Weak Subgame Perfect Equilibria and their Application to Quantitative Reachability*

Thomas Brihaye[1], Véronique Bruyère[2], Noémie Meunier†[2], and Jean-François Raskin‡[3]

1 **Département de mathématique, Université de Mons (UMONS) Mons, Belgium**
2 **Département d'informatique, Université de Mons (UMONS) Mons, Belgium**
3 **Département d'informatique, Université Libre de Bruxelles (U.L.B.) Brussels, Belgium**

Abstract

We study n-player turn-based games played on a finite directed graph. For each play, the players have to pay a cost that they want to minimize. Instead of the well-known notion of Nash equilibrium (NE), we focus on the notion of subgame perfect equilibrium (SPE), a refinement of NE well-suited in the framework of games played on graphs. We also study natural variants of SPE, named weak (resp. very weak) SPE, where players who deviate cannot use the full class of strategies but only a subclass with a finite number of (resp. a unique) deviation step(s).

Our results are threefold. Firstly, we characterize in the form of a Folk theorem the set of all plays that are the outcome of a weak SPE. Secondly, for the class of quantitative reachability games, we prove the existence of a finite-memory SPE and provide an algorithm for computing it (only existence was known with no information regarding the memory). Moreover, we show that the existence of a constrained SPE, i.e. an SPE such that each player pays a cost less than a given constant, can be decided. The proofs rely on our Folk theorem for weak SPEs (which coincide with SPEs in the case of quantitative reachability games) and on the decidability of MSO logic on infinite words. Finally with similar techniques, we provide a second general class of games for which the existence of a (constrained) weak SPE is decidable.

1998 ACM Subject Classification B.6.3 [Design Aids] Automatic Synthesis, F.1.2 [Modes of Computation] Interactive and Reactive Computation

Keywords and phrases multi-player games on graphs, quantitative objectives, Nash equilibrium, subgame perfect equilibrium, quantitative reachability

Digital Object Identifier 10.4230/LIPIcs.CSL.2015.504

1 Introduction

Two-player zero-sum infinite duration games played on graphs are a mathematical model used to formalize several important problems in computer science. Reactive system synthesis is one such important problem. In this context, see e.g. [14], the vertices and the edges of the graph represent the states and the transitions of the system; one player models the system to synthesize, and the other player models the (uncontrollable) environment of the system.

* This work has been partly supported by European project Cassting (FP7-ICT-601148).
† Author supported by F.R.S.-FNRS fellowship.
‡ Author supported by ERC Starting Grant (279499: inVEST).

© Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin; licensed under Creative Commons License CC-BY
24th EACSL Annual Conference on Computer Science Logic (CSL 2015).
Editor: Stephan Kreutzer; pp. 504–518
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In the classical setting, the objectives of the two players are opposite, i.e. the environment is *adversarial*. Modeling the environment as fully adversarial is usually a bold abstraction of reality and there are recent works that consider the more general setting of non zero-sum games which allow to take into account the different objectives of each player. In this latter setting the environment has its own objective which is most often *not* the negation of the objective of the system. The concept of *Nash equilibrium* (NE) [12] is central to the study of non zero-sum games and can be applied to the general setting of n player games. A strategy profile is a NE if no player has an incentive to deviate unilaterally from his strategy, since he cannot strictly improve on the outcome of the strategy profile by changing his strategy only.

However in the context of sequential games (such as games played on graphs), it is well-known that NEs present a serious weakness: a NE allows for *non-credible threats* that rational players should not carry out [16]. Hence, for sequential games, the notion of NE has been strengthened into the notion of *subgame perfect equilibrium* (SPE): a strategy profile is an SPE if it is a NE in all the subgames of the original game. While the notion of SPE is rather well understood for finite state game graphs with ω-regular objectives or for games in finite extensive form (finite game trees), less is known for game graphs with *quantitative objectives* in which players encounter costs that they want to minimize, like in classical quantitative objectives such as mean-payoff, discounted sum, or quantitative reachability.

Several natural and important questions arise for such games: Can we decide the existence of an SPE, and more generally the *constrained* existence of an SPE (i.e. an SPE in which each player encounters a cost less than some fixed value)? Can we compute such SPEs that use finite-memory strategies only? Whereas several authors have studied what the hypotheses are to impose on games in a way to guarantee the existence of an SPE, the previous algorithmic questions are still wide open. In this article, we provide progress in the understanding of the notion of SPE. We study some variants of SPEs and establish a theorem that characterizes their possible outcomes in quantitative games. We derive from this characterization interesting algorithms and information on the strategies for two important classes of quantitative games. Our contributions are detailed in the next paragraph.

Contributions. First, we formalize a notion of *deviation step* from a strategy profile that allows us to define two natural variants of NEs. While a NE must be resistant to the unilateral deviation of one player for any number of deviation steps, a *weak* (resp. *very weak*) NE must be resistant to the unilateral deviation of one player for any *finite* number of (resp. a unique) deviation step(s). Then we use those variants to define the corresponding notions of *weak* and *very weak* SPE. The latter notion is very close to the one-step deviation property [13]. Any very weak SPE is also a weak SPE, and there are games for which there exists a weak SPE but no SPE. Also, for games with upper-semicontinuous cost functions and for games played on finite game trees, the three notions are equivalent.

Second, we characterize in the form of a Folk theorem[1] all the possible outcomes of weak SPEs. The characterization is obtained starting from all possible plays of the game and the application of a nonincreasing operator that removes plays that cannot be outcome of a weak SPE. We show that the limit of the nonincreasing chain of sets always exists and contains exactly all the possible outcomes of weak SPEs. Furthermore, we show how for each such outcome, we can associate a strategy profile that generates it and which is a weak SPE.

[1] We do not consider our result as folklore, but we use this terminology, as also done in [6], in reference to the "classical folk theorems" for repeated games which characterize the payoff profiles of NEs and SPEs in repeated games (see for instance Chapter 8 in [13]).

