), \Delta)$, we construct a new stochastic arena $\overline{G}$ as specified previously in Section 4.2.1. We have:

$$\overline{G} = ((\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{E}), (\overline{V}_{\exists}, \overline{V}_{\forall}, \overline{V}_R \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}), \overline{\Delta})$$

Since all newly introduced vertices, namely $\widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$, are random vertices, a strategy of either player in G is still valid in $\overline{G}$, and thus we do not distinguish between strategies in G and $\overline{G}$.

Fixing a pair of strategies of Eve and Adam, denoted with $\sigma, \gamma \in \Sigma \times \Gamma$, in G, we obtain a sub-arena $G_{\sigma,\gamma}$. Similarly, we obtain a sub-arena $\overline{G}_{\sigma,\gamma}$. We introduce some notations for $\overline{G}_{\sigma,\gamma}$.

- If U is a BSCC in $G_{\sigma,\gamma}$, we call $\widetilde{U}$ an *rBSCC* in $\overline{G}_{\sigma,\gamma}$. If U is an even or odd BSCC, we call $\widetilde{U}$ an *even rBSCC* or an *odd rBSCC* respectively.
- $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})$ denotes the probability for a play starting from $\overline{v} \in \overline{V}$ to reach an rBSCC in $\overline{G}$, formally written as:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) = \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(\{\overline{u} \mid \overline{u} \text{ is in an rBSCC}\}))$$

- $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven})$ denotes the probability for a play starting from $\overline{v} \in \overline{V}$ to reach an even rBSCC, formally written as:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven}) = \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(\{\overline{u} \mid \overline{u} \text{ is in an even rBSCC}\}))$$

  Naturally, we define $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterOdd})$ in a similar way.
- $\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven})$ denotes the minimum probability for a play starting from a vertex $\overline{v} \in \overline{V}$ in an even rBSCC to reach the winning sink. Formally,

$$\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) = \min_{\overline{v} \text{ is in an even rBSCC}} \{\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))\}$$

- $\overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd})$ denotes the maximum probability for a play starting from a vertex $\overline{v} \in \overline{V}$ in an odd rBSCC to reach the winning sink. Formally,

$$\overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) = \max_{\overline{v} \text{ is in an odd rBSCC}} \{\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))\}$$

4.2.3 From DPGs: A Special Case

A reduction from solving DPGs to solving SSGs quantitatively was proposed in [7], using the gadgets specified in Section 4.2.1. We sketch this reduction as follows.

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena where $V_R = \emptyset$, $p : V \to \mathbb{N}$ be a priority function, and $(G, \mathit{PA}(p))$ be a DPG. We construct an SSG $(\overline{G}, \mathit{RE}(v_{\text{win}}))$ as specified previously in Section 4.2.1, and we have:

$$\overline{G} = ((\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{E}), (\overline{V}_\exists, \overline{V}_\forall, \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}), \overline{\Delta})$$

Note that since $\overline{V}_R = \emptyset$, the set of random vertices in $\overline{G}$ is $\widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$.

We fix a pair of strategies $\sigma, \gamma \in \Sigma \times \Gamma$ in G. Since there are no random vertices in G, BSCCs in $G_{\sigma,\gamma}$ are reduced to cycles, and a play is reduced to a simple path leading to an infinitely repeated cycle. Correspondingly, in $\overline{G}_{\sigma,\gamma}$, rBSCCs are reduced to cycles with random vertices from $\widehat{V}$ leading to sinks v_{win} and v_{lose}. We present the following Figure 4.7 for some intuition.

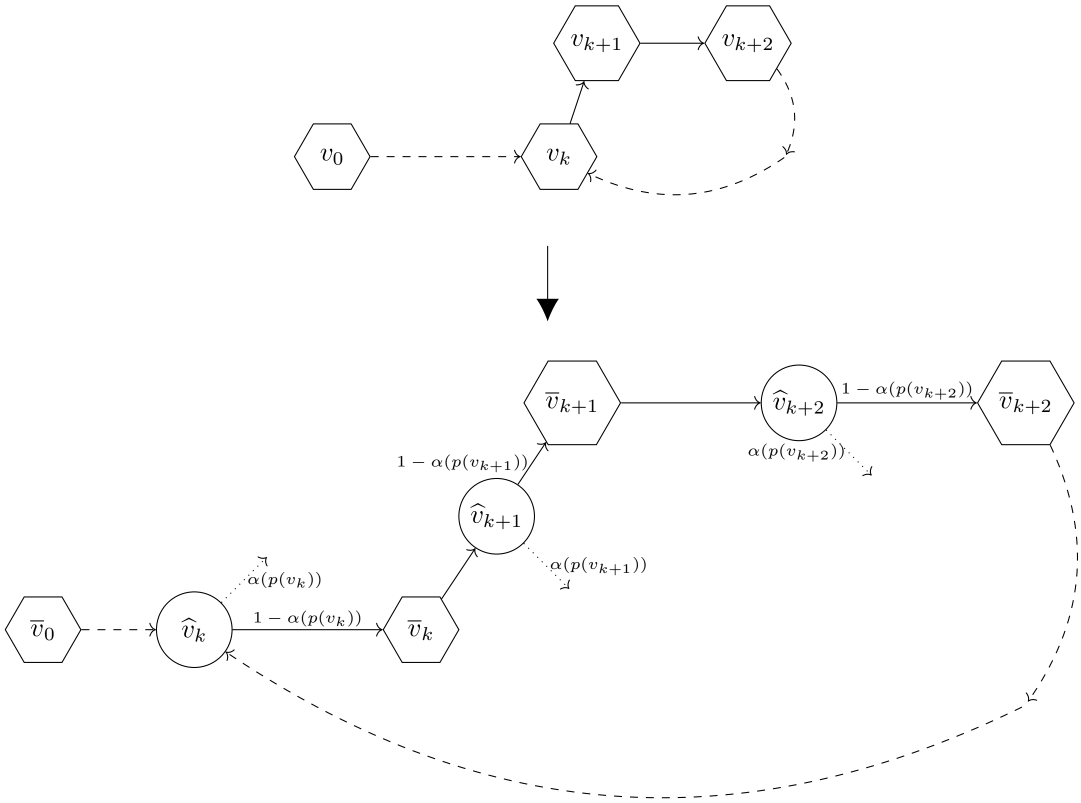

Figure 4.7: Deterministic scenario of a play in $G_{\sigma,\gamma}$ and $\overline{G}_{\sigma,\gamma}$

Note that we leave out the winning and losing sinks in Figure 4.7, and we represent the transitions into sinks with dotted arrows for visual neatness.

As a result, the probabilities in $\overline{G}_{\sigma,\gamma}$ are reduced to the following:

- $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})$ is the probability for a play starting from $\overline{v} \in \overline{V}$ to reach a cycle.
- $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven})$ is the probability for a play starting from $\overline{v} \in \overline{V}$ to reach an even cycle. Naturally, we define $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterOdd})$ in a similar way.
- $\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven})$ is the minimum probability for a play starting from a vertex

$\overline{v} \in \overline{V}$ on an even cycle to reach the winning sink.

- $\overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd})$ is the maximum probability for a play starting from a vertex $\overline{v} \in \overline{V}$ on an odd cycle to reach the losing sink.

In the following, we present the theorem that states the correctness of this reduction and give a proof.

Theorem 4.3 (Reducing DPGs to quantitative SSGs)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena with n vertices where $V_R = \emptyset$, $p : V \to \mathbb{N}$ be a priority function, and $(G, \mathit{PA}(p))$ be a DPG. We construct an SSG $(\overline{G}, \mathit{RE}(v_{\text{win}}))$ as in Section 4.2.1.

Let $J^{>\overline{v}}_{even} = \{\overline{u} \mid p(u) \text{ is even}, p(u) > p(v)\}$ and $J^{>\overline{v}}_{odd} = \{\overline{u} \mid p(u) \text{ is odd}, p(u) > p(v)\}$. If the following conditions hold:

(A_0) $\alpha(0) \leq \frac{1}{6n}$

(A_1) $\forall\, \overline{v} \in \overline{V},\ \sum_{\overline{u} \in J^{>\overline{v}}_{odd}} \alpha(p(u)) \leq \frac{5}{9}\alpha(p(v))$

(A_2) $\forall\, \overline{v} \in \overline{V},\ \sum_{\overline{u} \in J^{>\overline{v}}_{even}} \alpha(p(u)) \leq \frac{5}{9}\alpha(p(v))$

then for all vertices $v \in V$, $\langle E\rangle(\mathit{PA}(p))(v) = 1$ in the DPG $(G, \mathit{PA}(p))$ if and only if $\langle E\rangle(\mathit{RE}(v_{\text{win}}))(\overline{v}) > \frac{1}{2}$ in the SSG $(\overline{G}, \mathit{RE}(v_{\text{win}}))$. ■

Remark 4.4 (A Specific Instance of α)
When we fix α as $\alpha(k) = (\frac{1}{2})^{(2n+3)(k+1)}$, the conditions (A_0), (A_1) and (A_2) are trivially satisfied. ■

Proof. We assume conditions (A_0), (A_1) and (A_2) hold. We fix an arbitrary pair of strategies $\sigma, \gamma \in \Sigma \times \Gamma$.

We first show that condition (A_1) implies:

$$\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) \geq \frac{3}{5}$$

We take an arbitrary even cycle from $\overline{G}_{\sigma,\gamma}$, and denote it with $c = (\widehat{v}_1\overline{v}_1 \cdots \widehat{v}_m\overline{v}_m)^\omega$ for some integer $m > 0$. Without loss of generality, we start from $\widehat{v}_1$, and we assume that v_k has the smallest priority among $\{v_i \mid i = 1, 2, \cdots, m\}$. For simplicity, for all $i = 1, 2, \cdots, m$, we denote with s_i the probability of starting from $\widehat{v}_1$ and reaching either sink via the edge $(\widehat{v}_i, v_{\text{win}})$ or $(\widehat{v}_i, v_{\text{lose}})$ within the first traversal of the cycle. Formally, for all $i = 1, 2, \cdots, m$, we define s_i as follows:

$$\begin{aligned} s_i &= \alpha(p(v_i)) \cdot \prod_{j=1}^{i-1}(1 - \alpha(p(v_j))) \\ &\leq \alpha(p(v_i)) \end{aligned} \tag{4.1}$$

It follows that the probability of reaching neither sink through a traversal of the whole cycle can be written as $1 - \sum_{i=1}^{m} s_i$.

We denote with p_k the probability of starting from $\widehat{v}_1$ and reaching v_{win} via the edge $(\widehat{v}_k, v_{\text{win}})$, and obtain the following:

$$\begin{aligned}
p_k &= s_k \sum_{l=0}^{\infty}(1 - \sum_{i=1}^{m} s_i)^l \\
&= \frac{s_k}{\sum_{i=1}^{m} s_i} \\
&\geq \frac{s_k}{s_k + \sum_{i=1}^{k-1} \alpha(p(v_i)) + \sum_{i=k+1}^{m} \alpha(p(v_i))} && \textit{follows from (4.1)} \\
&\geq \frac{s_k}{s_k + \sum_{\overline{u} \in J_{odd}^{>\overline{v}_k}} \alpha(p(u))} \\
&= \frac{s_k}{s_k + \frac{5}{9}\alpha(p(v_k))} && \textit{follows from } (A_1) \\
&= \frac{\alpha(p(v_k)) \cdot \prod_{j=1}^{k-1}(1 - \alpha(p(v_j)))}{\alpha(p(v_k)) \cdot \prod_{j=1}^{k-1}(1 - \alpha(p(v_j))) + \frac{5}{9}\alpha(p(v_k))} \\
&= \frac{\prod_{j=1}^{k-1}(1 - \alpha(p(v_j)))}{\prod_{j=1}^{k-1}(1 - \alpha(p(v_j))) + \frac{5}{9}} \\
&\geq \frac{(1 - \alpha(0))^n}{(1 - \alpha(0))^n + \frac{5}{9}} \\
&\geq \frac{1 - n\alpha(0)}{1 - n\alpha(0) + \frac{5}{9}} && \textit{follows from Bernoulli Inequality} \\
&\geq \frac{1 - n \cdot \frac{1}{6n}}{1 - n \cdot \frac{1}{6n} + \frac{5}{9}} && \textit{follows from } (A_0) \\
&= \frac{3}{5}
\end{aligned}$$

Note that c is an arbitrary even cycle and k is an arbitrary position on c. It follows immediately that:

$$\overline{\Pr}_{\sigma,\gamma}^{\min}(\mathit{winEven}) \geq \frac{3}{5} \qquad (A_1')$$

In a similar way, we obtain that (A_2) implies:

$$\overline{\Pr}_{\sigma,\gamma}^{\max}(\mathit{winOdd}) \leq \frac{2}{5} \qquad (A_2')$$

What is more, we obtain for $\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(\mathit{crossPath})$ the following:

$$\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(\mathit{crossPath}) \geq 1 - \sum_{\overline{v} \in \overline{V}} \alpha(p(v)) > 1 - n\alpha(0) \geq \frac{5}{6} \qquad (A_0')$$

Now we show that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) = 1$ in the DPG $(G, PA(p))$ if and only if $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) > \frac{1}{2}$ in the SSG $(\overline{G}, RE(v_{\text{win}}))$. Let $U_\exists$ and $U_\forall$ be the winning regions of Eve and Adam in $(G, PA(p))$ respectively. It follows from Theorem 2.27 that $U_\exists \uplus U_\forall = V$.

- We first prove that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) = 1$ in the DPG implies $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) > \frac{1}{2}$ in the SSG.

  We look at the DPG $(G, PA(p))$. According to Lemma 3.1, Eve has an optimal strategy $\sigma^* \in \Sigma$, such that for all strategies $\gamma \in \Gamma$ of Adam, all cycles in $G_{\sigma^*,\gamma} \restriction U_\exists$ are even. We take an arbitrary winning vertex $v \in U_\exists$ of Eve, or equivalently an arbitrary vertex $v \in V$ where $\langle E\rangle(PA(p))(v) = 1$. According to Lemma 3.1, a play starting from v does not leave $G_{\sigma^*,\gamma} \restriction U_\exists$.

  We now look at the SSG $(\overline{G}, RE(v_{\text{win}}))$. We define $\widetilde{U_\exists} = \overline{U_\exists} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$ such that $\overline{G}_{\sigma^*,\gamma} \restriction \widetilde{U_\exists}$ is well defined. It follows from the construction of $\overline{G}$ that all rBSCCs in $\overline{G}_{\sigma^*,\gamma} \restriction \widetilde{U_\exists}$ are even, and a play starting from $\overline{v}$ does not leave $\overline{G}_{\sigma^*,\gamma} \restriction \widetilde{U_\exists}$. Thus it follows from (A'_0) and (A'_1) that:

  $$\begin{aligned}\overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{Reach}(v_{\text{win}})) &\geq \overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{crossPath}) \cdot \overline{\Pr}^{\min}_{\sigma^*,\gamma}(\mathit{winEven}) \\ &> \frac{5}{6} \cdot \frac{3}{5} \\ &= \frac{1}{2}\end{aligned}$$

  We obtain for the value function $\langle E\rangle(RE(v_{\text{win}}))(\overline{v})$ of the SSG that:

  $$\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) = \inf_{\gamma \in \Gamma} \overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{Reach}(v_{\text{win}})) > \frac{1}{2}$$

- For the other direction, we prove that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) \neq 1$ in the DPG implies $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) \leq \frac{1}{2}$ in the SSG.

  Similarly, in the DPG $(G, PA(p))$, Adam has an optimal strategy $\gamma^* \in \Gamma$, such that for all strategies $\sigma \in \Sigma$ of Eve, all cycles in $G_{\sigma,\gamma^*} \restriction U_\forall$ are odd.

  We take an arbitrary vertex $v \notin U_\exists$ that is not winning for Eve, or equivalently an arbitrary vertex $v \in V$ where $\langle E\rangle(PA(p))(v) \neq 1$. Since $U_\exists \uplus U_\forall = V$, we have $v \in U_\forall$, and therefore a play starting from v does not leave $G_{\sigma,\gamma^*} \restriction U_\forall$.

  We now look at the SSG $(\overline{G}, RE(v_{\text{win}}))$. We define $\widetilde{U_\forall} = \overline{U_\forall} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$ such that $\overline{G}_{\sigma^*,\gamma} \restriction \widetilde{U_\forall}$ is well defined. It follows from the construction of $\overline{G}$ that all rBSCCs in $\overline{G}_{\sigma,\gamma^*} \restriction \widetilde{U_\forall}$ are odd, and a play starting from $\overline{v}$ does not leave

$\overline{G}_{\sigma,\gamma^*} \upharpoonright \widetilde{U}_\forall$. Thus it follows from (A_0') and (A_2') that:

$$\begin{aligned}\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma^*}(\mathit{Reach}(v_{\mathrm{win}})) &\leq 1 - \overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma^*}(\mathit{crossPath}) + \overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma^*}(\mathit{crossPath}) \cdot \overline{\mathrm{Pr}}^{\max}_{\sigma,\gamma^*}(\mathit{winOdd}) \\ &< \frac{1}{6} + \frac{5}{6} \cdot \frac{2}{5} \\ &= \frac{1}{2}\end{aligned}$$

We obtain for the value function $\langle E\rangle(RE(v_{\mathrm{win}}))(\overline{v})$ of the SSG that:

$$\langle E\rangle(RE(v_{\mathrm{win}}))(\overline{v}) = \sup_{\sigma\in\Sigma} \overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma^*}(\mathit{Reach}(v_{\mathrm{win}})) \leq \frac{1}{2}$$

□

4.2.4 Extension to Qualitative SPGs

We now present a direct reduction from qualitative SPGs to quantitative SSGs, using the gadgets from Section 4.2.1.

Theorem 4.5 (Reducing Qualitative SPGs to Quantitative SSGs)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\mathrm{win}}))$ as in Section 4.2.1. We introduce a constant b in $\overline{G}$ defined as follows:

$$b = \min_{\substack{\sigma\in\Sigma,\gamma\in\Gamma \\ \hat{\pi}\text{ contains no cycle}}} \overline{\mathrm{Pr}}_{\sigma,\gamma}(\hat{\pi})$$

Intuitively, b is the smallest probability of a simple path in $\overline{G}_{\sigma,\gamma}$ for all $\sigma, \gamma \in \Sigma \times \Gamma$.

If for all $\sigma, \gamma \in \Sigma \times \Gamma$ and $v \in V$, the following conditions hold:

(B_0) $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > 1 - \frac{1}{2}b$

(B_1) $\overline{\mathrm{Pr}}^{\min}_{\sigma,\gamma}(\mathit{winEven}) \geq \frac{2}{2+b}$, $\overline{\mathrm{Pr}}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) \leq 1 - \frac{2}{2+b}$

then $\langle E\rangle(PA(p))(v) = 1$ in the SPG $(G, PA(p))$ if and only if $\langle E\rangle(RE(v_{\mathrm{win}}))(\overline{v}) > \frac{2-b}{2+b}$ in the SSG $(\overline{G}, RE(v_{\mathrm{win}}))$. ■

Note that both $1 - \frac{1}{2}b$ and $\frac{2}{2+b}$ are close to 1. With fixed strategies from both players, the key of an SPG is what happens inside BSCCs. In the newly constructed SSG, we divide a play into two stages: before and after reaching an rBSCC. The intuition of the reduction is: if condition (B_0) is satisfied, then the outcome of the SSG is highly likely to depend on what happens inside rBSCCs; if condition (B_1) is satisfied, then the outcome after reaching an rBSCC is similar to the outcome of the corresponding BSCC in the SPG. Thus the SSG is an approximate simulation of the SPG. We present the formal proof as follows.