Additionally, to illustrate the potential of our Folk theorem, we show how it can be refined and used to answer open questions about two classes of quantitative games. The first class of games that we consider are *quantitative reachability games*, such that each player aims at reaching his own set of target states as soon as possible. As the cost functions in those games are continuous, our Folk theorem characterizes precisely the outcomes of SPEs and not only weak SPEs. In [1, 7], it has been shown that quantitative reachability games always have SPEs. The proof provided for this theorem is non constructive since it relies on topological arguments. Here, we strengthen this existential result by proving that there always exists, not only an SPE but, a *finite-memory* SPE. Furthermore, we provide an algorithm to construct such a finite memory SPE. This algorithm is based on a constructive version of our Folk Theorem for the class of quantitative reachability games: we show that the nonincreasing chain of sets of potential outcomes stabilizes after a finite number of steps and that each intermediate set is an ω-regular set that can be effectively described using MSO sentences. The second class of games that we consider is the class of games with cost functions that are *prefix-independent*, whose range of values is *finite*, and for which each value has an *ω-regular* pre-image. For this general class of games, with similar techniques as for quantitative reachability games, we show how to construct an effective representation of all possible outcomes compatible with a weak SPE, and consequently that the existence of a weak SPE is decidable. In those two applications, we show that our construction also allow us to answer the question of existence of a constrained (weak) SPE, i.e. a (weak) SPE in which players pays a cost which is bounded by a given value.

Related work. The concept of SPE has been first introduced and studied by the game theory community. The notion of SPE has been first introduced by Kuhn in finite extensive form games [10]. For such games, backward induction can be used to prove that there always exists an SPE. By inspecting the backward induction proof, it is not difficult to realize that the notion of very weak SPE and SPE are equivalent in this context.

SPEs for infinite trees defined as the unfolding of finite graphs with *qualitative*, i.e. win-lose, ω-regular objectives, have been studied by Ummels in [19]: it is proved that such games always have an SPE, and that the existence of a constrained SPE is decidable.

In [9], the authors provide an effective representation of the outcomes of NEs in concurrent priced games by constructing a Büchi automaton accepting the language of outcomes of all NEs satisfying a bound vector. The existence of NEs in quantitative games played on graphs is studied in [3]; it is shown that for a large class of games, there always exists a finite-memory NE. This result is extended in [4] for two-player games and secure equilibria (a refinement of NEs); additionally the constrained existence problem for secure equilibria is also shown decidable for a large range of cost functions. None of these articles consider SPEs.

In [6], the authors prove that for quantitative games with cost functions that are upper-semicontinous and with finite range, there always exists an SPE. This result also relies on a nonincreasing chain of sets of possible outcomes of SPEs. The main differences with our work is that we obtain a Folk theorem that *characterizes* all possible outcomes of weak SPEs with no restriction on the cost functions. Moreover we have shown that our Folk theorem can be made effective for two classes of quantitative games of interest. Effectiveness issues are not considered in [6]. Prior to this work, Mertens shows in [11] that if the cost functions are bounded and Borel measurable then there always exists an ϵ-NE. In [7], Fudenberg et al. show that if the cost functions are all continuous, then there always exists an SPE. Those results were recently extended in [15] by Le Roux and Pauly.

Organization of the article. In Section 2, we present the notions of quantitative game, classical NE and SPE, and their variants. In Section 3, we propose our Folk Theorem for weak SPEs. In Section 4, we provide an algorithm for computing a finite-memory SPE for quantitative reachability games, and a second algorithm to decide the constrained existence of an SPE for this class of games. We also show that the existence of a (constrained) weak SPE is decidable for another class of games. A conclusion and future work are given in the last section.

2 Preliminaries and Variants of Equilibria

In this section, we recall the notions of quantitative game, Nash equilibrium, and subgame perfect equilibrium. We also introduce variants of Nash and subgame perfect equilibria, and compare them with the classical notions.

2.1 Quantitative Games

We consider multi-player turn-based non zero-sum quantitative games in which, for each infinite play, players pay a cost that they want to minimize.[2]

▶ **Definition 1.** A *quantitative game* is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda})$ where:

- Π is a finite set of players,
- V is a finite set of vertices,
- $(V_i)_{i\in\Pi}$ is a partition of V such that V_i is the set of vertices controlled by player $i \in \Pi$,
- $E \subseteq V \times V$ is a set of edges, such that[3] for all $v \in V$, there exists $v' \in V$ with $(v, v') \in E$,
- $\bar{\lambda} = (\lambda_i)_{i\in\Pi}$ is a cost function such that $\lambda_i : V^\omega \to \mathbb{R} \cup \{+\infty\}$ is player i cost function.

A *play* of $\mathcal{G}$ is an infinite sequence $\rho = \rho_0\rho_1 \ldots \in V^\omega$ such that $(\rho_i, \rho_{i+1}) \in E$ for all $i \in \mathbb{N}$. *Histories* of $\mathcal{G}$ are finite sequences $h = h_0 \ldots h_n \in V^+$ defined in the same way. The *length* $|h|$ of h is the number n of its edges. We denote by $\text{First}(h)$ (resp. $\text{Last}(h)$) the first vertex h_0 (resp. last vertex h_n) of h. Usually histories are non-empty, but in specific situations it will be useful to consider the empty history ϵ. The set of all histories (ended by a vertex in V_i) is denoted by Hist (by Hist_i). A *prefix* (resp. *suffix*) of a play ρ is a finite sequence $\rho_0 \ldots \rho_n$ (resp. infinite sequence $\rho_n\rho_{n+1} \ldots$) denoted by $\rho_{\leq n}$ or $\rho_{<n+1}$ (resp. $\rho_{\geq n}$). We use notation $h < \rho$ when a history h is prefix of a play ρ. Given two distinct plays ρ and ρ', their longest common prefix is denoted by $\rho\hat{\ }\rho'$.

When an initial vertex $v_0 \in V$ is fixed, we call $(\mathcal{G}, v_0)$ an *initialized* quantitative game. A play (resp. a history) of $(\mathcal{G}, v_0)$ is a play (resp. history) of $\mathcal{G}$ starting in v_0. The set of histories $h \in \text{Hist}$ (resp. $h \in \text{Hist}_i$) with $\text{First}(h) = v_0$ is denoted by $\text{Hist}(v_0)$ (resp. $\text{Hist}_i(v_0)$). In the figures of this article, we will often *unravel* the graph of the game $(\mathcal{G}, v_0)$ from the initial vertex v_0, which ends up in an infinite tree.

Given a play $\rho \in V^\omega$, its *cost* is given by $\bar{\lambda}(\rho) = (\lambda_i(\rho))_{i\in\Pi}$. In this article, we are particularly interested in quantitative reachability games in which $\lambda_i(\rho)$ is equal to the number of edges to reach a given set of vertices.