Proof. We assume that conditions (B_0) and (B_1) hold. We show that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) = 1$ in the SPG $(G, PA(p))$ if and only if $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) > \frac{2-b}{2+b}$ in the SSG $(\overline{G}, RE(v_{\text{win}}))$. Let $U_\exists$ and $U_\forall$ be the almost-sure winning regions of Eve and Adam in $(G, PA(p))$ respectively, and we define $\widetilde{U}_\exists$ and $\widetilde{U}_\forall$ in the same way as 4.2.3.

- We first prove that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) = 1$ in the SPG implies $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) > \frac{2-b}{2+b}$ in the SSG.

  We look at the SPG $(G, PA(p))$. According to Lemma 3.1, Eve has an optimal strategy $\sigma^* \in \Sigma$, such that for all strategies $\gamma \in \Gamma$ of Adam, all BSCCs in $G_{\sigma^*,\gamma} \upharpoonright U_\exists$ are even.

  We take an arbitrary almost-sure winning vertex $v \in U_\exists$ of Eve, or equivalently an arbitrary vertex $v \in V$ where $\langle E\rangle(PA(p))(v) = 1$. According to Lemma 3.1, a play starting from v does not leave $G_{\sigma^*,\gamma} \upharpoonright U_\exists$.

  We now look at the SSG $(\overline{G}, RE(v_{\text{win}}))$. It follows from the construction of $\overline{G}$ that all rBSCCs in $\overline{G}_{\sigma^*,\gamma} \upharpoonright \widetilde{U}_\exists$ are even, and a play starting from $\overline{v}$ does not leave $\overline{G}_{\sigma^*,\gamma} \upharpoonright \widetilde{U}_\exists$. Thus it follows that:

$$\begin{aligned}\overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{Reach}(v_{\text{win}})) &\geq \overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{crossPath}) \cdot \overline{\Pr}^{\min}_{\sigma^*,\gamma}(\mathit{winEven}) \\ &> (1 - \frac{1}{2}b) \cdot \frac{2}{2+b} \\ &= \frac{2-b}{2+b}\end{aligned}$$

  We can obtain for the value function of $(\overline{G}, RE(v_{\text{win}}))$ that:

$$\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) \geq \inf_{\gamma\in\Gamma} \overline{\Pr}^{\overline{v}}_{\sigma^*,\gamma}(\mathit{Reach}(v_{\text{win}})) > \frac{2-b}{2+b}$$

- For the other direction, we prove that for all vertices $v \in V$, $\langle E\rangle(PA(p))(v) \neq 1$ in the SPG implies $\langle E\rangle(RE(v_{\text{win}}))(\overline{v}) \leq \frac{2-b}{2+b}$ in the SSG.

  We look at the SPG $(G, PA(p))$. We take an arbitrary vertex $v \notin U_\exists$ that is not almost-sure winning for Eve, or equivalently an arbitrary vertex $v \in V$ where $\langle E\rangle(PA(p))(v) \neq 1$. As a result, if Adam follows an optimal strategy $\gamma^* \in \Gamma$, then a play starting from v in G_{σ,γ^*} has a positive probability of reaching an odd BSCC.

  We now look at the SSG $(\overline{G}, RE(v_{\text{win}}))$. It follows from the construction of $\overline{G}$ that a play starting from $\overline{v}$ has a positive probability of reaching an odd rBSCC, and $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterOdd) \geq b$. Thus it follows that:

$$
\begin{aligned}
&\overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(Reach(v_{\text{win}}))\\
&< (1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(crossPath)) + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterEven) \cdot 1 + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterOdd) \cdot \overline{\Pr}^{\max}_{\sigma,\gamma^*}(winOdd)\\
&\leq (1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(crossPath)) + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterEven) \cdot 1 + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterOdd) \cdot (1 - \frac{2}{2+b})\\
&= 1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(enterOdd) \cdot \frac{2}{2+b}\\
&\leq 1 - b \cdot \frac{2}{2+b}\\
&= \frac{2-b}{2+b}
\end{aligned}
$$

We obtain for the value function of $(\overline{G}, RE(v_{\text{win}}))$ that

$$\langle E \rangle (RE(v_{\text{win}}))(\overline{v}) \leq \sup_{\sigma \in \Sigma} \overline{\Pr}^{\overline{v}}_{\sigma,\gamma^*}(Reach(v_{\text{win}})) \leq \frac{2-b}{2+b} \qquad \square$$

4.3 Lower Bounds of *crossPath* and *winEven* Probabilities

In this section, we show that conditions (B_0) and (B_1) are satisfiable if we arrange the monotonically decreasing function $\alpha : \mathbb{N} \to [0,1]$ properly. We first present a useful lemma, and then show how to arrange α to satisfy (B_0) and (B_1) respectively.

Lemma 4.6 (Lower Bound of Sink Reachability)
Let $\mathcal{M} = (V \uplus \{v_f, v_s, v_b\}, \delta, v_I)$ be a Markov Chain with m states. If the following conditions hold:

(1) There are only two BSCCs, namely $\{v_s\}$ and $\{v_b\}$, and $\delta(v_s, v_s) = \delta(v_b, v_b) = 1$.

(2) For all states $v \in V$,

- $\delta(v, v_s) = \alpha$, $\delta(v, V \uplus \{v_f\}) = 1 - \alpha$, $\delta(v, v_b) = 0$.
- $\Pr^v(\mathit{Reach}(v_b)) > 0$.
- For all $v' \in V \uplus \{v_f\}$, $\delta(v, v') > 0$ implies that $\delta(v, v') \geq s$.

(3) For v_f, $\delta(v_f, V \uplus \{v_f\}) = k$ and $\delta(v_f, v_b) = l$.

Then for all $v \in V$, we can obtain a lower bound for the reachability probability of v_b as follows:

$$\Pr^v(\mathit{Reach}(v_b)) \geq \frac{(1-s) \cdot s^m \cdot l}{1 - (s+t) + k \cdot s^{m+1} + t \cdot s^m - k \cdot s^m}$$

where $t = 1 - \alpha - s$. ■

Proof. We assume all conditions in Lemma 4.6 hold.

For all $v \in V$, a play starting from v has to reach v_f first to reach v_b, and therefore:

$$\Pr^v(\mathit{Reach}(v_b)) = \Pr^v(\mathit{Reach}(v_f)) \cdot \Pr^{v_f}(\mathit{Reach}(v_b)) \leq \Pr^{v_f}(\mathit{Reach}(v_b))$$

We rename all $v \in V$ in the order of the probability for a play starting from v to reach v_b, such that:

$$0 < \Pr^{v_1}(\mathit{Reach}(v_b)) \leq \cdots \leq \Pr^{v_m}(\mathit{Reach}(v_b)) \leq \Pr^{v_f}(\mathit{Reach}(v_b)) \qquad (4.2)$$

For all $i = 1, 2, \cdots, m$, we denote $\Pr^{v_i}(\mathit{Reach}(v_b))$ with x_i, and we denote $\Pr^{v_f}(\mathit{Reach}(v_b))$ with x_{m+1}, so we write inequality 4.2 as:

$$0 < x_1 \leq x_2 \leq \cdots \leq x_m \leq x_{m+1}$$

It follows from condition (2) that for all $i = 1, 2, \cdots, m$:

$$\delta(v_i, v_s) = \alpha, \;\; \delta(v_i, v_b) = 0$$

$$\sum_{j=1}^{m+1} \delta(v_i, v_j) = 1 - \alpha$$

We recall Theorem 2.8 for the linear equation system for calculating readability probabilities in MCs. As there is no direct transition from v_i to v_b, we write x_i as:

$$x_i = \sum_{j=1}^{m+1} (\delta(v_i, v_j) \cdot x_j) \tag{4.3}$$

It follows from condition (2) that $\sum_{j=1}^{m+1} \delta(v_i, v_j) = 1 - \alpha$ and $x_i > 0$. We claim that for equation 4.3 to hold, there must exist $j > i$ such that $\delta(v_i, v_j) > 0$. Therefore we obtain the following for x_i:

$$\begin{aligned} x_i &= \sum_{j=1}^{m+1} (\delta(v_i, v_j) \cdot x_j) \\ &= \sum_{j=1}^{i} (\delta(v_i, v_j) \cdot x_j) + \sum_{j=i+1}^{m+1} (\delta(v_i, v_j) \cdot x_j) \\ &\geq (\sum_{j=1}^{i} \delta(v_i, v_j)) \cdot x_1 + (\sum_{j=i+1}^{m+1} \delta(v_i, v_j)) \cdot x_{i+1} \\ &\geq tx_1 + sx_{i+1} \end{aligned} \tag{4.4}$$

where $t = 1 - \alpha - s$.

We rearrange inequality 4.4 into the following form:

$$x_i - \frac{t}{1-s} x_1 \geq s(x_{i+1} - \frac{t}{1-s} x_1) \tag{4.5}$$

Similarly, for x_{m+1} we have the following:

$$\begin{aligned} x_{m+1} &= \sum_{j=1}^{m+1} (\delta(v_f, v_j) \cdot x_j) + \delta(v_f, v_b) \\ &\geq (\sum_{j=1}^{m+1} \delta(v_f, v_j)) \cdot x_1 + \delta(v_f, v_b) \\ &= kx_1 + l \end{aligned} \tag{4.6}$$

We combine inequalities 4.5 and 4.6 together to obtain the following:

$$\begin{aligned} x_1 - \frac{t}{1-s} x_1 &\geq s(x_2 - \frac{t}{1-s} x_1) \\ &\geq s^2(x_3 - \frac{t}{1-s} x_1) \\ &\geq \dots \\ &\geq s^m(x_{m+1} - \frac{t}{1-s} x_1) \\ &\geq s^m(kx_1 + l - \frac{t}{1-s} x_1) \end{aligned} \tag{4.7}$$

We rearrange the inequality 4.7 to obtain the final result:

$$x_1 \geq \frac{(1-s) \cdot s^m \cdot l}{1-(s+t)+k \cdot s^{m+1}+t \cdot s^m - k \cdot s^m} \qquad \square$$

Remark 4.7 (Actual Worst Case)
With all inequalities being equations in the previous proof, we get the intuitively 'worst case' Markov Chain. We give the figure below as an example for the scenario where $m = 5$. We draw the transitions to v_s with dotted arrows for visual neatness. Note that Lemma 4.6 applies no matter $\delta(v_f, v_s) > 0$ or not, so we draw this transition in a dashed arrow.

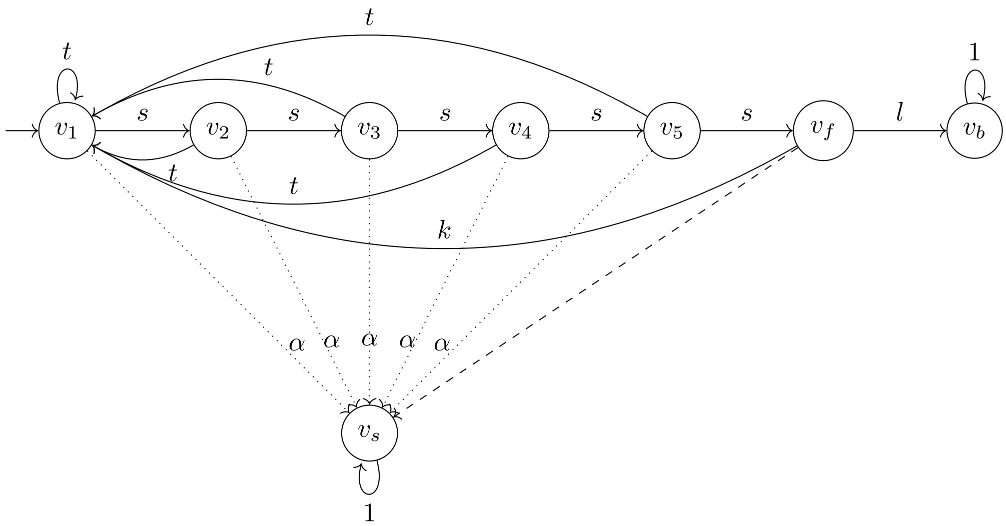

Figure 4.8: 'Worst case' Markov Chain

4.3.1 Before Entering An rBSCC

In this section, we focus on what happens before a play reaches an rBSCC in the newly constructed arena $\overline{G}$. Specifically, we show that condition (B_0) holds when $\alpha(0)$ is sufficiently small by proving the following theorem.

Theorem 4.8 (Lower Bound of *crossPath* Probability)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena with n vertices, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\text{win}}))$ as in Section 4.2.1. We introduce a constant b defined as in Theorem 4.5. We denote with $\delta_{\min}$ the smallest probability of a transition given by Δ, formally $\delta_{\min} = \min_{u\in V_R, v\in V} \Delta(u, v)$.

For the monotonically decreasing function $\alpha : \mathbb{N} \to [0, 1]$, if $\alpha(0) = \delta_{\min}^{2n+3}$, then for all $\sigma \in \Sigma$, $\gamma \in \Gamma$ and $v \in V$, the following holds:

$$\overline{\Pr}{}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > 1 - \frac{1}{2}b$$

■

Remark 4.9 (The Range of $\delta_{\min}$)
Note that if $\delta_{\min} = 1$, it indicates that $(G, PA(p))$ is a DPG, and we have discussed this scenario in Section 4.2.3. Therefore we assume $\delta_{\min} \neq 1$ in the sequel, indicating that $\delta_{\min} \in (0, \frac{1}{2}]$. ■

To prove Theorem 4.8, we first fix an arbitrary pair of strategies $\sigma, \gamma \in \Sigma \times \Gamma$, and obtain sub-arenas $G_{\sigma,\gamma}$ and $\overline{G}_{\sigma,\gamma}$.

We consider the Markov Chain views of $G_{\sigma,\gamma}$ and $\overline{G}_{\sigma,\gamma}$, namely $\mathcal{M}_{\sigma,\gamma}$ and $\overline{\mathcal{M}}_{\sigma,\gamma}$ respectively. Note that in the following part of this proof, we leave out the initial state of a Markov Chain for the reasons stated in Remark 2.22. Hence we have $\mathcal{M}_{\sigma,\gamma} = (V, \delta)$ and $\overline{\mathcal{M}}_{\sigma,\gamma} = (\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{\delta})$. Additionally, we denote the set of vertices that belongs to BSCCs in $\mathcal{M}_{\sigma,\gamma}$ with V_B and let $V_T = V \backslash V_B$. Correspondingly, the set of vertices that belong to rBSCCs in $\overline{\mathcal{M}}_{\sigma,\gamma}$ is $\overline{V}_B \uplus \widehat{V}_B$.

We present an example in Figure 4.9. For all vertices v_i in the figure, we let $p(v_i) = i$. We leave out v_{win} , v_{lose} and their incoming transitions, and represent newly introduced vertices with dotted circles for visual neatness. It is easy to observe that $V_B = \{v_5\}$ and $V_T = \{v_0, v_1, v_2, v_3, v_4\}$. We also use this Markov Chain as a running example in the following proof.

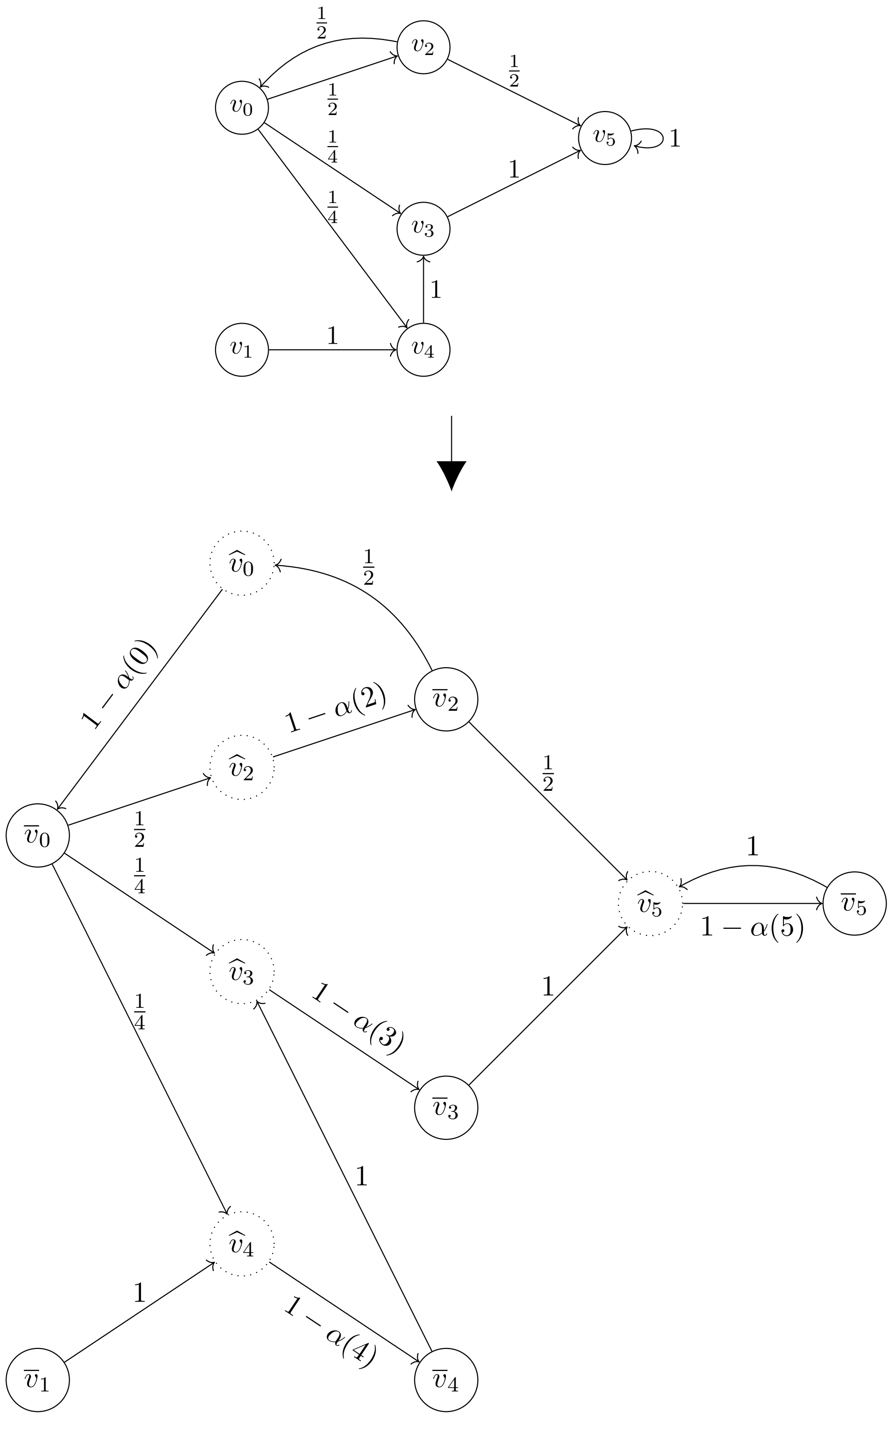

Figure 4.9: From $\mathcal{M}_{\sigma,\gamma}$ to $\overline{\mathcal{M}}_{\sigma,\gamma}$

The proof of Theorem 4.8 consists of two parts:

1. We make a 4-step transformation to $\overline{\mathcal{M}}_{\sigma,\gamma}$. We denote the resulting Markov Chain of the i-th step with $\mathcal{M}_i = (V_i, \delta_i)$. In the i-th Markov Chain, we denote with $\Pr_i^v(\mathit{Reach}(V_i'))$ the probability of reaching $V_i' \subseteq V_i$ from v. We show that for all $\overline{v} \in \overline{V}_T$, we have:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) \geq \Pr_4^v(\mathit{Reach}(v_b))$$

   where $v_b \in V_4$ is a specific vertex in $\mathcal{M}_4$.