▶ **Definition 2.** A *quantitative reachability game* is a quantitative game $\mathcal{G}$ such that the cost function $\bar{\lambda} : V^\omega \to (\mathbb{N} \cup \{+\infty\})^\Pi$ is defined as follows. Each player i has a *target set* $T_i \subseteq V$, and for each play $\rho = \rho_0\rho_1 \ldots$ of $\mathcal{G}$, the cost $\lambda_i(\rho)$ is the least index n such that $\rho_n \in T_i$ if it exists, and $+\infty$ otherwise.

[2] Alternatively, players could receive a payoff that they want to maximize.
[3] Each vertex has at least one outgoing edge.

Notice that the cost function $\bar{\lambda}$ of a quantitative game is often defined from $|\Pi|$-tuples of weights labeling the edges of the game. For instance, in inf games, $\lambda_i(\rho)$ is equal to the infimum of player i weights seen along ρ. Some other classical examples are liminf, limsup, mean-payoff, and discounted sum games [5]. In case of quantitative reachability on graphs with weighted edges, the cost $\lambda_i(\rho)$ for player i is replaced by the sum of the weights seen along ρ until his target set is reached. We do not consider this extension here. Notice that when weights are positive integers, replacing each edge with cost c by a path of length c composed of c new edges allows to recover Definition 2.

Let us recall the notions of prefix-independent, continuous, and lower- (resp. upper-) semicontinuous cost functions. Since V is endowed with the discrete topology, and thus V^ω with the product topology, a sequence of plays $(\rho_n)_{n\in\mathbb{N}}$ converges to a play $\rho = \lim_{n\to\infty}\rho_n$ if every prefix of ρ is prefix of all ρ_n except, possibly, of finitely many of them.

▶ **Definition 3.** Let λ_i be a player i cost function. Then

- λ_i is *prefix-independent* if $\lambda_i(h\rho) = \lambda_i(\rho)$ for any history h and play ρ.
- λ_i is *continuous* if whenever $\lim_{n\to\infty}\rho_n = \rho$, then $\lim_{n\to\infty}\lambda_i(\rho_n) = \lambda_i(\rho)$.
- λ_i *upper-semicontinuous* (resp. *lower-semicontinuous*) if whenever $\lim_{n\to\infty}\rho_n = \rho$, then $\limsup_{n\to\infty}\lambda_i(\rho_n) \le \lambda_i(\rho)$ (resp. $\liminf_{n\to\infty}\lambda_i(\rho_n) \ge \lambda_i(\rho)$).

For instance, the cost functions used in liminf and mean-payoff games are prefix-independent, contrarily to the case of inf games. Clearly, if λ_i is continuous, then it is upper- and lower-semicontinuous. The cost functions of liminf and mean-payoff games are neither upper-semicontinuous nor lower-semicontinuous, whereas they are continuous in discounted sum games. The cost functions λ_i used in quantitative reachability games can be transformed into continuous ones as follows [1]: $\lambda'_i(\rho) = 1 - \frac{1}{\lambda_i(\rho)+1}$ if $\lambda_i(\rho) < +\infty$, and $\lambda'_i(\rho) = 1$ otherwise.

2.2 Strategies and Deviations

A *strategy* σ for player $i \in \Pi$ is a function $\sigma : \mathrm{Hist}_i \to V$ assigning to each history[4] $hv \in \mathrm{Hist}_i$ a vertex $v' = \sigma(hv)$ such that $(v, v') \in E$. In an initialized game $(\mathcal{G}, v_0)$, σ is restricted to histories starting with v_0. A player i strategy σ is *positional* if it only depends on the last vertex of the history, i.e. $\sigma(hv) = \sigma(v)$ for all $hv \in \mathrm{Hist}_i$. It is a *finite-memory* strategy if it needs only finite memory of the history (recorded by a finite strategy automaton, also called a Moore machine). A play ρ is *consistent* with a player i strategy σ if $\rho_{k+1} = \sigma(\rho_{\le k})$ for all k such that $\rho_k \in V_i$. A *strategy profile* of $\mathcal{G}$ is a tuple $\bar{\sigma} = (\sigma_i)_{i\in\Pi}$ of strategies, where each σ_i is a player i strategy. It is called *positional* (resp. *finite-memory*) if all σ_i, $i \in \Pi$, are positional (resp. finite-memory). Given an initial vertex v_0, such a strategy profile determines a unique play of $(\mathcal{G}, v_0)$ that is consistent with all the strategies. This play is called the *outcome* of $\bar{\sigma}$ and is denoted by $\langle\bar{\sigma}\rangle_{v_0}$.

Given σ_i a player i strategy, we say that player i *deviates* from σ_i if he does not stick to σ_i and prefers to use another strategy σ'_i. Let $\bar{\sigma}$ be a strategy profile. When all players stick to their strategy σ_i except player i that shifts to σ'_i, we denote by (σ'_i, σ_{-i}) the derived strategy profile, and by $\langle\sigma'_i, \sigma_{-i}\rangle_{v_0}$ its outcome in $(\mathcal{G}, v_0)$. In the next definition, we introduce the notion of deviation step of a strategy σ'_i from a given strategy profile $\bar{\sigma}$.

[4] In this article we often write a history in the form hv with $v \in V$ to emphasize that v is the last vertex of this history.

▶ **Definition 4.** Let $(\mathcal{G}, v_0)$ be an initialized game, $\bar{\sigma}$ be a strategy profile, and σ'_i be a player i strategy. We say that σ'_i has a *hv-deviation step* from $\bar{\sigma}$ for some history $hv \in \text{Hist}_i(v_0)$ with $v \in V_i$, if $hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ and $\sigma_i(hv) \neq \sigma'_i(hv)$.

Notice that the previous definition requires that hv is a prefix of the outcome $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$; it says nothing about σ'_i outside of this outcome. A strategy σ'_i can have a finite or an infinite number of deviation steps in the sense of Definition 4. A strategy with three deviation steps is depicted in Figure 1 such that each $h_k v_k$-deviation step from $\bar{\sigma}$, $1 \leq k \leq 3$, is highlighted with a dashed edge.

In light of Definition 4, we introduce the following classes of strategies.