2. We apply Lemma 4.6 to $\mathcal{M}_4$ to obtain a lower bound of $\Pr_4^v(\mathit{Reach}(v_b))$, and hence a lower bound of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})$. We show that when $\alpha(0) = \delta_{\min}^{2n+3}$, the following holds:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > 1 - \frac{1}{2}b$$

We now present the proof as follows.

First Part: Transforming the Markov Chain

We use the Markov Chains from Figure 4.9 as a running example.

STEP 1 In the first step, we eliminate $\widehat{V}$ according to Lemma 1 in [21]. Formally we have $V_1 = \overline{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$, and the resulting transition function δ_1 can be described as follows:

- For all $\overline{u}, \overline{v} \in \overline{V} \times \overline{V}$, $\delta_1(\overline{u}, \overline{v}) = (1 - \alpha(p(v))) \cdot \overline{\delta}(\overline{u}, \widehat{v})$.
- For all $\overline{u} \in \overline{V}$,

$$\delta_1(\overline{u}, v_{\text{win}}) = \sum_{p(v) \text{ is even}} \alpha(p(v)) \cdot \overline{\delta}(\overline{u}, \widehat{v})$$
$$\delta_1(\overline{u}, v_{\text{lose}}) = \sum_{p(v) \text{ is odd}} \alpha(p(v)) \cdot \overline{\delta}(\overline{u}, \widehat{v}).$$

- $\delta_1(v_{\text{win}}, v_{\text{win}}) = \delta_1(v_{\text{lose}}, v_{\text{lose}}) = 1$.

Note that this transformation does not change the probability for a play starting from $\overline{v} \in \overline{V}$ to reach v_{win} or v_{lose}. We refer to [21] for the details of this transformation. In the rest of this proof, we rename all vertices $\overline{v} \in \overline{V}$ as v for visual neatness. It follows that $\mathcal{M}_1 = (V \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \delta_1)$, and for all vertices $v \in V_T$:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) = \Pr_1^v(\mathit{Reach}(V_B))$$

In the following Figure 4.10, we present the resulting $\mathcal{M}_1$ for the running example. For visual neatness, we only draw the transitions from v_0 to v_{win}

and v_{lose}, which are the most complicated ones, as a demonstration. They are represented with dotted arrows.

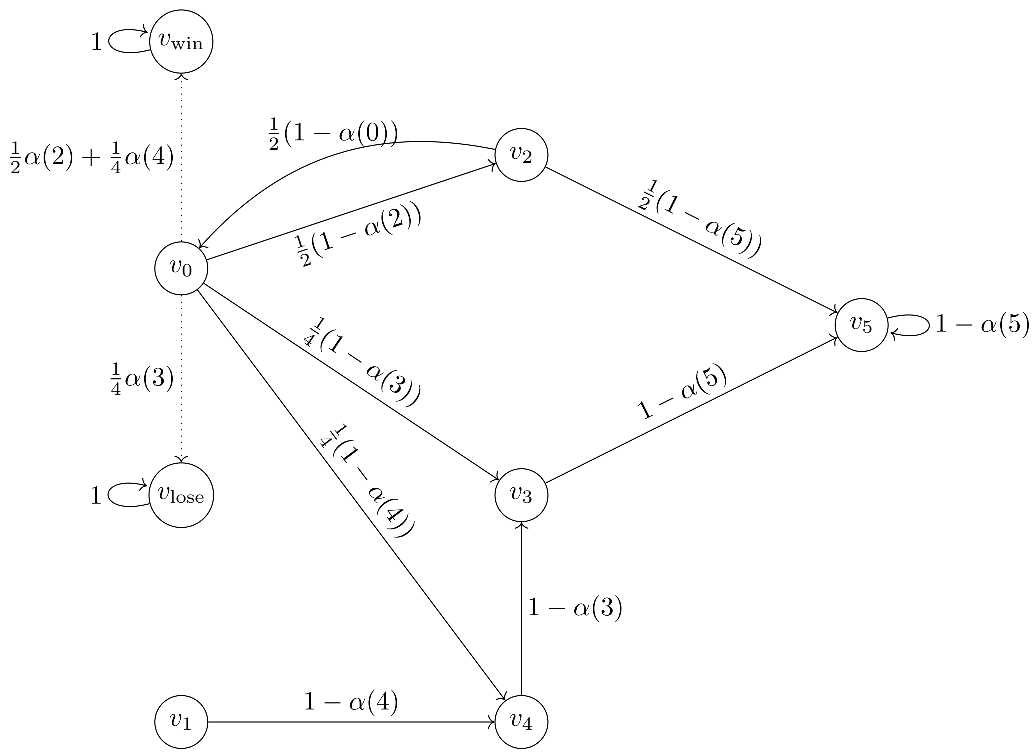

Figure 4.10: Before entering an rBSCC: eliminating $\widehat{V}$

Observation 4.10
Let us recall the original Markov Chain view $\mathcal{M}_{\sigma,\gamma} = (V, \delta)$. An Effect of this transformation is that for all pairs of vertices $u, v \in V \times V$, $\delta_1(u, v) = (1-\alpha(p(v)))\delta(u, v)$. Intuitively, if we do not consider v_{win} and v_{lose}, $\mathcal{M}_1$ has the same transitions as $\mathcal{M}_{\sigma,\gamma}$ with discounted probabilities, and it follows that for all $v \in V_T$, $\Pr_1^v(\mathit{Reach}(V_B)) > 0$. ■

STEP 2 In this step, we scale up the values of α to its maximum. Specifically, for all $n \in \mathbb{N}$, we let $\alpha(n) = \alpha(0)$. The vertices of the Markov Chain remain the same, hence $V_2 = V \uplus \{v_{\text{win}}, v_{\text{lose}}\}$. We obtain the resulting transition function δ_2 by substituting every occurrence of $\alpha(p(v))$ in δ_1 with $\alpha(0)$. Since we scale up the values of α, it is trivially observable that for all $v \in V_T$, the probability of reaching v_{win} or v_{lose} from v does not decrease, and thus:

$$\Pr_1^v(\mathit{Reach}(V_B)) \geq \Pr_2^v(\mathit{Reach}(V_B))$$

In the following Figure 4.11, we present v_0 and its outgoing transitions in the resulting $\mathcal{M}_2$ for the running example. The same idea applies to all other vertices.

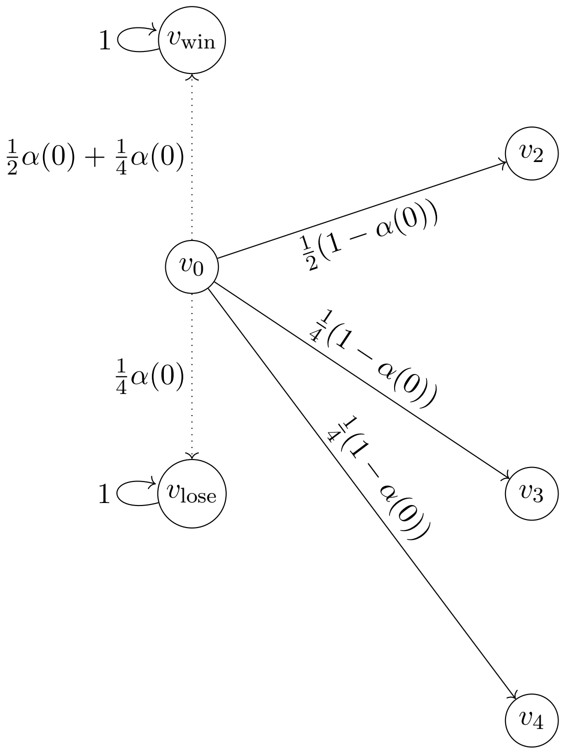

Figure 4.11: Before entering an rBSCC: scaling up α

Observation 4.11

- For all vertices $v \in V$, we have $\delta_2(v, \{v_{\text{win}}, v_{\text{lose}}\}) = \alpha(0)$ and $\delta_2(v, V) = 1 - \alpha(0)$.
- For all pairs of vertices $u, v \in V \times V$, $\delta_2(u, v) = (1 - \alpha(0))\delta(u, v)$. If we do not consider v_{win} and v_{lose}, $\mathcal{M}_2$ still has the same transitions as $\mathcal{M}_{\sigma,\gamma}$ with discounted probabilities, and it follows that for all $v \in V_T$, $\Pr_2^v(\mathit{Reach}(V_B)) > 0$. ■

STEP 3 In this step, we merge v_{win} and v_{lose} as a single sink v_s with a self-loop of probability 1. The subscript s here stands for *sinks.* All incoming transitions to either v_{win} or v_{lose} now go to v_s.

Additionally, we merge all vertices in V_B as a single vertex. To give a clear intuition, we consider this as two sub-steps:

- We first duplicate v_s, and denote the new one with v_s'. For all vertices $v \in V_B$, we replace the transitions (v, v_s) with transitions (v, v_s').
- It follows from the construction that there is no transition leaving $V_B \uplus$

$\{v'_s\}$. We collapse $V_B \uplus \{v'_s\}$ into one single vertex v_b with a self-loop of probability 1. The subscript b here stands for *bottom*. For all vertices $u \in V_T$, all transitions from u into V_B are merged into transition (u, v_b).

Formally the vertices now become $V_3 = V_T \uplus \{v_b, v_s\}$, and the new transition function δ_3 can be described as follows:

- For all $u, v \in V_T \times V_T$, we have $\delta_3(u, v) = \delta_2(u, v)$.
- For all $u \in V_T$, $\delta_3(u, v_s) = \delta_2(u, v_{\text{win}}) + \delta_2(u, v_{\text{lose}}) = \alpha(0)$.
- For all $u \in V_T$, $\delta_3(u, v_b) = \sum_{v \in V_B} \delta_2(u, v)$.
- $\delta_3(v_s, v_s) = \delta_3(v_b, v_b) = 1$.

It follows trivially from the transformation that for all vertices $v \in V_T$:

$$\Pr_2^v(Reach(V_B)) = \Pr_3^v(Reach(v_b))$$

In the following Figure 4.12, we present the resulting $\mathcal{M}_3$ for the running example. Note that we draw the transitions from V_T to v_s with dotted arrows and leave out their probabilities for visual neatness. All dotted transitions have probability $\alpha(0)$.

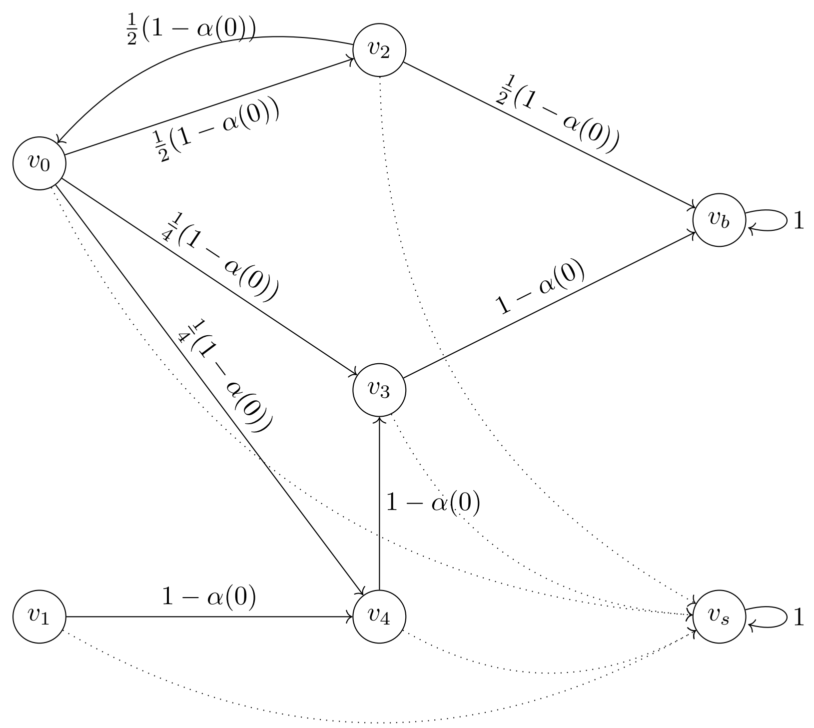

Figure 4.12: Before entering an rBSCC: merging sinks and rBSCCs

Observation 4.12
If we do not consider v_{win} and v_{lose}, the non-BSCC part of $\mathcal{M}_3$ still has the same transitions as the non-BSCC part of $\mathcal{M}_{\sigma,\gamma}$ with discounted probabilities, and it follows that for all $v \in V_T$, $\Pr_3^v(\mathit{Reach}(v_b)) > 0$. ■

STEP 4 Before going into the transformation, we introduce the notion of *frontier vertex.* In the resulting Markov Chain $\mathcal{M}_3$ of previous steps, we call $v_f \in V_T$ a *frontier vertex* if $\delta_3(v_f, v_b) > 0$, and we denote the set of frontier vertices with V_F. It is easy to observe in the running example that $V_F = \{v_2, v_3\}$.

Since all paths of a finite-state Markov Chain reach a BSCC, for all vertices $v \in V_T$, there is at least one path from v to v_b. We fix a starting vertex $v_0 \in V_T$. Without loss of generality, we assume all other vertices are reachable from v_0. Otherwise, we can always eliminate the unreachable fragment. In the running example, we fix a starting vertex v_0, and v_4 can be ignored as it is not reachable from v_0.

In this step, we apply the following procedure to $\mathcal{M}_3 = (V_T \uplus \{v_b, v_s\}, \delta_3)$ to *'defrontierize'* all but one frontier vertices:

```
1: procedure DFV(M = (V_T ⊎ {v_b, v_s}, δ))// M results from STEP 3
2:     while |V_F| ≥ 2 do// M has more than one frontier vertices
3:         Take an arbitrary vertex v_f ∈ V_F.
4:         Take a vertex v_p ∈ V_T\V_F, s.t. there exists π = v_p ··· v'_f v_b where v'_f ≠ v_f.
5:         δ(v_f, v_p) ← δ(v_f, v_p) + δ(v_f, v_b)
6:         δ(v_f, v_b) ← 0
7:     end while
8:     return M
9: end procedure
```

Remark 4.13 (Defrontierize Frontier Vertices)

- Regarding line 4, there must exist a vertex $v_p \in V_T \backslash V_F$, such that there exists a path $\pi = v_p \cdots v'_f v_b$, where $v'_f \neq v_f$. Otherwise, it indicates that v_f is the only frontier vertex. The subscript p here stands for *pivot.* Moreover, this ensures that if there is a path $v \cdots v_f v_b$ before an iteration, there is a path $v \cdots v_f v_p \cdots v_b$ after the iteration.
- Regarding lines $5-6$, we consider them as the following equivalent version for clarity:

$$\begin{aligned}
&\delta' \leftarrow \delta \\
&\delta'(v_f, v_p) \leftarrow \delta(v_f, v_p) + \delta(v_f, v_b) \\
&\delta'(v_f, v_b) \leftarrow 0 \\
&\delta \leftarrow \delta'
\end{aligned}$$

It follows that v_f is no longer a frontier vertex at the end of the iteration. ■

In the following Figure 4.13, we present the resulting Markov Chain for the running example. Here we ignore v_1 as it is not reachable from v_0, and defrontierize v_3 by substituting transition (v_3, v_b) with transition (v_3, v_0). Note that all dotted transitions still have probability $\alpha(0)$, and all deleted transitions are represented by dashed arrows.

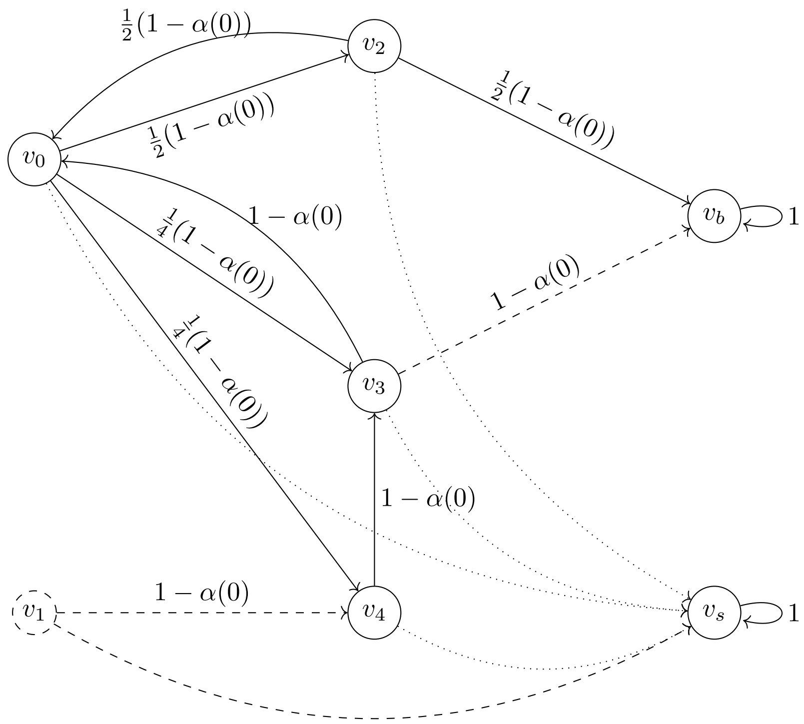

Figure 4.13: Before entering an rBSCC: defrontierizing all but one frontier vertices

Regarding this procedure, we present the following lemma, which claims that the probability of reaching v_b from v_0 does not increase after each iteration. We attach the proof of Lemma 4.14 later.

Lemma 4.14 (Non-Increasing Reachability)
For each loop iteration, we denote with Pr and Pr′ the probabilities associated with transition function δ and δ' respectively, and it holds that $\Pr'^{v_0}(\mathit{Reach}(v_b)) \leq \Pr^{v_0}(\mathit{Reach}(v_b))$. ■

Furthermore, we apply a similar transformation as before to the last frontier vertex v_f so that $\delta_4(v_f, v_b) = \delta_{\min}\alpha(0)$. The redundant probability is transferred to transition (v_f, v_0). We present the following Figure 4.14 for the intuition in the running example, where v_2 is the last frontier vertex. We consider the transition (v_2, v_b) as two separate ones, and substitute one of them with a transition (v_2, v_0) for clarity. Note that the deleted and added transitions are represented by dashed and dotted arrows respectively.

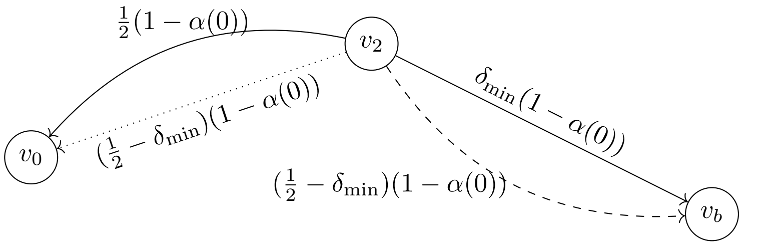

Figure 4.14: Before entering an rBSCC: the final action

As this transformation is essentially the same as 'defrontierizing', we obtain that the probability of reaching v_b from v_0 does not increase through a similar analysis to the proof of Lemma 4.14.

After applying the procedure, only one frontier vertex remains, and we denote it with v_f. The resulting Markov Chain is $\mathcal{M}_4 = (V'_T \uplus \{v_f, v_s, v_b\}, \delta_4)$, where $V'_T = V_T \backslash \{v_f\}$. It follows from Lemma 4.14 that:

$$\Pr_3^{v_0}(Reach(v_b)) \geq \Pr_4^{v_0}(Reach(v_b))$$

Observation 4.15

- It follows from previous observations and Remark 4.13 that for all vertices $v \in V'_T$, we have $\Pr_4^v(Reach(v_b)) > 0$.
- It follows from previous transformation steps that:
  - $\{v_s\}$ and $\{v_b\}$ are the only BSCCs, and $\delta_4(v_s, v_s) = \delta_4(v_b, v_b) = 1$.
  - For all vertices $v \in V'_T$, we have $\delta_4(v, v_s) = \alpha(0)$, $\delta_4(v, V'_T \uplus \{v_f\}) = 1 - \alpha(0)$ and $\delta_4(v, v_b) = 0$.
  - For all $v, v' \in (V'_T \uplus \{v_f\}) \times (V'_T \uplus \{v_f\})$, $\delta_4(v, v') > 0$ implies that $\delta_4(v, v') > \delta_{\min}(1 - \alpha(0))$.
- For the frontier vertex, we have $\delta_4(v_f, v_b) = \delta_{\min}\alpha(0)$ and $\delta_4(v_f, V'_T \uplus \{v_f\})) = \delta_{\min}(1 - \alpha(0))$. ■

Combining the reasoning of the 4-step transformation, we obtain that for all $\overline{v} \in \overline{V}_T$, $\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(crossPath) \geq \Pr_4^v(Reach(v_b))$, which concludes the first part of the proof of Theorem 4.8.

Second Part: Arranging $\alpha(0)$

We first apply Lemma 4.6 to $\mathcal{M}_4 = (V'_T \uplus \{v_f, v_s, v_b\}, \delta_4)$ from the previous part to obtain a lower bound of $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})$.

Lemma 4.16 (Lower Bound of *crossPath* Probability)
For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all $\overline{v} \in \overline{V}$, the following holds:

$$\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) \geq \frac{(1-s)s^n}{(1-s)-(1-s^n)t} \tag{4.8}$$

where $s = \delta_{\min}(1 - \alpha(0))$ and $t = (1 - \delta_{\min})(1 - \alpha(0))$. ■

Proof. We look at $\mathcal{M}_4 = (V'_T \uplus \{v_f, v_s, v_b\}, \delta_4)$ from the previous part. It follows from Observation 4.15 that Lemma 4.6 can be applied to $\mathcal{M}_4$, where:

- $s = l = \delta_{\min}(1 - \alpha(0))$
- $k = t = (1 - \delta_{\min})(1 - \alpha(0))$

Therefore we obtain a lower bound of $\mathrm{Pr}^{v}_{4}(\mathit{Reach}(v_b))$ as follows:

$$\mathrm{Pr}^{v}_{4}(\mathit{Reach}(v_b)) \geq \frac{(1-s)s^n}{(1-s)-(1-s^n)t}$$

where $s = \delta_{\min}(1 - \alpha(0))$ and $t = (1 - \delta_{\min})(1 - \alpha(0))$. We thus obtain a lower bound of $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})$ as well. □

Remark 4.17 (An Achievable Lower Bound of crossPath Probability)
Note that this value is achievable. It corresponds to the *crossPath* probability of $\overline{v}$ in the scenario given by Figure 4.15, where $\{\widehat{v}_b, \overline{v}_b\}$ is the only rBSCC, and all vertices have priority 0 in the original arena. ■

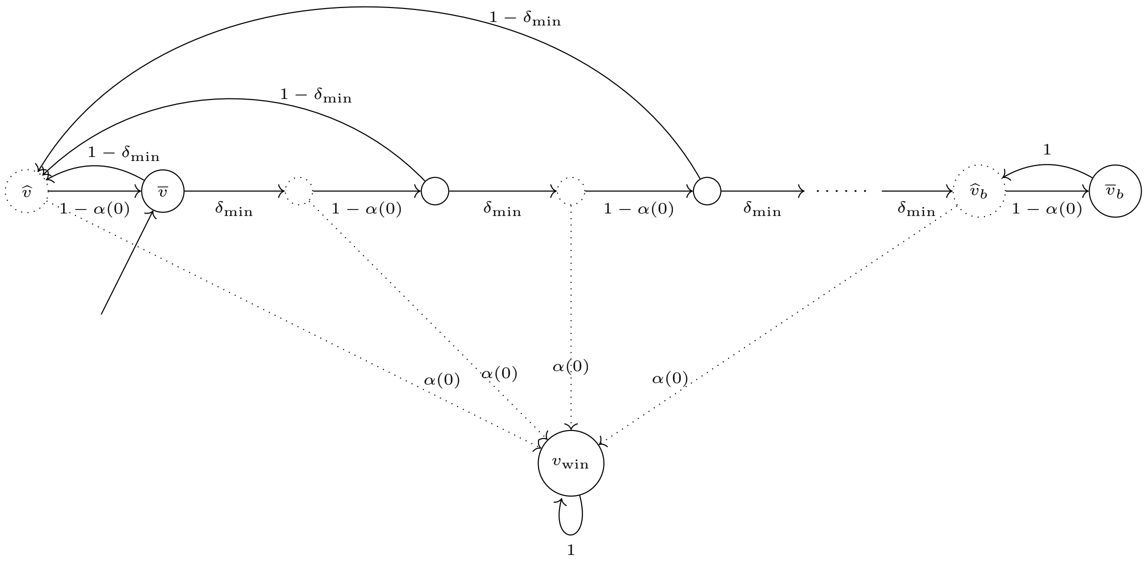

Figure 4.15: An achievable lower bound of *crossPath* probability

In the following calculation, we write $\alpha(0)$ as α for visual neatness.

Corollary 4.18 (An Easier Lower Bound of *crossPath* Probability)
For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all $\overline{v} \in \overline{V}$, the following holds:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > \frac{\delta_{\min}^n (1-\alpha)^{n+1}}{2\alpha + \delta_{\min}^n (1-\alpha)^{n+1}} \qquad \blacksquare$$

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

$$\begin{aligned}
&\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) \\
\geq\ & \frac{(1-s)s^n}{(1-s)-(1-s^n)t} && \textit{Lemma 4.16} \\
=\ & \frac{(1-\delta_{\min}(1-\alpha))\delta_{\min}^n(1-\alpha)^n}{\alpha + (1-\delta_{\min})\delta_{\min}^n(1-\alpha)^{n+1}} \\
>\ & \frac{(1-\delta_{\min})(1-\alpha)\delta_{\min}^n(1-\alpha)^n}{\alpha + (1-\delta_{\min})\delta_{\min}^n(1-\alpha)^{n+1}} && \textit{since } 1-\delta_{\min}(1-\alpha) > (1-\delta_{\min})(1-\alpha) \\
=\ & \frac{\delta_{\min}^n(1-\alpha)^{n+1}}{\frac{\alpha}{1-\delta_{\min}} + \delta_{\min}^n(1-\alpha)^{n+1}} \\
\geq\ & \frac{\delta_{\min}^n(1-\alpha)^{n+1}}{2\alpha + \delta_{\min}^n(1-\alpha)^{n+1}} && \textit{since } 1-\delta_{\min} \geq \tfrac{1}{2}
\end{aligned}$$

This concludes the proof. $\square$

We recall that b is the minimum probability of a simple path in $\overline{G}$, and it follows that:

$$b \geq (\delta_{\min}(1-\alpha))^n \tag{$*$}$$

To prove $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > 1 - \frac{1}{2}b$, it suffices to show that:

$$\frac{1}{2}(\delta_{\min}(1-\alpha))^n \geq \frac{2\alpha}{2\alpha + \delta_{\min}^n(1-\alpha)^{n+1}} \tag{4.9}$$

Now we show that if $\alpha = \delta_{\min}^{2n+3}$, the inequality 4.9 holds. We substitute α with $\delta_{\min}^{2n+3}$ and rewrite inequality 4.9 as follows:

$$\delta_{\min}^{3n+3}(1-\delta_{\min}^{2n+3})^n + \frac{1}{2}\delta_{\min}^{2n}(1-\delta_{\min}^{2n+3})^{2n+1} \geq 2\delta_{\min}^{2n+3} \tag{4.10}$$

For the left side of inequality 4.10, we obtain the following:

$$\begin{aligned}
& \frac{1}{2}\delta_{\min}^{2n}(1-\delta_{\min}^{2n+3})^{2n+1} & \\
\geq\ & \frac{1}{2}\delta_{\min}^{2n}(1-(2n+1)\delta_{\min}^{2n+3}) & \textit{follows from Bernoulli inequality} \\
=\ & \frac{1}{2}\delta_{\min}^{2n} - \frac{(2n+1)}{2}\delta_{\min}^{4n+3} & \\
\geq\ & \frac{1}{2}\delta_{\min}^{2n} - \frac{3}{64}\delta_{\min}^{2n} & \textit{follows from } n \geq 1, \delta_{\min} \leq \tfrac{1}{2} \\
=\ & \frac{29}{64}\delta_{\min}^{2n} = 2\delta_{\min}^{2n+3} \cdot \frac{29}{128\delta_{\min}^3} & \\
\geq\ & 2\delta_{\min}^{2n+3} & \textit{since } 128\delta_{\min}^3 \leq 16
\end{aligned}$$

Therefore when $\alpha(0) = \delta_{\min}^{2n+3}$, condition (B_0) is satisfied:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > 1 - \frac{1}{2}b$$

This concludes the proof of Theorem 4.8.

Proof of Lemma 4.14

We recall Lemma 4.14, and present its proof here.

Proof. Without loss of generality, we assume that there is no transition (v_f, v_p) before the loop iteration. Otherwise, we can consider it as a separate transition from the newly added one.

- Before the loop iteration, we can classify the finite paths from v_0 to v_b into two sets, namely:

- the paths that reach v_b via the transition (v_f, v_b);
- the paths that reach v_b via a transition (v'_f, v_b), where $v'_f \neq v_f$.

We write the finite paths from v_0 to v_b formally as:

$$\{v_0 \cdots v_f v_b\} \uplus \{v_0 \cdots v'_f v_b \mid v'_f \neq v_f\}$$

Note that before the loop iteration, the transition (v_f, v_p) is not available in either case, so we also write them as:

$$\{\overbrace{v_0 \cdots}^{\text{no } v_f v_p} v_f v_b\} \uplus \{\overbrace{v_0 \cdots}^{\text{no } v_f v_p} v'_f v_b \mid v'_f \neq v_f\}$$

- After the iteration, we can classify the finite paths from v_0 to v_b into two sets:
  - the paths that take the transition (v_f, v_p) at least once;
  - the paths never take the transition (v_f, v_p).

  We write them formally as:

$$\{\overbrace{v_0 \cdots}^{\textit{no } v_f v_p} v_f v_p \cdots v_b\} \uplus \{\overbrace{v_0 \cdots}^{\textit{no } v_f v_p} v'_f v_b \mid v'_f \neq v_f\}$$

  Note that in the latter set $v'_f \neq v_f$ is ensured implicitly since the transition (v_f, v_b) has been removed.

Therefore the difference in reachability probabilities lies in the former set of finite paths, and we obtain the following:

$$\begin{aligned}
&\Pr'^{v_0}(\mathit{Reach}(v_b)) - \Pr^{v_0}(\mathit{Reach}(v_b)) \\
&= \Pr'(\{\overbrace{v_0 \cdots}^{\textit{no } v_f v_p} v_f v_p \cdots v_b\}) - \Pr(\{\overbrace{v_0 \cdots}^{\textit{no } v_f v_p} v_f v_b\}) \\
&= \Pr'(\{\pi \mid \pi = \overbrace{v_0 \ldots}^{\textit{no } v_f v_p} v_f\}) \cdot \delta'(v_f, v_p) \cdot \Pr'(\{\pi \mid \pi = v_p \ldots v_f\}) \\
&\quad - \Pr(\{\pi \mid \pi = v_0 \ldots v_f\}) \cdot \delta(v_f, v_b) \\
&= \Pr(\{\pi \mid \pi = v_0 \ldots v_f\}) \cdot \delta(v_f, v_b) \cdot \Pr'(\{\pi \mid \pi = v_p \ldots v_f\}) \\
&\quad - \Pr(\{\pi \mid \pi = v_0 \ldots v_f\}) \cdot \delta(v_f, v_b) \\
&= \Pr(\{\pi \mid \pi = v_0 \ldots v_f\}) \cdot \delta(v_f, v_b) \cdot (\Pr'(\{\pi \mid \pi = v_p \ldots v_f\}) - 1) \\
&\leq 0
\end{aligned}$$

We thus conclude that in each loop iteration, the probability of reaching v_b from v_0 does not increase. $\square$

4.3.2 Inside An rBSCC

In this section, we focus on what happens after a play reaches an rBSCC in the newly constructed arena $\overline{G}$. Specifically, we show that condition (B_1) holds when $\alpha(k+1)/\alpha(k)$ is sufficiently small for all $k \in \mathbb{N}$ by proving the following theorem.

Theorem 4.19 (Lower Bound of Minimum *winEven* Probability)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena with n vertices, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\text{win}}))$ as in Section 4.2.1. We introduce a constant b in $\overline{G}$ defined as in 4.5. We denote with $\delta_{\min}$ the smallest probability of a transition given by Δ, formally $\delta_{\min} = \min_{u \in V_R, v \in V} \Delta(u, v)$.

For the monotonically decreasing function $\alpha : \mathbb{N} \to [0, 1]$, if $\alpha(k) = \delta_{\min}^{(2k+1)(2n+3)}$ for all $k \in \mathbb{N}$, then for all $\sigma \in \Sigma$, $\gamma \in \Gamma$ and $v \in V$, the following holds:

$$\overline{\Pr}_{\sigma,\gamma}^{\min}(\mathit{winEven}) \geq \frac{2}{2+b}$$

This indicates that the maximum probability for Adam to win in an even rBSCC is $1 - \frac{2}{2+b}$, and it follows from the duality that $\overline{\Pr}_{\sigma,\gamma}^{\max}(\mathit{winOdd}) \leq 1 - \frac{2}{2+b}$. ■

To prove Theorem 4.19, we fix an arbitrary pair of strategies $\sigma, \gamma \in \Sigma \times \Gamma$, and obtain sub-arenas $G_{\sigma,\gamma}$ and $\overline{G}_{\sigma,\gamma}$.

We consider the Markov Chain views of $G_{\sigma,\gamma}$ and $\overline{G}_{\sigma,\gamma}$, namely $\mathcal{M}_{\sigma,\gamma}$ and $\overline{\mathcal{M}}_{\sigma,\gamma}$ respectively. Note that in the following part of this proof, we leave out the initial state of a Markov Chain for the reasons stated in Remark 2.22. Hence we have $\mathcal{M}_{\sigma,\gamma} = (V, \delta)$ and $\overline{\mathcal{M}}_{\sigma,\gamma} = (\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{\delta})$.

We take an arbitrary even BSCC C in $\mathcal{M}_{\sigma,\gamma}$, and we denote the smallest priority of C with k. The counterpart of C in $\overline{\mathcal{M}}_{\sigma,\gamma}$ is an even rBSCC denoted with $\overline{C} \uplus \widehat{C}$. It follows from the construction of $\overline{G}$ that no transitions are leaving $\overline{C} \uplus \widehat{C}$ except for entering the winning and losing sinks. Therefore, $\overline{C} \uplus \widehat{C} \uplus \{v_{\text{win}}, v_{\text{lose}}\}$ and their internal transitions can be viewed as an independent Markov Chain, and we denote it with $\overline{\mathcal{M}}_c = (\overline{C} \uplus \widehat{C} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{\delta}_c)$. We denote all vertices of priority k in C with C_k, and let $C_{>k} = C \backslash C_k$.

Similar to the proof of Theorem 4.8, the proof of Theorem 4.19 also consists of two parts:

1. We make a 4-step transformation to $\overline{\mathcal{M}}_c$. We denote the resulting Markov Chain of the i-th step with $\mathcal{M}_{ci} = (V_{ci}, \delta_{ci})$. In the i-th Markov Chain, we denote with $\Pr_{ci}^{v}(\mathit{Reach}(V_i'))$ the probability of reaching $V_i' \subseteq V_i$ from v. We show that for all $\overline{v} \in \overline{C}$, we have:

$$\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(\mathit{Reach}(v_{\text{win}})) \geq \Pr_{c4}^{v}(\mathit{Reach}(v_{\text{win}}))$$

2. We apply Lemma 4.6 to $\mathcal{M}_{c4}$ to obtain a lower bound of $\Pr^{v}_{c4}(\mathit{Reach}(v_{\text{win}}))$, and hence a lower bound of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$. We generalize this to all rBSCCs, and show that when $\alpha(k) = \delta_{\min}^{(2k+1)(2n+3)}$ for all $k \in \mathbb{N}$, the following holds:

$$\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) \geq \frac{2}{2+b}$$

It follows from the duality that $\overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) \leq 1 - \frac{2}{2+b}$.

We now present the proof as follows.

First Part: Transforming the Markov Chain

The First Step: We first apply the same transformation as **STEP 1** of the previous proof to $\overline{\mathcal{M}}_c$, eliminating $\widehat{C}$ and renaming all $\overline{v} \in \overline{C}$ as v. We have $\mathcal{M}_{c1} = (C \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \delta_{c1})$, and for all $\overline{v} \in \overline{C}$, it holds that $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) = \Pr^{v}_{c1}(\mathit{Reach}(v_{\text{win}}))$.

The Second Step: For all vertex pairs $u, v \in C \times C_{>k}$ with $p(v) = 2m$ for some $m \in \mathbb{N}$, we transfer probability $\delta(u,v)\alpha(2m)$ from transition (u, v_{win}) to transition (u, v_{lose}). We give Figure 4.16 for an intuition. The deleted and added transitions are represented with a dashed arrow and a dotted arrow respectively, and the transition (u, v_{lose}) has probability $y + \delta(u,v)\alpha(2m)$ with two separate ones combined. It follows trivially that for all $v \in C$, we have $\Pr^{v}_{c1}(\mathit{Reach}(v_{\text{win}})) \geq \Pr^{v}_{c2}(\mathit{Reach}(v_{\text{win}}))$.

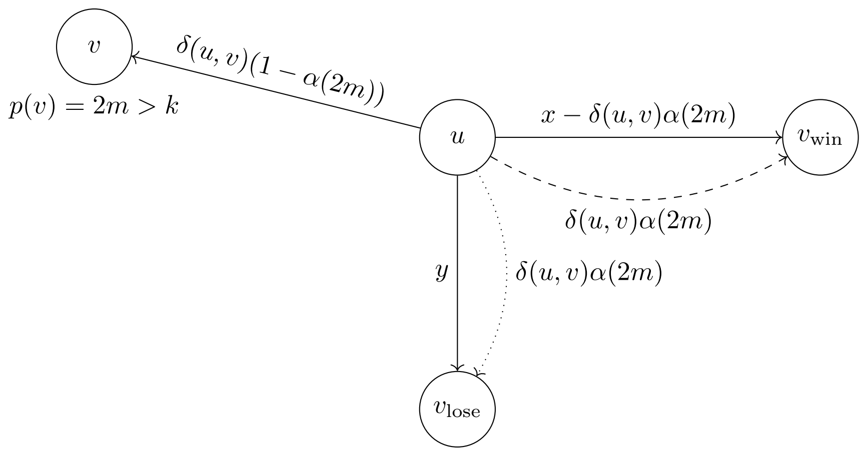

Figure 4.16: Inside an rBSCC: the second step

After the second step, we observe that all transitions entering v_{win} have the probability $(\sum_{v \in C_k} \delta(u,v))\alpha(k)$ for some $u \in C$.