▶ **Definition 5.** Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar{\sigma}$ be a strategy profile.

- A strategy σ'_i is *finitely deviating* from $\bar{\sigma}$ if it has a finite number of deviation steps from $\bar{\sigma}$.
- It is *one-shot deviating* from $\bar{\sigma}$ if it has a v_0-deviation step from $\bar{\sigma}$, and no other deviation step.

In other words, a strategy σ'_i is finitely deviating from $\bar{\sigma}$ if there exists a history $hv < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$ such that for all $h'v'$, $hv \leq h'v' < \langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$, we have $\sigma'_i(h'v') = \sigma_i(h'v')$ (σ'_i acts as σ_i from hv along $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$). The strategy σ'_i is one-shot deviating from $\bar{\sigma}$ if it differs from σ_i at the initial vertex v_0, and after v_0 acts as σ_i along $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$. As for Definition 4, the previous definition says nothing about σ'_i outside of $\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}$. Clearly any one-shot deviating strategy is finitely deviating. The strategy of Figure 1 is finitely deviating but not one-shot deviating.

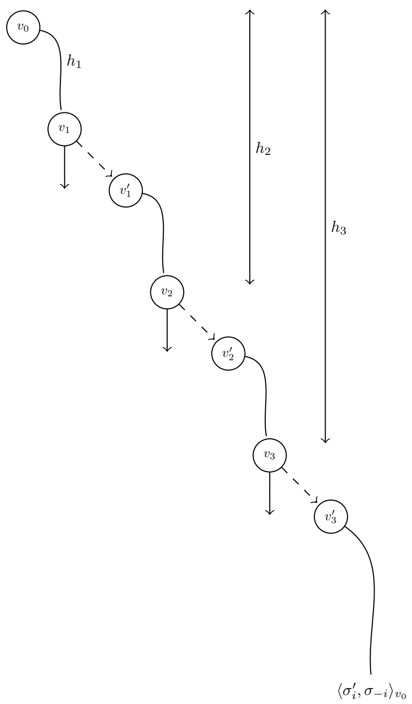

Figure 1 A strategy σ'_i with a finite number of deviation steps.

2.3 Nash and Subgame Perfect Equilibria, and Variants

In this paper, we focus on subgame perfect equilibria and their variants. Let us first recall the classical notion of Nash equilibrium. A strategy profile $\bar{\sigma}$ in an initialized game is a Nash equilibrium if no player has an incentive to deviate unilaterally from his strategy, since he cannot strictly decrease his cost when using any other strategy.

▶ **Definition 6.** Given an initialized game $(\mathcal{G}, v_0)$, a strategy profile $\bar{\sigma} = (\sigma_i)_{i \in \Pi}$ of $(\mathcal{G}, v_0)$ is a *Nash equilibrium (NE)* if for all players $i \in \Pi$, for all player i strategies σ'_i, we have $\lambda_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}) \geq \lambda_i(\langle \bar{\sigma} \rangle_{v_0})$.

We say that a player i strategy σ'_i is a *profitable deviation* for i w.r.t. $\bar{\sigma}$ if $\lambda_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}) < \lambda_i(\langle \bar{\sigma} \rangle_{v_0})$. Therefore $\bar{\sigma}$ is a NE if no player has a profitable deviation w.r.t. $\bar{\sigma}$.

Let us propose the next variants of NE.

▶ **Definition 7.** Let $(\mathcal{G}, v_0)$ be an initialized game. A strategy profile $\bar{\sigma}$ is a *weak NE* (resp. *very weak NE*) in $(\mathcal{G}, v_0)$ if, for each player $i \in \Pi$, for each finitely deviating (resp. one-shot deviating) strategy σ'_i of player i, we have $\lambda_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0}) \geq \lambda_i(\langle \bar{\sigma} \rangle_{v_0})$.

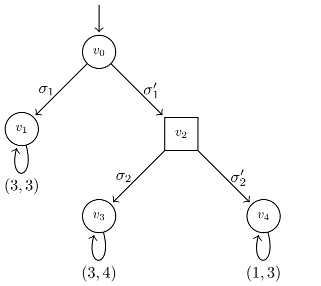

Figure 2 A simple two-player quantitative game.

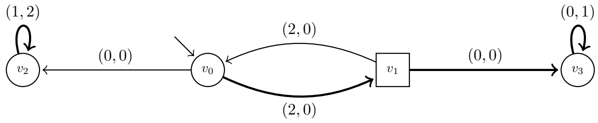

Figure 3 A two-player game with a (very) weak SPE and no SPE. For each player, the cost of a play is his unique weight seen in the ending cycle.

▶ **Example 8.** Consider the two-player quantitative game depicted in Figure 2. Circle (resp. square) vertices are player 1 (resp. player 2) vertices. The edges are labeled by couples of weights such that weights $(0, 0)$ are not specified. For each player i, the cost $\lambda_i(\rho)$ of a play ρ is the weight of its ending loop. In this simple game, each player i have two positional strategies that are respectively denoted by σ_i and σ'_i (see Figure 2).

The strategy profile (σ_1, σ'_2) is not a NE since σ'_1 is a profitable deviation for player 1 w.r.t. (σ_1, σ'_2) (player 1 pays cost 1 instead of cost 3). This strategy profile is neither a weak NE nor a very weak NE because in this simple game, player 1 can only deviate from σ_1 by using the one-shot deviating strategy σ'_1. On the contrary, the strategy profile (σ_1, σ_2) is a NE with outcome $v_0 v_1^\omega$ of cost $(3, 3)$. It is also a weak NE and a very weak NE.

By definition, any NE is a weak NE, and any weak NE is a very weak NE. The contrary is false: in the previous example, (σ'_1, σ_2) is a very weak NE, but not a weak NE. We will see later an example of game with a weak NE that is not an NE (see Example 12).

The notion of subgame perfect equilibrium is a refinement of NE. In order to define it, we need to introduce the following notions. Given a quantitative game $\mathcal{G} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda})$ and a history h of $\mathcal{G}$, we denote by $\mathcal{G}_{\restriction h}$ the game $\mathcal{G}_{\restriction_h} = (\Pi, V, (V_i)_{i\in\Pi}, E, \bar{\lambda}_{\restriction_h})$ where $\bar{\lambda}_{\restriction h}(\rho) = \bar{\lambda}(h\rho)$ for any play of $\mathcal{G}_{\restriction h}$[5], and we say that $\mathcal{G}_{\restriction h}$ is a *subgame* of $\mathcal{G}$. Given an initialized game $(\mathcal{G}, v_0)$, and a history $hv \in \text{Hist}(v_0)$, the initialized game $(\mathcal{G}_{\restriction h}, v)$ is called the subgame of $(\mathcal{G}, v_0)$ with history hv. Notice that $(\mathcal{G}, v_0)$ can be seen as a subgame of itself with history hv_0 such that $h = \epsilon$. Given a player i strategy σ in $(\mathcal{G}, v_0)$, we define the strategy $\sigma_{\restriction h}$ in $(\mathcal{G}_{\restriction h}, v)$ as $\sigma_{\restriction h}(h') = \sigma(hh')$ for all histories $h' \in \text{Hist}_i(v)$. Given a strategy profile $\bar{\sigma} = (\sigma_i)_{i\in\Pi}$, we use notation $\bar{\sigma}_{\restriction h}$ for $(\sigma_{i\restriction h})_{i\in\Pi}$, and $\langle\bar{\sigma}_{\restriction h}\rangle_v$ is its outcome in the subgame $(\mathcal{G}_{\restriction h}, v)$.

We can now recall the classical notion of subgame perfect equilibrium: it is a strategy profile in an initialized game that induces a NE in each of its subgames. In particular, a subgame perfect equilibrium is a NE.

▶ **Definition 9.** Given an initialized game $(\mathcal{G}, v_0)$, a strategy profile $\bar{\sigma}$ of $(\mathcal{G}, v_0)$ is a *subgame perfect equilibrium (SPE)* if $\bar{\sigma}_{\restriction h}$ is a NE in $(\mathcal{G}_{\restriction h}, v)$, for every history $hv \in \text{Hist}(v_0)$.

As for NE, we propose the next variants of SPE.

▶ **Definition 10.** Let $(\mathcal{G}, v_0)$ be an initialized game. A strategy profile $\bar{\sigma}$ is a *weak SPE* (resp. *very weak SPE*) if $\bar{\sigma}_{\restriction h}$ is a weak NE (resp. very weak NE) in $(\mathcal{G}_{\restriction h}, v)$, for all histories $hv \in \text{Hist}(v_0)$.

---

[5] In this article, we will always use notation $\bar{\lambda}(h\rho)$ instead of $\bar{\lambda}_{\restriction h}(\rho)$.

▶ **Example 11.** We come back to the game depicted in Figure 2. We have seen before that the strategy profile (σ_1, σ_2) is a NE. Notice that this NE uses a non-credible threat of player 2 that prefers to pay a cost of 4 instead of 3 (by using σ_2'). Such a threat is not allowed for SPEs. Indeed consider the subgame $(\mathcal{G}_{\restriction v_0}, v_2)$ of $(\mathcal{G}, v_0)$ with history v_0v_2. In this subgame, σ_2' is a profitable deviation for player 2, showing that (σ_1, σ_2) is not an SPE. One can easily verify that the strategy profile (σ_1', σ_2') is an SPE, as well as a weak SPE and a very weak SPE, due to the simple form of the game.

The previous example is too simple to show the differences between classical SPEs and their variants. The next example presents a game with a (very) weak SPE but no SPE.

▶ **Example 12.** Consider the two-player game $(\mathcal{G}, v_0)$ in Figure 3. The edges are labeled by couples of weights, and for each player i the cost $\lambda_i(\rho)$ of a play ρ is the unique weight seen in its ending cycle. With this definition, $\lambda_i(\rho)$ can also be seen as either the mean-payoff, or the liminf, or the limsup, of the weights of ρ. It is known that this game has no SPE [17].

Let us show that the strategy profile $\bar{\sigma}$ depicted with thick edges is a very weak SPE. Due to the simple form of the game, only two cases are to be treated. Consider the subgame $(\mathcal{G}_{\restriction h}, v_0)$ with $h \in (v_0v_1)^*$, and the one-shot deviating strategy σ_1' of player 1 such that $\sigma_1'(v_0) = v_2$. Then $\langle \bar{\sigma}_{\restriction h} \rangle_{v_0} = v_0v_1v_3^\omega$ and $\langle \sigma_1', \sigma_{2\restriction h} \rangle_{v_0} = v_0v_2^\omega$, showing that σ_1' is not a profitable deviation for player 1. One also checks that in the subgame $(\mathcal{G}_{\restriction h}, v_1)$ with $h \in (v_0v_1)^*v_0$, the one-shot deviating strategy σ_2' of player 2 such that $\sigma_2'(v_1) = v_0$ is not profitable for him.

Similarly, one can prove that $\bar{\sigma}$ is a weak SPE (see also Proposition 13 hereafter). Notice that $\bar{\sigma}$ is not an SPE. Indeed the strategy σ_2' such that $\sigma_2'(hv_1) = v_0$ for all h, is a profitable deviation for player 2 in $(\mathcal{G}, v_0)$. This strategy is (of course) not finitely deviating. Finally notice that $\bar{\sigma}$ is a weak NE that is not an NE.

From Definition 10, any SPE is a weak SPE, and any weak SPE is a very weak SPE. The next proposition states that weak SPE and very weak SPE are equivalent notions, but this is no longer true for SPE and weak SPE as shown previously by Example 12. The first part of the proof is based on arguments from the one-step deviation property used to prove Kuhn's theorem [10]. The second part follows from Example 12 [17].

▶ **Proposition 13.**
- *Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar{\sigma}$ be a strategy profile. Then $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE.*
- *There exists an initialized game $(\mathcal{G}, v_0)$ with a weak SPE but no SPE.*

Under the next hypotheses on the game or the costs, the equivalence between SPE, weak SPE, and very weak SPE holds. The first case, when the cost functions are continuous, is a classical result in game theory, see for instance [8]; the second case appears as a part of the proof of Kuhn's theorem [10].