The Third Step: We scale up α in a similar manner as in **STEP 2** of the previous proof. For all integers $n \geq k+1$, we let $\alpha(n) = \alpha(k+1)$. It follows trivially that for all $v \in C$, the probability of reaching v_{lose} from v does not decrease, and thus we have $\Pr^v_{c2}(\mathit{Reach}(v_{\text{win}})) \geq \Pr^v_{c3}(\mathit{Reach}(v_{\text{win}}))$.

The Fourth Step: We look at an vertex pair $v, v_k \in C \times C_k$.

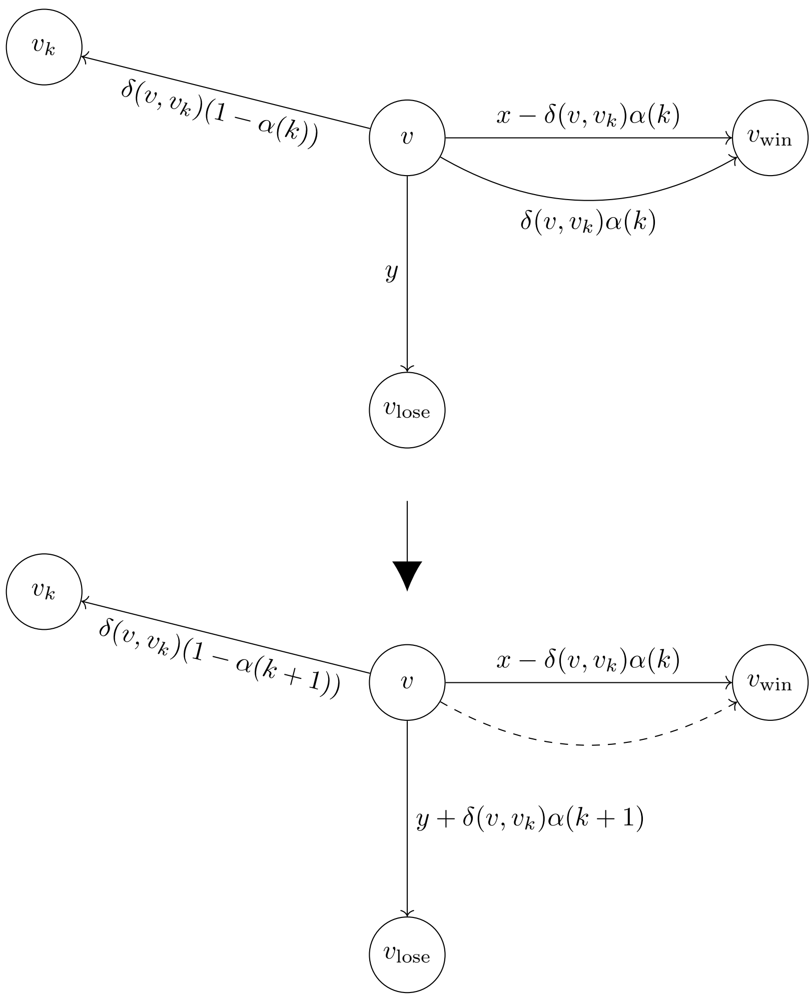

Figure 4.17: Inside an rBSCC: the fourth step

If we make the transformation given in Figure 4.17, with a similar analysis as in the proof of Lemma 4.14, we obtain the following for all $v_0 \in C$, where Pr and Pr′ are the probabilities associated with the transition functions before and after the transformation respectively.

$$
\begin{aligned}
&\mathrm{Pr}'^{v_0}(\mathit{Reach}(v_{\mathrm{win}})) - \mathrm{Pr}^{v_0}(\mathit{Reach}(v_{\mathrm{win}})) \\
= &\mathrm{Pr}(\{\pi \mid \pi = v_0 \cdots v\})\delta(v, v_k)(\alpha(k) - \alpha(k+1))\mathrm{Pr}'(\mathit{Reach}(v_{\mathrm{win}})) \\
&- \mathrm{Pr}(\{\pi \mid \pi = v_0 \cdots v\})\delta(v, v_k) \\
= &\mathrm{Pr}(\{\pi \mid \pi = v_0 \cdots v\})\delta(v, v_k)((\alpha(k) - \alpha(k+1))\mathrm{Pr}'(\mathit{Reach}(v_{\mathrm{win}})) - 1) \\
\leq &\, 0
\end{aligned}
$$

We apply this transformation to all but one vertex pairs $v, v_k \in C \times C_k$. After this there exists only one vertex pair $v, v_k \in C \times C_k$, where transition (v, v_k) and transition (v, v_{win}) have probabilities $x(1-\alpha(k))$ and $x\alpha(k)$ respectively. For clarity, we consider these two transitions as four, and apply the transformation given in Figure 4.18 to finish the last step.

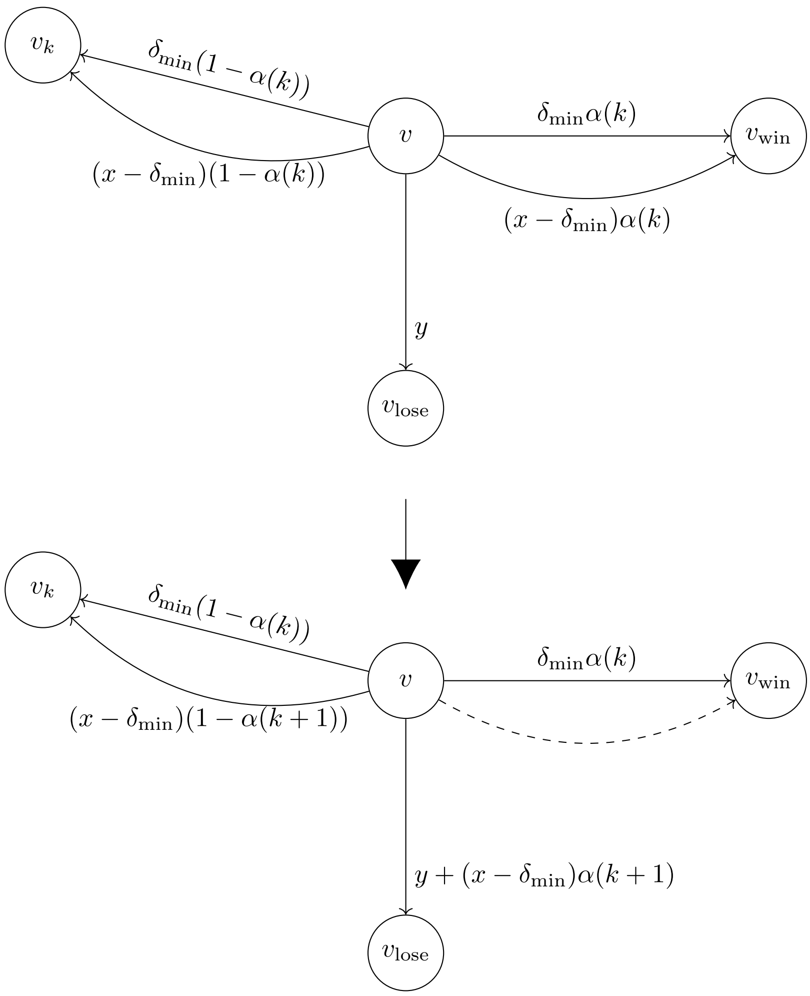

Figure 4.18: Inside an rBSCC: the final action

As a result, we observe that $\delta_{c4}(v, v_{\text{win}}) = \delta_{\min}\alpha(k)$ and $\delta_{c4}(v, v_{\text{lose}}) = (1-\delta_{\min})\alpha(k+1)$. As this last action is essentially the same as the transformation in Figure 4.17, we claim that for all $v \in C$, we have $\Pr^{v}_{c3}(\mathit{Reach}(v_{\text{win}})) \geq \Pr^{v}_{c4}(\mathit{Reach}(v_{\text{win}}))$. We also make the following observation regarding $\mathcal{M}_{c4}$.

Observation 4.20

- There are only two BSCCs in $\mathcal{M}_{c4}$, namely $\{v_{\text{win}}\}$ and $\{v_{\text{lose}}\}$, and $\delta_{c4}(v_{\text{win}}) = \delta_{c4}(v_{\text{lose}}) = 1$.
- There is only one vertex, denoted with $v_f \in C$, such that $\delta_{c4}(v_f, v_{\text{win}}) > 0$. We also have $\delta_{c4}(v_f, C) = \delta_{\min}(1-\alpha(k)) + (1-\delta_{\min})(1-\alpha(k+1))$ and $\delta_{c4}(v_f, v_{\text{win}}) = \delta_{\min}\alpha(k)$.
- Since we never change the connectivity inside C during the transformations, for all $v \in C$, we have $\Pr^{v}_{c4}(\mathit{Reach}(v_{\text{win}})) > 0$. For all $v \in C\backslash\{v_f\}$, we have $\delta_{c4}(v, v_{\text{lose}}) = \alpha(k+1)$, $\delta_{c4}(v, C) = 1 - \alpha(k+1)$ and $\delta_{c4}(v, v_{\text{win}}) = 0$; for all $v' \in C$, $\delta_{c4}(v, v') > 0$ implies that $\delta_{c4}(v, v') > \delta_{\min}(1 - \alpha(k+1))$. ■

Second Part: Arranging $\alpha(k+1)/\alpha(k)$

We first apply Lemma 4.6 to $\mathcal{M}_{c4}$ from the previous part to obtain a lower bound of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$.

Lemma 4.21 (Lower Bound of Minimum *winEven* Probability)
For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all $\overline{v} \in \overline{V}$ in an even rBSCC in $\overline{G}_{\sigma,\gamma}$ with smallest priority k, the following holds:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) \geq \frac{(1-s)\cdot s^{n-1}\cdot l}{1-(s+t)+k\cdot s^{n}+t\cdot s^{n-1}-k\cdot s^{n-1}}$$

where:

- $s = \delta_{\min}(1-\alpha(k+1))$
- $t = (1-\delta_{\min})(1-\alpha(k+1))$
- $l = \delta_{\min}\alpha(k)$
- $k = \delta_{\min}(1-\alpha(k)) + (1-\delta_{\min})(1-\alpha(k+1))$ ■

In the following reasoning, we write $\alpha(k)$ and $\alpha(k+1)$ as α_k and α_{k+1} respectively for visual neatness.

Remark 4.22 (A Tight Lower Bound of Minimum winEven Probability)
Note that this value corresponds to the winning probability of $\overline{v}$ in the scenario given by Figure 4.19, but the outgoing transitions of $\overline{v}_f$ is not achievable in $\overline{G}$. In the original arena G, this requires a transition of probability $\delta_{\min}$ from v_f to a vertex with

priority k, and a transition of probability $1 - \delta_{\min}$ from v_f to a vertex with priority $k+1$. These two transitions are merged as one from v_f to v in our scenario, but v has a unique priority $k+1$. However, we argue that this is still a relatively tight lower bound of minimum *winEven* probability. ■

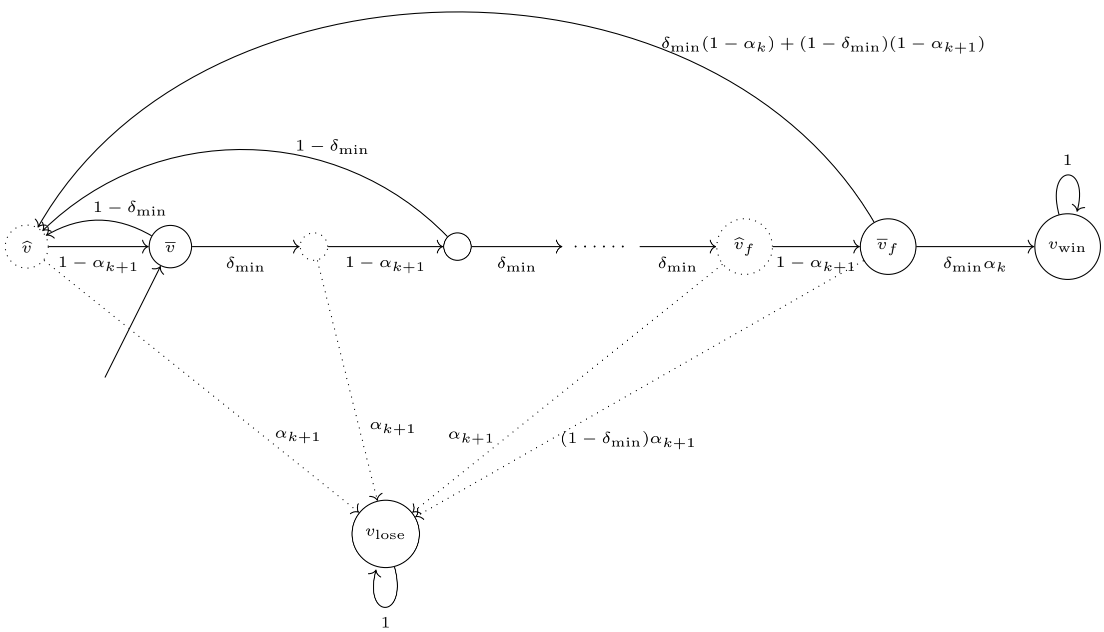

Figure 4.19: A tight lower bound of minimum *winEven* probability

Corollary 4.23 (An Easier Lower Bound of Minimum *winEven* Probability) For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all $\overline{v} \in \overline{V}$ in an even rBSCC in $\overline{G}_{\sigma,\gamma}$ with smallest priority k, the following holds:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(Reach(v_{\text{win}})) > \frac{\delta^n_{\min}(1-\delta_{\min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta^n_{\min}(1-\delta_{\min}) + \frac{\alpha_{k+1}}{\alpha_k}} \quad ■$$

Proof. We scale down the right side as follows:

$$
\begin{aligned}
&\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) \\
\geq\ & \frac{(1-s)\cdot s^{n-1}\cdot l}{1-(s+t)+k\cdot s^{n}+t\cdot s^{n-1}-k\cdot s^{n-1}} && \textit{Lemma 4.21} \\
=\ & \frac{(1-s)\cdot s^{n-1}\cdot l}{1-(s+t)+s^{n-1}(ks+t-k)} \\
=\ & \frac{(1-\delta_{\min}+\delta_{\min}\alpha_{k+1})\cdot\delta_{\min}^{n-1}(1-\alpha_{k+1})^{n-1}\cdot\delta_{\min}\alpha_k}{\alpha_{k+1}+\delta_{\min}^{n-1}(1-\alpha_{k+1})^{n-1}(\delta_{\min}^2(-\alpha_k+\alpha_{k+1}+\alpha_k\alpha_{k+1}-\alpha_{k+1}^2)+\delta_{\min}(\alpha_k-2\alpha_{k+1}+\alpha_{k+1}^2))} \\
>\ & \frac{(1-\delta_{\min})\cdot\delta_{\min}^{n}\cdot(1-\alpha_{k+1})^{n}\cdot\alpha_k}{\alpha_{k+1}+\delta_{\min}^{n}\cdot(\delta_{\min}(-\alpha_k+\alpha_{k+1}+\alpha_k\alpha_{k+1}-\alpha_{k+1}^2)+(\alpha_k-2\alpha_{k+1}+\alpha_{k+1}^2))} \\
& \qquad \textit{since } 1-\delta_{\min}+\delta_{\min}\alpha_{k+1}>(1-\delta_{\min})(1-\alpha_{k+1}) \textit{ and } (1-\alpha_{k+1})^{n-1}<1 \\
\geq\ & \frac{(1-\delta_{\min})\delta_{\min}^{n}\alpha_k(1-n\alpha_{k+1})}{\alpha_{k+1}+\delta_{\min}^{n}(\alpha_k(1-\delta_{\min})-\alpha_{k+1}(2-\alpha_{k+1}-\delta_{\min}(1+\alpha_k-\alpha_{k+1})))} \\
& && \textit{follows from Bernoulli inequality} \\
>\ & \frac{(1-\delta_{\min})\delta_{\min}^{n}\alpha_k-n\delta_{\min}^{n}(1-\delta_{\min})\alpha_k\alpha_{k+1}}{\alpha_{k+1}+(1-\delta_{\min})\delta_{\min}^{n}\alpha_k} \\
& && \textit{since } \alpha_{k+1}(2-\alpha_{k+1}-\delta_{\min}(1+\alpha_k-\alpha_{k+1}))>0 \\
>\ & \frac{\delta_{\min}^{n}\alpha_k(1-\delta_{\min})-\alpha_{k+1}}{\delta_{\min}^{n}\alpha_k(1-\delta_{\min})+\alpha_{k+1}} && \textit{since } n\delta_{\min}^{n}(1-\delta_{\min})\alpha_k<1 \\
=\ & \frac{\delta_{\min}^{n}(1-\delta_{\min})-\frac{\alpha_{k+1}}{\alpha_k}}{\delta_{\min}^{n}(1-\delta_{\min})+\frac{\alpha_{k+1}}{\alpha_k}}
\end{aligned}
$$

This concludes the proof. □

To prove $\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven})\geq\frac{2}{2+b}$, it suffices to show that for all $k\in\mathbb{N}$:

$$\frac{\delta_{\min}^{n}(1-\delta_{\min})-\frac{\alpha_{k+1}}{\alpha_k}}{\delta_{\min}^{n}(1-\delta_{\min})+\frac{\alpha_{k+1}}{\alpha_k}}>\frac{2}{2+b} \tag{4.11}$$

If we fix $\alpha(k)=\delta_{\min}^{(2k+1)(2n+3)}$ for all $k\in\mathbb{N}$, the left side is simplified to:

$$\frac{\delta_{\min}^{n}(1-\delta_{\min})-\frac{\alpha_{k+1}}{\alpha_k}}{\delta_{\min}^{n}(1-\delta_{\min})+\frac{\alpha_{k+1}}{\alpha_k}}=\frac{\delta_{\min}^{n}(1-\delta_{\min})-\delta_{\min}^{4n+6}}{\delta_{\min}^{n}(1-\delta_{\min})+\delta_{\min}^{4n+6}}=\frac{1-\delta_{\min}-\delta_{\min}^{3n+6}}{1-\delta_{\min}+\delta_{\min}^{3n+6}} \tag{4.12}$$

Therefore it suffices to show that:

$$(2+b)(1-\delta_{\min}-\delta_{\min}^{3n+6})>2(1-\delta_{\min}+\delta_{\min}^{3n+6}) \tag{4.13}$$

which can be further simplified to:

$$b(1-\delta_{\min}-\delta_{\min}^{3n+6})>4\delta_{\min}^{3n+6} \tag{4.14}$$

We start with the left side as follows:

$$
\begin{aligned}
& b(1-\delta_{\min}-\delta_{\min}^{3n+6}) \\
\geq\ & \delta_{\min}^{n}(1-\delta_{\min}^{2n+3})^{n}(1-\delta_{\min}-\delta_{\min}^{3n+6}) && \textit{follows from } (*) \\
\geq\ & \delta_{\min}^{n}(1-n\delta_{\min}^{2n+3})(1-\delta_{\min}-\delta_{\min}^{3n+6}) && \textit{follows from Bernoulli's inequality} \\
\geq\ & \delta_{\min}^{n}(1-\delta_{\min}^{n+3})(1-\delta_{\min}-\delta_{\min}^{3n+6}) && \textit{since } n\delta_{\min}^{n}<1 \\
\geq\ & \delta_{\min}^{n}\cdot\frac{7}{8}\cdot\frac{255}{512} && \textit{follows from the monotonicity and } \delta_{\min}\leq\tfrac{1}{2} \\
\geq\ & \delta_{\min}^{n}\cdot\frac{1}{64} \geq\ \delta_{\min}^{n}\cdot 4\delta_{\min}^{2n+6} = 4\delta_{\min}^{3n+6}
\end{aligned}
$$

Therefore when $\alpha(k)=\delta_{\min}^{(2k+1)(2n+3)}$, condition (B_1) is satisfied:

$$\overline{\Pr}_{\sigma,\gamma}^{\min}(\mathit{winEven}) \geq \frac{2}{2+b}$$

This concludes the proof of Theorem 4.19.