▶ **Proposition 14.** *Let $(\mathcal{G}, v_0)$ be an initialized game, and $\bar{\sigma}$ be a strategy profile.*
- *If all cost functions λ_i are continuous, or even upper-semicontinuous*[6]*, then $\bar{\sigma}$ is an SPE iff $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE.*
- *If $\mathcal{G}$ is a finite tree*[7]*, then $\bar{\sigma}$ is an SPE iff $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE.*

---

[6] In games where the players receive a payoff that they want to maximize, the hypothesis of upper-semicontinuity has to be replaced by lower-semicontinuity.
[7] In a finite tree game, the plays are finite sequences of vertices ending in a leaf and their cost is associated with the ending leaf. An example of such a game is depicted in Figure 2.

Recall that discounted sum games and quantitative reachability games are continuous. Thus for these games, the three notions of SPE, weak SPE and very weak SPE, are equivalent.

▶ **Corollary 15.** *Let $(\mathcal{G}, v_0)$ be an initialized quantitative reachability game, and $\bar{\sigma}$ be a strategy profile. Then $\bar{\sigma}$ is an SPE iff $\bar{\sigma}$ is a weak SPE iff $\bar{\sigma}$ is a very weak SPE.*

On the opposite, the initialized game of Figure 3 has a weak SPE but no SPE. Its cost function λ_2 is not upper-semicontinuous. Indeed, we have that $\lim_{n\to\infty}(v_0v_1)^n v_3^\omega = (v_0v_1)^\omega$ and $\lim_{n\to\infty} \lambda_2((v_0v_1)^n v_3^\omega) = 1 > 0 = \lambda_2((v_0v_1)^\omega)$.

3 Folk Theorem for Weak SPEs

In this section, we characterize in the form of a Folk Theorem the set of all outcomes of weak SPEs. To this end we define a nonincreasing sequence of sets of plays that initially contain all the plays, and then lose, step by step, some plays that for sure are not outcomes of a weak SPE, until finally reaching a fixpoint.

Let $(\mathcal{G}, v_0)$ be a game. For an ordinal α and a history $hv \in \text{Hist}(v_0)$, let us consider the set $\mathbf{P}_\alpha(hv) = \{\rho \mid \rho \text{ is a } \textit{potential} \text{ outcome of a weak NE in } (\mathcal{G}_{\restriction h}, v) \text{ at step } \alpha\}$. This set is defined by induction on α as follows:

▶ **Definition 16.** Let $(\mathcal{G}, v_0)$ be a quantitative game. The set $\mathbf{P}_\alpha(hv)$ is defined as follows for each ordinal α and history $hv \in \text{Hist}(v_0)$:

- For $\alpha = 0$,
  $$\mathbf{P}_\alpha(hv) = \{\rho \mid \rho \text{ is a play in } (\mathcal{G}_{\restriction h}, v)\}. \quad (1)$$
- For a successor ordinal $\alpha + 1$,
  $$\mathbf{P}_{\alpha+1}(hv) = \mathbf{P}_\alpha(hv) \setminus \mathbf{E}_\alpha(hv) \quad (2)$$
  such that $\rho \in \mathbf{E}_\alpha(hv)$ (see Figure 4) iff
  - there exists a history h', $hv \le h' < h\rho$, and $\text{Last}(h') \in V_i$ for some i,
  - there exists a vertex v', $h'v' \not< h\rho$,
  - such that $\forall \rho' \in \mathbf{P}_\alpha(h'v')$: $\lambda_i(h\rho) > \lambda_i(h'\rho')$.
- For a limit ordinal α:
  $$\mathbf{P}_\alpha(hv) = \bigcap_{\beta<\alpha} \mathbf{P}_\beta(hv). \quad (3)$$

Notice that an element ρ of $\mathbf{P}_\alpha(hv)$ is a play in $(\mathcal{G}_{\restriction h}, v)$ (and not in $(\mathcal{G}, v_0)$). Therefore it starts with vertex v, and $h\rho$ is a play in $(\mathcal{G}, v_0)$. For $\alpha + 1$ being a successor ordinal, play $\rho \in \mathbf{E}_\alpha(hv)$ is erased from $\mathbf{P}_\alpha(hv)$ when for all $\rho' \in \mathbf{P}_\alpha(h'v')$, player i pays a lower cost $\lambda_i(h'\rho') < \lambda_i(h\rho)$, meaning that ρ is no longer a potentiel outcome of a weak NE in $(\mathcal{G}_{\restriction h}, v)$.

The sequence $(\mathbf{P}_\alpha(hv))_\alpha$ is nonincreasing by definition, and reaches a fixpoint:

▶ **Proposition 17.** *There exists an ordinal α_* such that $\mathbf{P}_{\alpha_*}(hv) = \mathbf{P}_{\alpha_*+1}(hv)$ for all histories $hv \in \text{Hist}(v_0)$.*

In the sequel, α_* always refers to the ordinal mentioned in Proposition 17.

Our Folk Theorem for weak SPEs is the next one.

▶ **Theorem 18.** *Let $(\mathcal{G}, v_0)$ be a quantitative game. There exists a weak SPE in $(\mathcal{G}, v_0)$ with outcome ρ iff $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \text{Hist}(v_0)$, and $\rho \in \mathbf{P}_{\alpha_*}(v_0)$.*

The proof of Theorem 18 follows from Lemmas 19 and 20.

▶ **Lemma 19.** *If $(\mathcal{G}, v_0)$ has a weak SPE $\bar{\sigma}$, then $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \text{Hist}(v_0)$, and $\langle\bar{\sigma}\rangle_{v_0} \in \mathbf{P}_{\alpha_*}(v_0)$.*

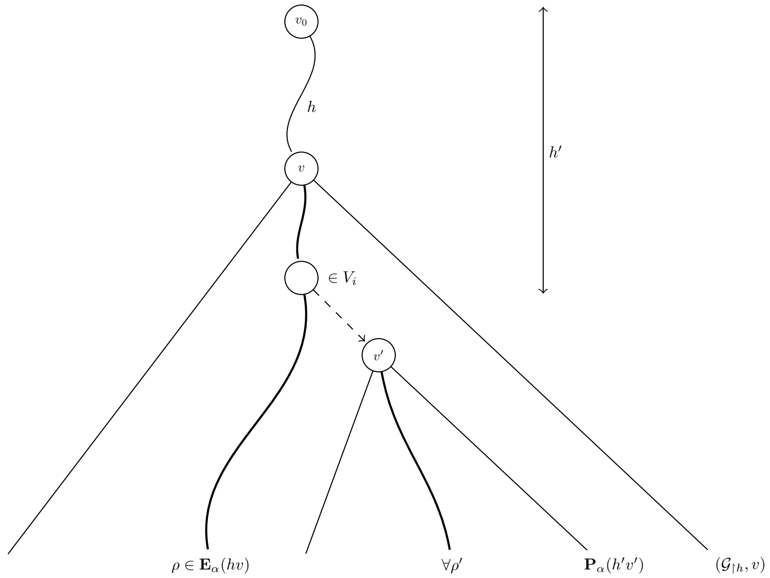

Figure 4 $\rho \in \mathbf{E}_\alpha(hv)$.

Proof. Let us show, by induction on α, that $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{P}_\alpha(hv)$ for all $hv \in \mathrm{Hist}(v_0)$.

For $\alpha = 0$, we have $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{P}_\alpha(hv)$ by definition of $\mathbf{P}_0(hv)$.

Let $\alpha + 1$ be a successor ordinal. By induction hypothesis, we have that $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{P}_\alpha(hv)$ for all $hv \in \mathrm{Hist}(v_0)$. Suppose that there exists hv such that $\langle \bar{\sigma}_{\restriction h} \rangle_v \notin \mathbf{P}_{\alpha+1}(hv)$, i.e. $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{E}_\alpha(hv)$. This means that there is a history $h' = hg \in \mathrm{Hist}_i$ for some $i \in \Pi$ with $hv \leq h' < h\rho$, and there exists a vertex v' with $h'v' \not< h\rho$, such that $\forall \rho' \in \mathbf{P}_\alpha(h'v')$, $\lambda_i(h \cdot \langle \bar{\sigma}_{\restriction h} \rangle_v) > \lambda_i(h'\rho')$. In particular, by induction hypothesis

$$\lambda_i(h \cdot \langle \bar{\sigma}_{\restriction h} \rangle_v) > \lambda_i(h' \cdot \langle \bar{\sigma}_{\restriction h'} \rangle_{v'}). \tag{4}$$

Let us consider the player i strategy σ'_i in $(\mathcal{G}_{\restriction h}, v)$ such that $g \cdot \langle \bar{\sigma}_{\restriction h'} \rangle_{v'}$ is consistent with σ'_i. Then σ'_i is a finitely deviating strategy with the (unique) g-deviation step from $\bar{\sigma}_{\restriction h}$. It is a profitable deviation for player i in $(\mathcal{G}_{\restriction h}, v)$ by (4), a contradiction with $\bar{\sigma}$ being a weak SPE.

Let α be a limit ordinal. By induction hypothesis $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{P}_\beta(hv), \forall \beta < \alpha$. Therefore $\langle \bar{\sigma}_{\restriction h} \rangle_v \in \mathbf{P}_\alpha(hv) = \bigcap_{\beta < \alpha} \mathbf{P}_\beta(hv)$. ◀

▶ **Lemma 20.** *Suppose that $\mathbf{P}_{\alpha_*}(hv) \neq \emptyset$ for all $hv \in \mathrm{Hist}(v_0)$, and let $\rho \in \mathbf{P}_{\alpha_*}(v_0)$. Then $(\mathcal{G}, v_0)$ has a weak SPE with outcome ρ.*

Proof. We are going to show how to construct a very weak SPE $\bar{\sigma}$ (and thus a weak SPE by Proposition 13) with outcome ρ. The construction of $\bar{\sigma}$ is done step by step thanks to a progressive labeling of the histories $hv \in \mathrm{Hist}(v_0)$. Let us give an intuitive idea of the construction of $\bar{\sigma}$. Initially, we partially construct $\bar{\sigma}$ such that it produces an outcome in $(\mathcal{G}, v_0)$ equal to $\rho \in \mathbf{P}_{\alpha_*}(v_0)$; we also label each non-empty prefix of ρ by ρ. Then we consider a shortest non-labeled history $h'v'$, and we correctly choose some $\rho' \in \mathbf{P}_{\alpha_*}(h'v')$ (we will see later how). We continue the construction of $\bar{\sigma}$ such that it produces the outcome ρ' in $(\mathcal{G}_{\restriction h'}, v')$, and for each non-empty prefix g of ρ', we label $h'g$ by ρ' (notice that the prefixes of h' have already been labeled by choice of h'). And so on. In this way, the labeling is a map $\gamma : \text{Hist}(v_0) \to \bigcup_{hv} \mathbf{P}_{\alpha_*}(hv)$ that allows to recover from $h'g$ the outcome ρ' of $\bar{\sigma}_{\restriction h'}$ in $(\mathcal{G}_{\restriction h'}, v')$ of which g is prefix. Let us now go into the details.

Initially, none of the histories is labeled. We start with history v_0 and the given play $\rho \in \mathbf{P}_{\alpha_*}(v_0)$. The strategy profile $\bar{\sigma}$ is partially defined such that $\langle\bar{\sigma}\rangle_{v_0} = \rho$, that is, if $\rho = \rho_0\rho_1 \ldots$, then $\sigma_i(\rho_{\leq n}) = \rho_{n+1}$ for all $\rho_n \in V_i$ and $i \in \Pi$. The non-empty prefixes h of ρ are all labeled with $\gamma(h) = \rho$.

At the following steps, we consider a history $h'v'$ that is not yet labeled, but such that h' has already been labeled. By induction, $\gamma(h') = \langle\bar{\sigma}_{\restriction h}\rangle_v$ such that $hv \leq h'$. Suppose that $\text{Last}(h') \in V_i$, we then choose a play $\rho' \in \mathbf{P}_{\alpha_*}(h'v')$ such that (see Figure 5)

$$\lambda_i(h \cdot \langle\bar{\sigma}_{\restriction h}\rangle_v) \leq \lambda_i(h'\rho'). \tag{5}$$

Such a play ρ' exists for the next reasons. By induction, we know that $\langle\bar{\sigma}_{\restriction h}\rangle_v \in \mathbf{P}_{\alpha_*}(hv)$. Since $\mathbf{P}_{\alpha_*}(hv) = \mathbf{P}_{\alpha_*+1}(hv)$ by Proposition 17, we have $\langle\bar{\sigma}_{\restriction h}\rangle_v \notin \mathbf{E}_{\alpha_*}(hv)$, and we get the existence of ρ' by definition of $\mathbf{E}_{\alpha_*}(hv)$. We continue to construct $\bar{\sigma}$ such that $\langle\bar{\sigma}_{\restriction h'}\rangle_{v'} = \rho'$, i.e. if $\rho' = \rho'_0\rho'_1 \ldots$, then $\sigma_i(h'\rho'_{\leq n}) = h'\rho'_{n+1}$ for all $\rho'_n \in V_i$ and $i \in \Pi$. For all non-empty prefixes g of ρ', we define $\gamma(h'g) = \rho'$ (notice that the prefixes of h' are already labeled).