Chapter 5

Reducing Quantitative Parity Games to Quantitative SSGs

The contents of this chapter are given in Figure 5.1:

- We first revisit the reductions from quantitative SPGs to quantitative SMPGs, then to quantitative SDPGs, finally to quantitative SSGs.
- We then show that the reduction idea proposed in the previous chapter can be leveraged to further reduce quantitative SPGs to quantitative SSGs 'directly'.

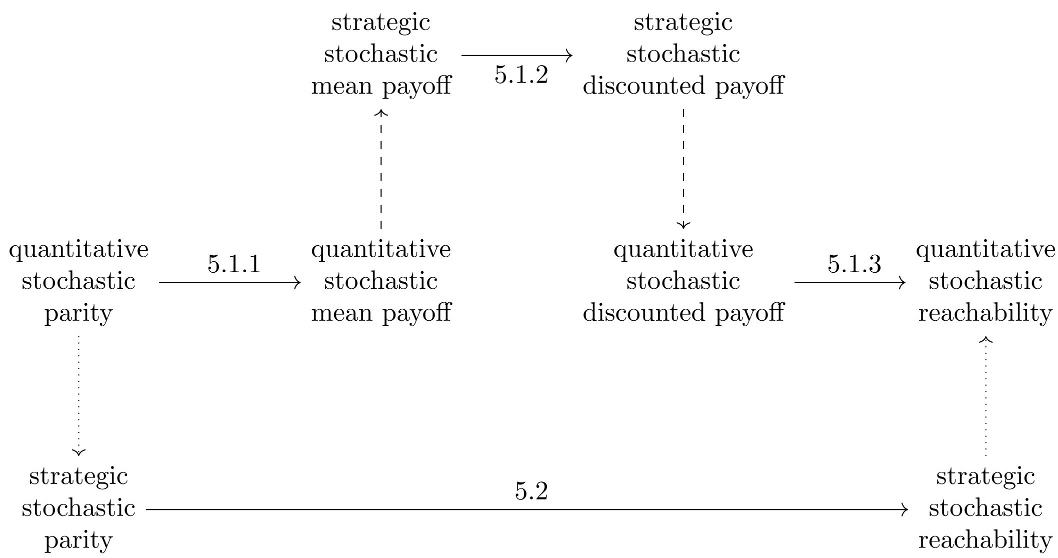

Figure 5.1: Reductions from quantitative parity to quantitative reachability

5.1 Previous Reduction Chain

In this section, we present the previous chain of reduction from quantitative SPGs to quantitative SSGs, as in the upper part of Figure 5.1.

As discussed in [2], the reductions drawn in dashed arrows are trivial for the following reasons, and we skip further details.

- From quantitative SMPGs to strategic SMPGs: once the strategies have been fixed, the values can be obtained by a complete analysis of the resulting Markov Chain.
- From strategic SDPGs to quantitative SDPGs: once the values are fixed, a strategy of Eve is optimal if and only if at each vertex of Eve it chooses an successor with the maximum value.

In the sequel, we sketch the other three reductions.

5.1.1 Quantitative SPGs to Quantitative SMPGs

We present the reduction from quantitative SPGs to quantitative SMPGs proposed in [9].

The first step of the reduction is from qualitative SPGs to quantitative SMPGs. We recall Section 4.1, where we present the reduction from qualitative SPGs to DPGs. Additionally, a reduction from DPGs to deterministic mean payoff games (DMPG), which is a special case of quantitative SMPGs, was proposed in [26]. We skip the details here. These two reductions complete the first step.

We now focus on the second step, that is, given an SPG and the almost-sure winning regions of Eve and Adam, we reduce the quantitative solution of this SPG to a quantitative SMPG.

Let $G = ((V,E),(V_\exists,V_\forall,V_R),\Delta)$ be a stochastic arena with n vertices, and $p : V \to \mathbb{N}$ be a priority function. We specify an SPG $(G, PA(p))$. Let the almost-sure winning regions of Eve and Adam be $U_\exists$ and $U_\forall$ respectively. We introduce a constant even number M such that the maximum priority given by p is either M or $M-1$. We further introduce a function $\bar{p} : V \to \mathbb{N}$ such that for all $v \in V$, we have $\bar{p}(v) = M - p(v)$. We denote with $\delta_{\min}$ the smallest probability of a transition given by Δ, formally $\delta_{\min} = \min_{u \in V_R, v \in V} \Delta(u,v)$. We construct a reward function $r : V \to \mathbb{R}$ as follows:

$$r(v) = \begin{cases} 1 & v \in U_\exists, \\ -1 & v \in U_\forall, \\ (-1)^{\bar{p}(v)} \cdot (2n)^{\bar{p}(v)} \cdot (\frac{1}{\delta_{\min}^n})^{\bar{p}(v)} & v \in V \setminus (U_\exists \uplus U_\forall). \end{cases}$$

Then for all $v \in V \setminus (U_\exists \uplus U_\forall)$ the following holds:

$$\langle E \rangle (PA(p))(v) = \frac{1}{2}(\langle E \rangle (MP(r))(v) + 1)$$

Before going into the details, we present a lemma that is used in the proof.

Lemma 5.1 (Long-Run Frequency)
Let $\mathcal{M} = (V, \delta, v_I)$ be a Markov Chain. We denote the minimum transition probability with p_m, formally written as:

$$p_m = \min_{u,v \in V \times V} \delta(u, v)$$

Let C be a BSCC in $\mathcal{M}$. We take an arbitrary $v_0 \in C$ as the starting state, and denote the state visited at time t with v_t. We define the *long-run frequency* of a state $v \in C$ from v_0 as:

$$\mathbf{freq}(v, v_0) = \lim_{n\to\infty} \frac{1}{n} \sum_{t=0}^{n} \Pr_{v_0}(v_t = v)$$

Then for all $v, v_0 \in C \times C$ we have:

$$\mathbf{freq}(v, v_0) \geq \frac{1}{n} \cdot p_m^n \qquad \blacksquare$$

We refer to [9] for the proof of Lemma 5.1.

We apply the notion of long-run frequency to SMPG as follows. If we fix arbitrary strategies $\sigma, \gamma \in \Sigma \times \Gamma$, all play π in $G_{\sigma,\gamma}$ enters a BSCC and stays till infinitum according to Theorem 2.12. We denote the BSCC with C and take an arbitrary vertex $v_0 \in C$, and we have $MP(r)(\pi) = \sum_{v \in C}(\mathbf{freq}(v, v_0) \cdot r(v))$.

We now present a proof sketch of the reduction.

Proof Sketch. We first show that $\langle E \rangle(PA(p))(v) \leq \frac{1}{2}(\langle E \rangle(MP(r))(v) + 1)$.

Let $\sigma^1 \in \Sigma$ be an optimal strategy of Eve for $PA(p)$ and $\gamma^* \in \Gamma$ be an optimal strategy of Adam for $MP(r)$. We fix σ^1 and γ^* to obtain G_{σ^1,γ^*}. We take an arbitrary BSCC from G_{σ^1,γ^*} and denote it with C. We make the following case analysis:

- If C is odd, it follows from Lemma 3.1 that $C \subseteq U_\forall$, and for all $v \in C$ we have $r(v) = -1$. As a result, if a play π enters C, we have $MP(r)(\pi) = -1$.
- If C is even, let z be a vertex with the minimum priority $p(z)$. Note that $\bar{p}(z)$ is also even, and we have $\max_{v \in C} \bar{p}(v) = \bar{p}(z)$.
  - If $\bar{p}(z) = 0$, it indicates that for all $v \in C$ we have $\bar{p}(v) = 0$, and thus $r(v) = 1$. As a result, if a play π enters C, we have $MP(r)(\pi) = 1$.
  - If $\bar{p}(z) \geq 2$, then we take an arbitrary $v_0 \in C$ and write the mean payoff of a play π entering C as:

$$MP(r)(\pi) = \mathbf{freq}(z, v_0) \cdot r(z) + \sum_{\substack{v \in C \\ v \neq z}} (\mathbf{freq}(v, v_0) \cdot r(v))$$

It follows from the construction of r that:

$$\begin{aligned}\mathbf{freq}(z, v_0) \cdot r(z) &= (\frac{1}{n} \cdot \delta_{\min}^n) \cdot ((2n)^{\bar{p}(z)} (\frac{1}{\delta_{\min}^n})^{\bar{p}(z)}) \\ &= 2 \cdot (2n)^{\bar{p}(z)-1} (\frac{1}{\delta_{\min}^n})^{\bar{p}(z)-1}\end{aligned}$$

For all vertices $v \in C$ and $v \neq z$, $r(v)$ reaches the smallest values when $\bar{p}(v) = \bar{p}(z) - 1$, and the following holds:

$$\begin{aligned}\sum_{\substack{v \in C \\ v \neq z}} (\mathbf{freq}(v, v_0) \cdot r(v)) &\geq (1 - \frac{1}{n} \cdot \delta_{\min}^n) \cdot (-(2n)^{\bar{p}(z)-1} (\frac{1}{\delta_{\min}^n})^{\bar{p}(z)-1}) \\ &\geq \; -(2n)^{\bar{p}(z)-1} (\frac{1}{\delta_{\min}^n})^{\bar{p}(z)-1}\end{aligned}$$

Therefore we obtain the following regarding the mean payoff of π:

$$MP(r)(\pi) \geq (2n)^{\bar{p}(z)-1} (\frac{1}{\delta_{\min}^n})^{\bar{p}(z)-1} \geq 1$$

We denote all BSCCs in G_{σ^1,γ^*} with B_E, and obtain the following:

$$\begin{aligned}\langle E \rangle(MP(r))(v) &= \sup_{\sigma \in \Sigma} \mathbb{E}^v_{\sigma,\gamma^*}(MP(r)) && \textit{since } \gamma^* \textit{ is optimal for } MP(r) \\ &\geq \mathbb{E}^v_{\sigma^1,\gamma^*}(MP(r)) && \textit{since } \sigma^1 \in \Sigma \\ &\geq 1 \cdot \Pr^v_{\sigma^1,\gamma^*}(\mathit{Reach}(B_E)) + (-1) \cdot (1 - \Pr^v_{\sigma^1,\gamma^*}(\mathit{Reach}(B_E))) \\ & && \textit{follows from previous calculation} \\ &= 2\mathbb{P}^v_{\sigma^1,\gamma^*}(\mathit{Reach}(B_E)) - 1 \\ &= 2\mathbb{P}^v_{\sigma^1,\gamma^*}(PA(p)) - 1 && \textit{follows from Corollary 2.28} \\ &\geq 2\langle E \rangle(PA(p))(v) - 1 && \textit{since } \sigma^1 \textit{ is optimal for } PA(p)\end{aligned}$$

Hence we obtain the desired results.

The other half of the proof is to show that $\langle E \rangle(PA(p))(v) \geq \frac{1}{2}(\langle E \rangle(MP(r))(v) + 1)$. It is obtained in a similar way by fixing $\sigma^* \in \Sigma$ optimal for $MP(r)$ and $\gamma^1 \in \Gamma$ optimal for $PA(p)$. We skip the details here. □

5.1.2 Strategic SMPGs to Strategic SDPGs

We present the reduction from strategic SMPGs to strategic SDPGs proposed in [2].

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena with n vertices. Let $r : V \to \mathbb{R}$ be a rewarding function, and $\beta \in (0, 1]$ be a discount factor. We specify an SMPG $(G, MP(r))$ and an SDPG $(G, DP(r, \beta))$. We first adjust the reward function r such that for all vertices $v \in V$, $r(v) \in [0, 1)$. We refer to [46] for the details of this process.

We assume for all $u \in V_R, v \in V$, $\Delta(u, v)$ is a rational number, which can be written as the fraction of two integers bounded by M. We also assume for all $v \in V$, $r(v)$ is a rational number, which can be written as the fraction of two integers with absolute values bounded by M.

We claim that if $\beta \in (1 - \frac{1}{2c}, 1]$, where $c = n(n!)^4 M^{2n^2+2n}$, then a strategy $\sigma \in \Sigma$ of Eve is optimal in $(G, MP(r))$ if and only if it is optimal in $(G, DP(r, \beta))$.

Proof Sketch. We look at the SDPG. If we fix a strategy pair $\sigma, \gamma \in \Sigma \times \Gamma$, the expected outcome is determined, and recall that the expected value of a play starting from vertex v is denoted with $\mathbb{E}^v_{\sigma,\gamma}(DP(r, \beta))$. In the sequel, we leave out the objective function as it is clear in the context.

We show that for two arbitrary strategy pairs $(\sigma, \gamma), (\sigma', \gamma') \in \Sigma \times \Gamma$, the sign of $\mathbb{E}^v_{\sigma,\gamma} - \mathbb{E}^v_{\sigma',\gamma'}$ remains the same for all $\beta \in (1 - \frac{1}{2c}, 1]$, and thus the optimal strategies remain the same as well.

We denote with P the transition matrix of the Markov view of $G_{\sigma,\gamma}$. We can calculate $\mathbb{E}^v_{\sigma,\gamma}$ by solving the following linear equation:

$$(I - \beta P)x = r \tag{5.1}$$

where x is the vector containing $\mathbb{E}^v_{\sigma,\gamma}$ for all vertices $v \in V$. Let $Q = I - \beta P$. We introduce some notations for clarity in the following steps.

- We denote the entry of a matrix, say A, at i-th row and j-th column with $A[i, j]$. We denote the i-th entry of a vector, say v, with $v[i]$.
- $P[i, j]$ can be written as $P[i, j] = \frac{c_{i,j}}{d_{i,j}}$, where$|c_{i,j}|$ and $|d_{i,j}|$ are both natural numbers bounded by M, and $|c_{i,j}| < |d_{i,j}|$.
- $r[i]$ can be written as $r[i] = \frac{c_{i,n+1}}{d_{i,n+1}}$, where $|c_{i,n+1}|$ and $|d_{i,n+1}|$ are both natural numbers bounded by M, and $|c_{i,s+1}| < |d_{i,s+1}|$.
- Let $\beta_f = 1 - \beta$. The subscript stands for *flip*.

Now we look at entries in the equation system $Qx = r$.

- For all entries on the diagonal of Q, namely for all $i = 1, 2, \cdots, n$, we have $Q[i, i] = \frac{c_{i,i}}{d_{i,i}}\beta_f + \frac{d_{i,i} - c_{i,i}}{d_{i,i}}$.
- For all other entries of Q, we have $Q[i, j] = \frac{c_{i,j}}{d_{i,j}}\beta_f - \frac{c_{i,j}}{d_{i,j}}$.
- For all entries of r, we have $r[i] = 0\beta_f - \frac{c_{i,n+1}}{d_{i,n+1}}$

Then for all $i = 1, 2, \cdots, n$, we multiply the i-th row with $\prod_{t=1}^{n+1} d_{i,t}$. In this way, we get an equivalent equation system, and we write it as:

$$P'x = r'$$

It follows from the previous transformations that all entries of P' are of the form $P'[i,j] = a_{i,j}\beta_f + b_{i,j}$, where both $a_{i,j}$ and $b_{i,j}$ have absolute values bounded by M^{n+1}. The same applies to all entries of r'.

According to Cramer's rule, we obtain that:

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

where P'_i is the matrix formed by replacing the i-th column of P' with the column vector r'. It follows from the calculation of determinants that both $det(P'_i)$ and $det(P)$ are polynomials in β_f of degree n, and the coefficients are of absolute values at most:

$$n! \cdot \binom{n}{\lceil n/2 \rceil} \cdot (M^{n+1})^n \leq (n!)^2 M^{n(n+1)}$$

We write $x[i]$ as p_1/p_2, and we apply the same calculation for the other strategy pair (σ', γ') to obtain $x'[i] = p_3/p_4$. The sign of $p_1/p_2 - p_3/p_4$ is the same as the sign of $p_1p_4 - p_2p_3$, which is a polynomial of β_f with at most $2n$ degrees, written as follows:

$$p_1p_4 - p_2p_3 = c_{2n}\beta_f^{2n} + c_{2n-1}\beta_f^{2n-1} + \cdots + c_1\beta_f^1 + c_0 \tag{5.2}$$

where the integer coefficients c_i for integer $i \in [0, 2n]$ has absolute values at most:

$$n \cdot ((n!)^2 M^{n(n+1)})^2 = n(n!)^4 M^{2n^2+2n} = c \tag{5.3}$$

When $\beta > 1 - \frac{1}{2c}$, the following holds:

(a) $c < \frac{1}{2\beta_f}$

(b) $1 - \beta_f > \frac{1}{2}$

We claim that the sign of $p_1p_4 - p_2p_3$ is the same as c_k, where k is the smallest subscript so that $c_k \neq 0$. The reasons are as follows:

- Regarding the absolute value of $c_k\beta_f^k$, we have $|c_k\beta_f^k| \geq \beta_f^k$.
- Regarding the absolute value of what comes after $c_k\beta_f^k$, we have:

$$\begin{aligned}
&|\sum_{i=k+1}^{2n}(c_i\beta_f^i)| \leq \sum_{i=k+1}^{2n}|c_i\beta_f^i| \leq c\sum_{i=k+1}^{2n}|\beta_f^i| \\
&< \frac{1}{2\beta_f} \cdot \beta_f^{k+1} \cdot \frac{1-\beta_f^{2n-k}}{1-\beta_f} && \textit{follows from (a)} \\
&= \beta_f^k \cdot \frac{1-\beta_f^{2n-k}}{2(1-\beta_f)} \\
&< \beta_f^k \cdot (1-\beta_f^{2n-k}) < \beta_f^k && \textit{follows from (b)}
\end{aligned}$$

We thus conclude that when $\beta > 1 - \frac{1}{2c}$, the sign of $\mathbb{E}^v_{\sigma,\gamma} - \mathbb{E}^v_{\sigma',\gamma'}$ remains the same, and the proof is complete. □

5.1.3 Quantitative SDPGs to Quantitative SSGs

We present the reduction from quantitative SDPGs to quantitative SSGs proposed also in [2].

Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, and $(G, DP(r, \beta))$ be an SDPG with reward function r and discount factor β. We first adjust the reward function r such that for all vertices $v \in V$, $r(v) \in [0, 1)$. We refer to [46] for the details of this process.

We construct a stochastic arena G' as given by Figure 5.2:

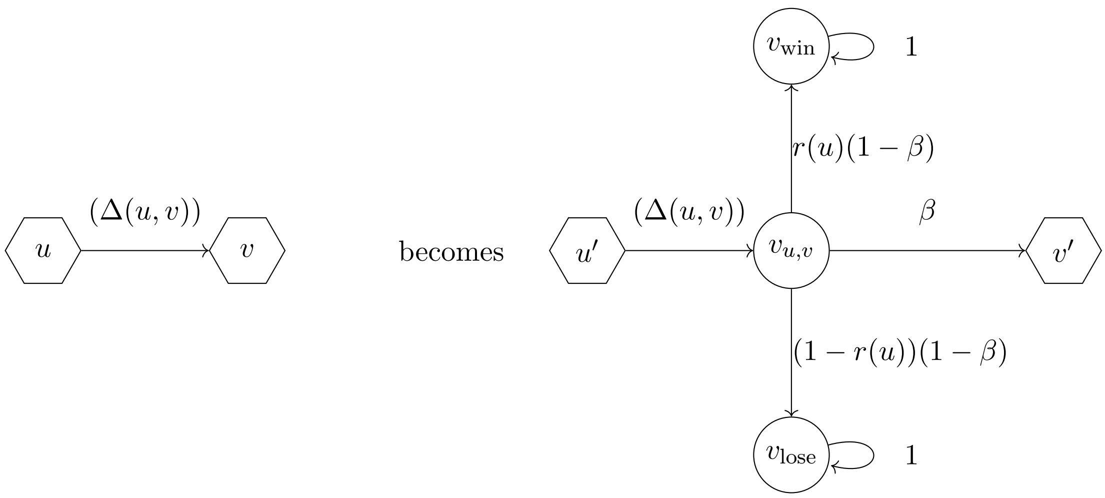

Figure 5.2: The gadget for reducing quantitative SDPGs to quantitative SSGs

Formally, we write $G' = ((V' \uplus V'_E \uplus \{v_{\text{win}}, v_{lose}\}, E'), (V'_\exists, V'_\forall, V'_R), \Delta')$ as follows:

$$\begin{aligned}
V' =& \{v' \mid v \in V\},\ V'_E = \{v_{(u,w)} \mid (u, w) \in E\} \\
V'_\exists =& \{v' \mid v \in V_\exists\},\ V'_\forall = \{v' \mid v \in V_\forall\} \\
V'_R =& \{v' \mid v \in V_R\} \uplus V'_E \uplus \{v_{\text{win}}, v_{lose}\} \\
E' =& \{(u', v_{(u,w)}), (v_{(u,w)}, w'), (v_{(u,w)}, v_{\text{win}}), (v_{(u,w)}, v_{lose}) \mid (u, w) \in E\}
\end{aligned}$$

We define the new transition function Δ' as follows:

- For all random vertices $u \in V_R$, we have:

$$\begin{aligned}
&\Delta'(u', v_{(u,w)}) = \Delta(u, w) \\
&\Delta'(v_{(u,w)}, w') = \beta \\
&\Delta'(v_{(u,w)}, v_{\text{win}}) = r(u)(1 - \beta) \\
&\Delta'(v_{(u,w)}, v_{lose}) = (1 - r(u))(1 - \beta)
\end{aligned}$$

- The sinks have only self-loops, namely $\Delta'(v_{\text{win}}, v_{\text{win}}) = 1$ and $\Delta'(v_{lose}, v_{lose}) = 1$.

For all $v \in V$, we claim that the value of v in $(G, DP(r, \beta))$ is the same as the value of v' in $(G', Reach(v_{\text{win}}))$.

Proof Sketch. We first look at the SDPG. For all vertices $u \in V$, we denote its value in $(G, DP(r, \beta))$ with x_u. The values satisfy the following equation system:

$$x_u = \begin{cases} (1-\beta)r(u) + \beta \max\limits_{(u,w)\in E}\{x_w\} & \text{if } u \in V_\exists \\ (1-\beta)r(u) + \beta \min\limits_{(u,w)\in E}\{x_w\} & \text{if } u \in V_\forall \\ (1-\beta)r(u) + \beta \sum\limits_{(u,w)\in E} \Delta(u,w)x_w & \text{if } u \in V_R \end{cases} \tag{5.4}$$

We then look at the SSG. For all vertices $u' \in V' \uplus V'_E$, we denote its value in $(G', Reach(v_{\text{win}}))$ with $x'_{u'}$. The values satisfy the following equation system:

$$x'_{u'} = \begin{cases} \max\limits_{(u',v_{(u,w)})\in E'}\{x_{v_{(u,w)}}\} & \text{if } u' \in V'_\exists \\ \min\limits_{(u',v_{(u,w)})\in E'}\{x_{v_{(u,w)}}\} & \text{if } u' \in V'_\forall \\ \sum\limits_{(u',v_{(u,w)})\in E'} \Delta'(u', v_{(u,w)}) \cdot x_{v_{(u,w)}} & \text{if } u' \in V'_R \end{cases} \tag{5.5}$$

$$x_{v_{(u,w)}} = (1-\beta)r(u) + \beta x_{w'} \quad \text{if } v_{(u,w)} \in V'_E \tag{5.6}$$

It follows from the construction of G' that:

- $(u', v_{(u,w)}) \in E'$ and $v_{(u,w)} \in V'_E$ if and only if $(u,w) \in E$.
- For all edges $(u,w) \in E$, $\Delta'(u', v_{(u,w)}) = \Delta(u,w)$.

Therefore we rewrite the previous equation system 5.5 and 5.6 as follows:

$$x'_{u'} = \begin{cases} \max\limits_{(u,w)\in E}\{x_{v_{(u,w)}}\} & \text{if } u' \in V'_\exists, \\ \min\limits_{(u,w)\in E}\{x_{v_{(u,w)}}\} & \text{if } u' \in V'_\forall, \\ \sum\limits_{(u,w)\in E} \Delta(u,w) \cdot x_{v_{(u,w)}} & \text{if } u' \in V'_R \end{cases} \tag{5.7}$$

$$x_{v_{(u,w)}} = (1-\beta)r(u) + \beta x_{w'} \quad \text{if } v_{(u,w)} \in V'_E \tag{5.8}$$

We substitute $x_{v_{(u,w)}}$ in the first three equations 5.7 according to 5.8, and obtain the following equivalent equation system:

$$x'_{u'} = \begin{cases} (1-\beta)r(u) + \beta \max_{(u,w)\in E} \{x_{w'}\} & \text{if } u' \in V'_{\exists}, \\ (1-\beta)r(u) + \beta \min_{(u,w)\in E} \{x_{w'}\} & \text{if } u' \in V'_{\forall}, \\ (1-\beta)r(u) + \beta \sum_{(u,w)\in E} \Delta(u,w)x_{w'} & \text{if } u' \in V'_R \end{cases} \tag{5.9}$$

We observe that the equation systems 5.4 and 5.9 are identical. We recall Definition 2.14, where we assume all arenas to be non-blocking. From the construction of G', we claim that if G is non-blocking, G' is *stopping*, which means there is a path from every vertex to the winning or losing sinks. As proved in [15], the equation system 5.9 has a unique solution since G' is stopping. We thus conclude the proof. □

5.2 Reducing Quantitative SPGs Directly to Quantitative SSGs

We recall Figure 5.1. In this section, we present the reductions in the lower half. The reductions drawn in dotted arrows are trivial for the following reasons, and we skip further details.

- From quantitative SPGs to strategic SPGs: once the strategies have been fixed, the values can be obtained by a complete analysis of the resulting Markov Chain.
- From strategic SSGs to quantitative SSGs: once the values are fixed, a strategy of Eve is optimal if and only if in each vertex of Eve it chooses an successor with the maximum value.

In the sequel, we present the reduction from strategic SPGs to strategic SSGs, using the same gadget from Section 4.2.1.

5.2.1 Two Useful Lemmas

Given an SPG $(G, PA(p))$, we construct an SSG $(\overline{G}, Reach(v_{\mathrm{win}}))$ as in Section 4.2.1. For all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, we define in $G_{\sigma,\gamma}$ the probability for a play from v to enter an even BSCC as $\mathrm{Pr}^v_{\sigma,\gamma}(enterEven)$. We define $\mathrm{Pr}^v_{\sigma,\gamma}(enterOdd)$ analogously.

First, we establish a connection between the winning probability in both SGs as follows.

Lemma 5.2 (Range of Values in the SSG)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\mathrm{win}}))$ as in Section 4.2.1.
If for some $x \in (0, 1)$ and $y \in (0, 1)$, $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(crossPath) > x$ and $\overline{\mathrm{Pr}}^{\min}_{\sigma,\gamma}(winEven) \geq y$, then the follwoing holds:

$$\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(Reach(v_{\mathrm{win}})) \in [y\mathrm{Pr}^v_{\sigma,\gamma}(enterEven) - y + xy, \mathrm{Pr}^v_{\sigma,\gamma}(enterEven) + 1 - xy] \qquad \blacksquare$$

Proof. For all vertices $v \in V$, we make the following observations:

(1) $\mathrm{Pr}^v_{\sigma,\gamma}(enterEven) + \mathrm{Pr}^v_{\sigma,\gamma}(enterOdd) = 1$

(2) $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(enterEven) + \overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(enterOdd) = \overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(crossPath)$

(3) $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(enterEven) \leq \mathrm{Pr}^v_{\sigma,\gamma}(enterEven)$ and $\overline{\mathrm{Pr}}^{\overline{v}}_{\sigma,\gamma}(enterOdd) \leq \mathrm{Pr}^v_{\sigma,\gamma}(enterOdd)$

(4) Regarding $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven})$, we obtain the following:

$$
\begin{aligned}
& \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven}) \\
= & \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterOdd}) && \text{follows from (2)} \\
\geq & \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) - \Pr^{v}_{\sigma,\gamma}(\mathit{enterOdd}) && \text{follows from (3)} \\
= & \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) + \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) - 1 && \text{follows from (1)}
\end{aligned}
$$

We thus obtain a lower bound of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$ as follows:

$$
\begin{aligned}
& \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) \\
\geq & \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven}) \cdot \overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) \\
\geq & (\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) + \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) - 1) \cdot \overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) && \textit{follows from (4)} \\
\geq & (x + \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) - 1) \cdot y && \textit{follows from assumptions } (C_0) \textit{ and } (C_1) \\
= & y\Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) - y + xy
\end{aligned}
$$

Similarly, we obtain an upper bound of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$ as follows:

$$
\begin{aligned}
& \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) \\
< & (1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})) + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterEven}) \cdot 1 + \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{enterOdd}) \cdot \overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) \\
\leq & (1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})) + \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) \cdot 1 + \\
& (\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) - \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven})) \cdot \overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) && \textit{follows from (2)} \\
= & (1 - \overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd})) \cdot \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) + 1 - \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath})(1 - \overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd})) \\
\leq & \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) + 1 - xy && \textit{follows from assumptions } (C_0) \textit{ and } (C_1)
\end{aligned}
$$

Therefore we get a range of $\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$ as follows:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}})) \in [y\Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) - y + xy, \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven}) + 1 - xy] \qquad \square$$

Remark 5.3 (Winning Probabilities As Reachability)
We recall Corollary 2.29, and the following statements follow:

- In the SPG $(G, \mathit{PA}(p))$, $\mathbb{P}^{v}_{\sigma,\gamma}(\mathit{PA}(p)) = \Pr^{v}_{\sigma,\gamma}(\mathit{enterEven})$.
- In the SSG $(\overline{G}, \mathit{RE}(v_{\text{win}}))$, $\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma}(\mathit{RE}(v_{\text{win}})) = \overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{Reach}(v_{\text{win}}))$.

The objectives are clear under the context, so we leave them out. We write Lemma 5.2 as follows:

$$\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} \in [y\mathbb{P}^{v}_{\sigma,\gamma} - y + xy, \mathbb{P}^{v}_{\sigma,\gamma} + 1 - xy] \qquad \blacksquare$$

Intuitively, with a fixed strategy pair, the value of the SSG falls into a range around the value of the SPG, and the range size depends on *crossPath* and minimum *winEven* probabilities.

Second, given a vertex $v \in V$, we show in $(G, PA(p))$ the minimum gap between two unequal winning probabilities from v.

Lemma 5.4 (Difference Between Values Under Different Strategy Pairs) Let $G = ((V,E),(V_\exists, V_\forall, V_R),\Delta)$ be a stochastic arena, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We assume for all $u \in V_R, v \in V$, $\Delta(u,v)$ is a rational number, which can be written as the fraction of two integers $\frac{a_{u,v}}{b_{u,v}}$, where $a_{u,v} \in \mathbb{N}, b_{u,v} \in \mathbb{N}$. There exists a constant $M \in \mathbb{N}_+$, such that for all $u \in V_R, v \in V$, $M \geq b_{u,v} \geq a_{u,v}$.
Let the cardinality of V be n. For all $\sigma, \gamma \in \Sigma \times \Gamma$ and $\sigma', \gamma' \in \Sigma \times \Gamma$, and for all vertices $v \in V$, if $\mathbb{P}_{\sigma,\gamma} > \mathbb{P}_{\sigma',\gamma'}$, then the following holds:

$$\mathbb{P}^v_{\sigma,\gamma} - \mathbb{P}^v_{\sigma',\gamma'} > \frac{1}{(n!)^2 M^{2n^2}}$$

■

Proof. It follows from Remark 5.3 that for all $\sigma, \gamma \in \Sigma \times \Gamma$ and $v \in V$, we have $\mathbb{P}^v_{\sigma,\gamma} = \mathrm{Pr}^v_{\sigma,\gamma}(\mathit{enterEven})$. We can obtain $\mathrm{Pr}^v_{\sigma,\gamma}(\mathit{enterEven})$ by setting all vertices in even BSCCs as the target set, and calculating the reachability probability. Calculating $\mathbb{P}^v_{\sigma,\gamma}$ is thus reduced to solving a linear equation system $(I - A)x = b$ according to Theorem 2.8. We do not repeat the details of obtaining A and b here. 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:

- We denote the number of equations with s. It follows that $s < n$ since there is at least one vertex in a BSCC.
- Let $Q = I - A$. We denote the i-th row of Q with $Q[i]$, and we denote 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 both natural numbers bounded by M, and $|c_{i,j}| \leq |d_{i,j}|$.
- 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 both natural numbers bounded by M, and it also holds that $|c_{i,s+1}| \leq |d_{i,s+1}|$.
- We denote the i-th entry of x with $x[i]$.

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:

$$\begin{bmatrix} \frac{c_{i,1}}{d_{i,1}} & \frac{c_{i,2}}{d_{i,2}} & \cdots & \frac{c_{i,s}}{d_{i,s}} \end{bmatrix} \cdot x = \frac{c_{i,s+1}}{d_{i,s+1}} \tag{5.10}$$

Then we multiply equation 5.10 with $\prod_{t=1}^{s+1} d_{i,t}$, and the effect are as follows:

- For all $j = 1, 2, \cdots, s$, $Q[i, j]$ becomes $(\prod_{t=1}^{s+1} d_{i,t})\frac{c_{i,j}}{d_{i,j}}$, which are integers with absolute values bounded by M^{s+1}.
- For all $i = 1, 2, \cdots, s$, $b[i]$ now becomes $(\prod_{t=1}^{s} d_{i,t})c_{i,s+1}$, which is also 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:

$$Q'x = b'$$

Thus all entries of Q' and b' are integers with absolute values bounded by M^{s+1}.

With Cramer's rule, for all $i = 1, 2, \cdots, s$, we obtain that:

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

where Q'_i is the matrix formed 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 reachability probability, it holds that $x[i] \leq 1$. 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 $x'[i] > x[i]$, then the following holds:

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

□

5.2.2 The Strategic Reduction

We now present the direct reduction from strategic SPGs to strategic SSGs.

Theorem 5.5 (Reducing Strategic SPGs to Strategic SSGs)
Let $G = ((V, E), (V_\exists, V_\forall, V_R), \Delta)$ be a stochastic arena with n vertices, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. We construct an SSG $(\overline{G}, RE(v_{\text{win}}))$ as in Section 4.2.1.

We assume for all $u \in V_R, v \in V$, $\Delta(u, v)$ is a rational number, which can be written as the fraction of two integers $\frac{a_{u,v}}{b_{u,v}}$, where $a_{u,v} \in \mathbb{N}, b_{u,v} \in \mathbb{N}$. There exists a constant $M \in \mathbb{N}_+$, such that for all $u \in V_R, v \in V$, $M \geq b_{u,v} \geq a_{u,v}$.
Let $\epsilon = \frac{1}{(n!)^2 M^{2n^2}}$. If for all $\sigma \in \Sigma$, $\gamma \in \Gamma$ and $v \in V$, the following conditions hold:

(C_0) $\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(crossPath) > \frac{4-\epsilon}{4}$

(C_1) $\overline{\Pr}^{\min}_{\sigma,\gamma}(\mathit{winEven}) \geq \frac{4}{4+\epsilon}$, $\overline{\Pr}^{\max}_{\sigma,\gamma}(\mathit{winOdd}) \leq 1 - \frac{4}{4+\epsilon}$

then a strategy $\sigma \in \Sigma$ of Eve is optimal in the SPG $(G, PA(p))$ if and only if σ is optimal in the SSG $(\overline{G}, RE(v_{\text{win}}))$. The same holds for Adam. ■

To give an intuition of this reduction, we first recall Lemma 5.2 and Lemma 5.4 here. When conditions (C_0) and (C_1) are satisfied in the SSG, the value ranges of different strategy pairs become sufficiently small so that they do not overlap, and hence the order of values is the same in both games. As a result, an optimal strategy for one game is also optimal in the other game. We present the formal proof as follows.

Proof. We assume conditions (C_0) and (C_1) hold. In the sequel, we show that a strategy $\sigma \in \Sigma$ of Eve is optimal in the SPG $(G, PA(p))$ if and only if σ is optimal in the SSG $(\overline{G}, \mathit{Reach}(v_{\text{win}}))$.

- We first show that an optimal strategy $\sigma \in \Sigma$ in the SSG is also optimal in the SPG. We make the proof by contradiction.

  We assume $\sigma \in \Sigma$ and $\gamma \in \Gamma$ are the optimal strategies of Eve and Adam in $(\overline{G}, \mathit{Reach}(v_{\text{win}}))$. Since assumptions ($C_0$) and ($C_1$) hold, it follows from Lemma 5.2 that for all $v \in V$, the following holds:

  $$\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} \in [y\mathbb{P}^{v}_{\sigma,\gamma} - y + xy, \mathbb{P}^{v}_{\sigma,\gamma} + 1 - xy]$$

  Substituting x and y, we obtain that:

  $$\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} \leq \mathbb{P}^{v}_{\sigma,\gamma} + \frac{2\epsilon}{4+\epsilon} \tag{5.11}$$

  If σ is not optimal in $(G, PA(p))$, there exists another strategy $\sigma' \in \Sigma$ and a vertex $v \in V$ such that $\mathbb{P}^{v}_{\sigma',\gamma} > \mathbb{P}^{v}_{\sigma,\gamma}$. It follows again from Lemma 5.2 that:

  $$\overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma} \in [y\mathbb{P}^{v}_{\sigma',\gamma} - y + xy, \mathbb{P}^{v}_{\sigma',\gamma} + 1 - xy] \tag{5.12}$$

  Furthermore, it follows from Lemma 5.4 that:

  $$\mathbb{P}^{v}_{\sigma',\gamma} > \mathbb{P}^{v}_{\sigma,\gamma} + \epsilon \tag{5.13}$$

As a result, we obtain the following:

$$
\begin{aligned}
&\overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma} \\
&\geq y\mathbb{P}^{v}_{\sigma',\gamma} - y + xy && \textit{follows from 5.12} \\
&> y(\mathbb{P}^{v}_{\sigma,\gamma} + \epsilon) - y + xy && \textit{follows from 5.13} \\
&= \frac{4}{4+\epsilon}(\mathbb{P}^{v}_{\sigma,\gamma} + \epsilon) - \frac{4}{4+\epsilon} \cdot \frac{\epsilon}{4} \\
&= \frac{4}{4+\epsilon}\mathbb{P}^{v}_{\sigma,\gamma} + \frac{3\epsilon}{4+\epsilon} \\
&= \mathbb{P}^{v}_{\sigma,\gamma} - \frac{\epsilon}{4+\epsilon}\mathbb{P}^{v}_{\sigma,\gamma} + \frac{3\epsilon}{4+\epsilon} \\
&\geq \mathbb{P}^{v}_{\sigma,\gamma} + \frac{2\epsilon}{4+\epsilon} \\
&\geq \overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} && \textit{follows from 5.11}
\end{aligned}
$$

It indicates that $\overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma} > \overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma}$, which contradicts the assumption that σ and γ are optimal in $(\overline{G}, \mathit{Reach}(v_{\text{win}}))$.

- For the other direction, we show that an optimal strategy $\sigma \in \Sigma$ in the SPG is also optimal in the SSG.

  We assume $\sigma \in \Sigma$ is an optimal strategy of Eve in $(G, \mathit{PA}(p))$. We take an arbitrary strategy $\gamma \in \Gamma$ of Adam. For all non-optimal strategies $\sigma' \in \Sigma$, there exists $v \in V$, such that $\mathbb{P}^{v}_{\sigma,\gamma} > \mathbb{P}^{v}_{\sigma',\gamma}$, and thus it follows from Lemma 5.4 that

$$\mathbb{P}^{v}_{\sigma,\gamma} > \mathbb{P}^{v}_{\sigma',\gamma} + \epsilon \tag{5.14}$$

  We first obtain the following for $\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma}$:

$$
\begin{aligned}
\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} &\geq y\mathbb{P}^{v}_{\sigma,\gamma} - y + xy && \textit{follows from Lemma 5.2} \\
&= \frac{4}{4+\epsilon}\mathbb{P}^{v}_{\sigma,\gamma} - \frac{4}{4+\epsilon} + \frac{4-\epsilon}{4+\epsilon} \\
&= \frac{4}{4+\epsilon}\mathbb{P}^{v}_{\sigma,\gamma} - \frac{\epsilon}{4+\epsilon} \\
&= \mathbb{P}^{v}_{\sigma,\gamma} - \frac{\epsilon}{4+\epsilon}(\mathbb{P}^{v}_{\sigma,\gamma} + 1) \\
&\geq \mathbb{P}^{v}_{\sigma,\gamma} - \frac{2\epsilon}{4+\epsilon} && \textit{since } \mathbb{P}^{v}_{\sigma,\gamma} < 1
\end{aligned}
$$

In the meantime, we obtain the following for $\overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma}$:

$$\begin{aligned}
\overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma} &\leq \mathbb{P}^{v}_{\sigma',\gamma} + 1 - xy && \textit{follows from Lemma 5.2}\\
&< \mathbb{P}^{v}_{\sigma,\gamma} - \epsilon + 1 - xy && \textit{follows from 5.14}\\
&= \mathbb{P}^{v}_{\sigma,\gamma} - \epsilon + 1 - \frac{4-\epsilon}{4+\epsilon}\\
&= \mathbb{P}^{v}_{\sigma,\gamma} - \frac{2\epsilon+\epsilon^2}{4+\epsilon}\\
&< \overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma}
\end{aligned}$$

Thus we have $\overline{\mathbb{P}}^{\overline{v}}_{\sigma,\gamma} > \overline{\mathbb{P}}^{\overline{v}}_{\sigma',\gamma}$, and it follows that σ is an optimal strategy of Eve in $(\overline{G}, \mathit{Reach}(v_{\mathrm{win}}))$. □

5.2.3 Arranging α

In the sequel, we show how to satisfy conditions (C_0) and (C_1) by arranging α properly, which is similar to Section 4.3.1 and Section 4.3.2.

Arranging $\alpha(0)$

We write $\alpha(0)$ as α for visual neatness. It follows from Corollary 4.18 that:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > \frac{\delta^n_{\min}(1-\alpha)^{n+1}}{2\alpha + \delta^n_{\min}(1-\alpha)^{n+1}}$$

Therefore to show:

$$\overline{\Pr}^{\overline{v}}_{\sigma,\gamma}(\mathit{crossPath}) > \frac{4-\epsilon}{4} \tag{5.15}$$

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

$$\frac{\delta^n_{\min}(1-\alpha)^{n+1}}{2\alpha + \delta^n_{\min}(1-\alpha)^{n+1}} \geq \frac{4-\epsilon}{4} \tag{5.16}$$

which can be further simplified as:

$$\epsilon \geq \frac{8\alpha}{2\alpha + \delta^n_{\min}(1-\alpha)^{n+1}} \tag{5.17}$$

We show that when $\alpha \leq \frac{\delta_{\min}^n}{8(n!)^2 M^{2n^2}}$, inequality 5.17 holds. We start with the right side:

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

Therefore we obtain that when $\alpha(0) \leq \frac{\delta_{\min}^n}{8(n!)^2 M^{2n^2}}$, inequality 5.15 holds, and thus condition (C_0) is satisfied.

Arranging $\alpha(k+1)/\alpha(k)$

We write $\alpha(k)$ and $\alpha(k+1)$ as α_k and α_{k+1} respectively for visual neatness. It follows from Corollary 4.23 that for all strategy pairs $\sigma, \gamma \in \Sigma \times \Gamma$, for all $\overline{v} \in \overline{V}$ in an even rBSCC in $\overline{G}_{\sigma,\gamma}$ with smallest priority k, the following holds:

$$\overline{\Pr}_{\sigma,\gamma}^{\overline{v}}(Reach(v_{\text{win}})) > \frac{\delta_{\min}^n(1-\delta_{\min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{\min}^n(1-\delta_{\min}) + \frac{\alpha_{k+1}}{\alpha_k}} \tag{5.18}$$

Therefore to show:

$$\overline{\Pr}_{\sigma,\gamma}^{\min}(winEven) \geq \frac{4}{4+\epsilon} \tag{5.19}$$

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

$$\frac{\delta_{\min}^n(1-\delta_{\min}) - \frac{\alpha_{k+1}}{\alpha_k}}{\delta_{\min}^n(1-\delta_{\min}) + \frac{\alpha_{k+1}}{\alpha_k}} \geq \frac{4}{4+\epsilon} \tag{5.20}$$

We denote $\frac{\alpha_{k+1}}{\alpha_k}$ with r in the following calculation. Inequality 5.20 can be further simplified to:

$$\frac{\epsilon}{4+\epsilon} \geq \frac{2r}{(1-\delta_{\min})\delta_{\min}^n + r} \tag{5.21}$$

and can be finally simplified to:

$$r \leq \frac{(1-\delta_{\min})\delta_{\min}^{n}}{8(n!)^2 M^{2n^2}+1} \tag{5.22}$$

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 5.19 holds, and thus condition (C_1) is satisfied.

Complexity Results

Let $G = ((V,E),(V_\exists, V_\forall, V_R),\Delta)$ be a stochastic arena where $|V| = n$, $p : V \to \mathbb{N}$ be a priority function, and $(G, PA(p))$ be an SPG. For all $u, v \in V_R \times V$, the value of $\Delta(u,v)$ can be written as a fraction of positive integers bounded by M.

In both qualitative and quantitative reductions, we construct an SSG $(\overline{G}, RE(v_{\text{win}}))$, where arena $\overline{G}$ is formally written as:

$$\overline{G} = ((\overline{V} \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}, \overline{E}), (\overline{V}_\exists, \overline{V}_\forall, \overline{V}_R \uplus \widehat{V} \uplus \{v_{\text{win}}, v_{\text{lose}}\}), \overline{\Delta})$$

We recall the qualitative reduction, where $\alpha(k) = \delta_{\min}^{(2k+1)(2n+3)}$ for all $k \in \mathbb{N}$. Under unary encoding, an exponential blow-up of the arena size is trivially observable in both qualitative and quantitative cases, due to the values of α.

However, under binary encoding, the reductions are polynomial. We only show the quantitative case, since the size of α is trivially smaller in the qualitative case.

Since $\delta_{\min} \geq \frac{1}{M}$ and $1 - \delta_{\min} \geq \frac{1}{2}$, a valid instance of α is for all $k \in \mathbb{N}$:

$$\alpha(k) = \left(\frac{1}{16(n!)^2 M^{2n^2+n}+1}\right)^{k+1}$$

Then the size of $\overline{G}$ is:

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

According to [2], quantitative SSGs under unary and binary encoding are in the same complexity class **NP** ∩ **coNP**. We thus obtain that both qualitative and quantitative SPGs under binary encoding are in **NP** ∩ **coNP** as well.

Chapter 6

Conclusion

6.1 Overview

In this thesis, we present direct reductions from both qualitative and quantitative SPGs to quantitative SSGs. We take inspiration from an existing reduction from DPGs to quantitative SSGs, whose gadget naturally extends to SPGs.

In Chapter 4, we discuss the qualitative case. With fixed strategies of both players, the key to an SPG is what happens inside BSCCs. In the newly constructed SSG, we call their counterparts rBSCCs, and we divide a play into two stages, namely before and after reaching an rBSCC.

We first identify an achievable lower bound of the probability for a play to reach an rBSCC (*crossPath*) through Markov Chain analysis. The lower bound is dependent on the size n and minimum transition probability $\delta_{\min}$ of the arena. We further show that when the reward function α is sufficiently small in terms of n and $\delta_{\min}$, the lower bound is sufficiently close to 1. Similarly, we find a tight lower bound of the minimum probability for Eve to win in an even rBSCC (*winEven*), which also depends on n and $\delta_{\min}$. When the ratios of α between large and small priorities are sufficiently small in terms of n and $\delta_{\min}$, we find the lower bound sufficiently close to 1. Due to the close-to-1 probability of reaching an rBSCC, and the dominance of the smallest priority in an rBSCC, the SSG is an approximate simulation of the SPG.

In Chapter 5, we discuss the quantitative case. With a fixed pair of strategies, the values of both the SPG and the SSG are determined, but the construction of the SSG makes it difficult to link them via equality.

For a fixed strategy pair, we show that the value of the SSG falls into a range around the value of the SPG, and the range is dependent on *crossPath* and minimum *winEven* probabilities. For different strategy pairs, we obtain a lower bound ϵ of their value differences in the SPG, by restricting transition probabilities to rational numbers and analyzing reachability equation systems of Markov Chains. We then show in the SSG that by arranging α properly in terms of n, $\delta_{\min}$, and ϵ, the value ranges of different strategy pairs can be narrowed so that they do not overlap. In this case, a strategic reduction from SPGs to SSGs is achieved. Since quantitative and strategic solutions of SPGs and SSGs are trivially interconvertible in polynomial time, this also yields a relatively direct quantitative reduction.

Although astronomical numbers can be introduced into the function α of the newly constructed SSGs, we finally showed that under binary encoding, both reductions are polynomial. This indicates that both qualitative and quantitative SSGs under binary encoding are in $\mathbf{NP} \cap \mathbf{coNP}$, substantiating the complexity results from previous works.

As the reductions are established, we obtain a theoretical solution to SPGs by first reducing them to SSGs and then applying algorithms for SSGs. However, it is generally not implementable due to the astronomical values of α.

6.2 Future Work

In our reductions, we assume the minimum transition probability to be $\delta_{\min} \in (0, \frac{1}{2}]$ in an SPG, but its value is 1 in a DPG. Therefore our reductions cannot capture DPGs to quantitative SSGs as a subcase. A future work direction is to bridge this gap.

Furthermore, the reductions cannot be leveraged to solve SPGs in general practice as discussed. We see the following directions for future work:

- We have not examined formally the optimal arrangement of α for potential implementation. It is possible to find the weakest requirements on α so that the reductions are correct, and to do optimization in this direction.
- We have identified some limit cases that require the values of α to be astronomical. It is possible to find subsets of SPGs, where the limit cases are avoided, or the arenas are structured in favor of implementable valuation of α.
- It is possible to do experiments on simple benchmarks, for example, stochastic Büchi games with limited sizes.

Bibliography

[1] R. Amir. Stochastic games in economics and related fields: an overview. *Stochastic games and applications*, pages 455–470, 2003. (Cited on page 2.)

[2] D. Andersson and P. B. Miltersen. The complexity of solving stochastic games on graphs. In *International Symposium on Algorithms and Computation*, pages 112–121. Springer, 2009. (Cited on pages 2, 4, 5, 70, 72, 75, and 86.)

[3] M. Azeem, A. Evangelidis, J. Křetínskỳ, A. Slivinskiy, and M. 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. (Cited on page 3.)

[4] C. Baier and J.-P. Katoen. *Principles of model checking*. MIT press, 2008. (Cited on pages 4, 9, 10, and 12.)

[5] C. Baier, J. Klein, L. Leuschner, D. Parker, and S. Wunderlich. Ensuring the reliability of your model checker: Interval iteration for markov decision processes. In *International Conference on Computer Aided Verification*, pages 160–180. Springer, 2017. (Cited on page 2.)

[6] C. S. Calude, S. Jain, B. Khoussainov, W. Li, and F. Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 252–263, 2017. (Cited on page 4.)

[7] K. Chatterjee and N. Fijalkow. A reduction from parity games to simple stochastic games. *Electronic Proceedings in Theoretical Computer Science*, 54:74–86, jun 2011. (Cited on pages 2, 4, 5, 20, 21, 34, and 36.)

[8] K. Chatterjee and T. 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. (Cited on page 4.)

[9] K. Chatterjee and T. A. Henzinger. Reduction of stochastic parity to stochastic mean-payoff games. *Information Processing Letters*, 106(1):1–7, 2008. (Cited on pages 5, 70, and 71.)

[10] K. Chatterjee and T. A. Henzinger. A survey of stochastic ω-regular games. *Journal of Computer and System Sciences*, 78(2):394–413, 2012. (Cited on page 3.)

[11] K. Chatterjee, T. A. Henzinger, B. Jobstmann, and A. Radhakrishna. Gist: A solver for probabilistic games. In *Computer Aided Verification: 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings 22*, pages 665–669. Springer, 2010. (Cited on page 4.)

[12] K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Simple stochastic parity games. In M. Baaz and J. A. Makowsky, editors, *Computer Science Logic*, pages 100–113, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. (Cited on pages 4 and 5.)

[13] K. Chatterjee, M. Jurdzinski, and T. A. Henzinger. Quantitative stochastic parity games. In *SODA*, volume 4, pages 121–130, 2004. (Cited on pages 4, 5, 28, and 33.)

[14] A. Condon. On algorithms for simple stochastic games. *Advances in computational complexity theory*, 13:51–72, 1990. (Cited on page 2.)

[15] A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992. (Cited on pages 4 and 77.)

[16] W. Czerwiński, L. Daviaud, N. Fijalkow, M. Jurdziński, R. Lazić, and P. 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. (Cited on page 4.)

[17] M. Darlington, K. D. Glazebrook, D. S. Leslie, R. Shone, and R. Szechtman. A stochastic game framework for patrolling a border. *European Journal of Operational Research*, 2023. (Cited on page 2.)

[18] L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *Journal of Computer and System Sciences*, 68(2):374–397, 2004. (Cited on page 3.)

[19] J. Eisentraut, E. Kelmendi, J. Křetínskỳ, and M. Weininger. Value iteration for simple stochastic games: Stopping criterion and learning algorithm. *Information and Computation*, 285:104886, 2022. (Cited on page 3.)

[20] E. A. Emerson, C. S. Jutla, and A. P. Sistla. On model-checking for fragments of μ-calculus. In *Computer Aided Verification: 5th International Conference, CAV'93 Elounda, Greece, June 28–July 1, 1993 Proceedings 5*, pages 385–396. Springer, 1993. (Cited on page 3.)

[21] E. M. Hahn, H. Hermanns, and L. Zhang. Probabilistic reachability for parametric markov models. *International Journal on Software Tools for Technology Transfer*, 13:3–19, 2011. (Cited on page 49.)

[22] E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. Model-free reinforcement learning for stochastic parity games. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020. (Cited on page 4.)

[23] A. Hartmanns and B. L. Kaminski. Optimistic value iteration. In *International Conference on Computer Aided Verification*, pages 488–511. Springer, 2020. (Cited on page 3.)

[24] C. Hensel, S. Junges, J.-P. Katoen, T. Quatmann, and M. Volk. The probabilistic model checker storm. *International Journal on Software Tools for Technology Transfer*, pages 1–22, 2021. (Cited on page 3.)

[25] M. Huth and M. Kwiatkowska. Quantitative analysis and model checking. In *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*, pages 111–122. IEEE, 1997. (Cited on page 3.)

[26] M. Jurdziński. Deciding the winner in parity games is in up$\cup$co-up. *Information Processing Letters*, 68(3):119–124, 1998. (Cited on pages 2 and 70.)

[27] M. Jurdziński and R. 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. (Cited on page 4.)

[28] M. Kattenbelt, M. Kwiatkowska, G. Norman, and D. Parker. A game-based abstraction-refinement framework for markov decision processes. *Formal Methods in System Design*, 36:246–280, 2010. (Cited on page 2.)

[29] D. Kozen. Results on the propositional μ-calculus. *Theoretical computer science*, 27(3):333–354, 1983. (Cited on page 3.)

[30] J. Křetínskỳ, E. Ramneantu, A. Slivinskiy, and M. Weininger. Comparison of algorithms for simple stochastic games. *Information and Computation*, 289:104885, 2022. (Cited on page 2.)

[31] M. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: Verification of probabilistic real-time systems. In *Computer Aided Verification: 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings 23*, pages 585–591. Springer, 2011. (Cited on page 3.)

[32] M. Kwiatkowska, G. Norman, D. Parker, and G. Santos. Prism-games 3.0: Stochastic game verification with concurrency, equilibria and time. In *Computer Aided Verification: 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21–24, 2020, Proceedings, Part II 32*, pages 475–487. Springer, 2020. (Cited on page 3.)

[33] K. 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. (Cited on page 4.)

[34] D. S. Leslie, S. Perkins, and Z. Xu. Best-response dynamics in zero-sum stochastic games. *Journal of Economic Theory*, 189:105095, 2020. (Cited on page 2.)

[35] D. A. Martin. The determinacy of blackwell games. *The Journal of Symbolic Logic*, 63(4):1565–1581, 1998. (Cited on page 21.)

[36] A. McIver and C. Morgan. Results on the quantitative μ-calculus qm μ. *ACM Transactions on Computational Logic (TOCL)*, 8(1):3–es, 2007. (Cited on page 3.)

[37] C. Morgan and A. McIver. A probabilistic temporal calculus based on expectations. In *Proc. Formal Methods Pacific*, volume 97. Citeseer, 1997. (Cited on page 3.)

[38] A. W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In *Computation Theory: Fifth Symposium, Zaborów, Poland December 3–8, 1984 Proceedings*, pages 157–168. Springer, 1985. (Cited on page 3.)

[39] P. Parys. Parity games: Zielonka's algorithm in quasi-polynomial time. *arXiv preprint arXiv:1904.12446*, 2019. (Cited on page 4.)

[40] K. Phalakarn, T. Takisaka, T. Haas, and I. Hasuo. Widest paths and global propagation in bounded value iteration for stochastic games. In *Computer Aided Verification: 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21–24, 2020, Proceedings, Part II 32*, pages 349–371. Springer, 2020. (Cited on page 3.)

[41] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014. (Cited on page 2.)

[42] T. Quatmann and J.-P. Katoen. Sound value iteration. In *International Conference on Computer Aided Verification*, pages 643–661. Springer, 2018. (Cited on page 2.)

[43] L. S. Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953. (Cited on page 1.)

[44] C. Stirling. Local model checking games. In *International conference on concurrency theory*, pages 1–11. Springer, 1995. (Cited on page 3.)

[45] T. Wilke. Alternating tree automata, parity games, and modal μ-calculus. *Bulletin of the Belgian Mathematical Society-Simon Stevin*, 8(2):359–391, 2001. (Cited on page 3.)

[46] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996. (Cited on pages 2, 72, and 75.)