Let us show that the constructed profile $\bar{\sigma}$ is a very weak SPE. Consider a history $hv \in \text{Hist}_i$ for some $i \in \Pi$, and a one-shot deviating strategy σ'_i from $\bar{\sigma}_{\restriction h}$ in the subgame $(\mathcal{G}_{\restriction h}, v)$. Let v' be such that $\sigma'_i(v) = v'$. By definition of $\bar{\sigma}$, we have $\gamma(hv) = \langle\bar{\sigma}_{\restriction g}\rangle_u$ for some history $gu \leq hv$ and $h \cdot \langle\bar{\sigma}_{\restriction h}\rangle_v = g \cdot \langle\bar{\sigma}_{\restriction g}\rangle_u$; and we have also $\gamma(hvv') = \langle\bar{\sigma}_{\restriction hv}\rangle_{v'}$. Moreover $\lambda_i(g \cdot \langle\bar{\sigma}_{\restriction g}\rangle_u) \leq \lambda_i(hv \cdot \langle\bar{\sigma}_{\restriction hv}\rangle_{v'})$ by (5), and $\lambda_i(hv \cdot \langle\bar{\sigma}_{\restriction hv}\rangle_{v'}) = \lambda_i(h \cdot \langle\sigma'_i, \sigma_{-i\restriction h}\rangle_v)$ because σ'_i is one-shot deviating. Therefore

$$\lambda_i(h \cdot \langle\bar{\sigma}_{\restriction h}\rangle_v) = \lambda_i(g \cdot \langle\bar{\sigma}_{\restriction g}\rangle_u) \leq \lambda_i(hv \cdot \langle\bar{\sigma}_{\restriction hv}\rangle_{v'}) = \lambda_i(h \cdot \langle\sigma'_i, \sigma_{-i\restriction h}\rangle_v)$$

which shows that $\bar{\sigma}_{\restriction h}$ is a very weak NE in $(\mathcal{G}_{\restriction h}, v)$. Hence $\bar{\sigma}$ is a very weak SPE, and thus also a weak SPE. ◀

4 Quantitative Reachability Games

In this section, we focus on quantitative reachability games, such that the player i cost of a play is the number of edges to reach his target set T_i (see Definition 2). Recall that for those games, SPEs, weak SPEs, and very weak SPEs, are equivalent notions (see Corollary 15). It is known that there always exists an SPE in quantitative reachability games [1, 7].

▶ **Theorem 21** ([1, 7]). *Each quantitative reachability game $(\mathcal{G}, v_0)$ has an SPE.*

As SPEs and weak SPEs coincide in quantitative reachability games, we get the next result by Theorem 18.

▶ **Corollary 22.** *Let $(\mathcal{G}, v_0)$ be a quantitative reachability game. The sets $\mathbf{P}_{\alpha_*}(hv)$ are non-empty, for all $hv \in \text{Hist}(v_0)$, and $\mathbf{P}_{\alpha_*}(v_0)$ is the set of outcomes of SPEs in $(\mathcal{G}, v_0)$.*

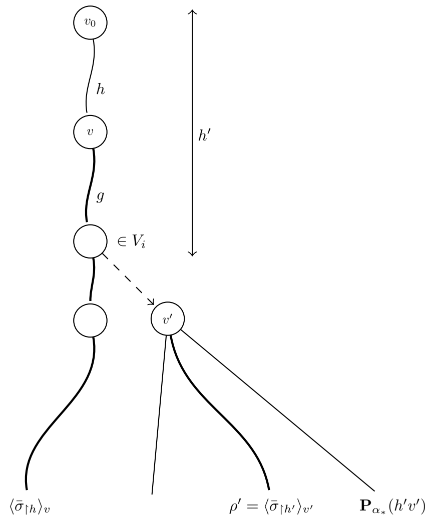

Figure 5 Construction of a very weak SPE $\bar{\sigma}$.

The proof provided for Theorem 21 is non constructive since it relies on topological arguments. Our main result is that one can algorithmically construct an SPE in $(\mathcal{G}, v_0)$ that is moreover finite-memory, thanks to the sets $\mathbf{P}_{\alpha_*}(hv)$.

▶ **Theorem 23.** *Each quantitative reachability initialized game $(\mathcal{G}, v_0)$ has a finite-memory SPE. Moreover there is an algorithm to construct such an SPE.*

We can also decide whether there exists a (finite-memory) SPE such that the cost of its outcome is component-wise bounded by a given constant vector.

▶ **Corollary 24.** *Let $(\mathcal{G}, v_0)$ be a quantitative reachability initialized game, and let $\bar{c} \in \mathbb{N}^{|\Pi|}$ be a given $|\Pi|$-tuple of integers. Then one can decide whether there exists a (finite-memory) SPE $\bar{\sigma}$ such that $\lambda_i(\langle\bar{\sigma}\rangle_{v_0}) \leq c_i$ for all $i \in \Pi$.*

The main ingredients of the proof of Theorem 23 are the next ones. They will be a little detailed in the sequel of this section, as well as the proof of Corollary 24.

1. Given α, the infinite number of sets $\mathbf{P}_\alpha(hv)$ can be replaced by the finite number of sets $\mathbf{P}^I_\alpha(v)$ where I is the set of players that did not reach their target set along history h.
2. The fixpoint of Proposition 17 is reached with some natural number $\alpha_* \in \mathbb{N}$.
3. Each $\mathbf{P}^I_\alpha(v)$ is a non-empty ω-regular set, thus containing a "lasso play" of the form $h{\cdot}g^\omega$.
4. The lasso plays of each $\mathbf{P}^I_\alpha(v)$ allow to construct a finite-memory SPE.

1. Sets $\mathbf{P}^I_\alpha(v)$. Let $(\mathcal{G}, v_0)$ be a quantitative reachability game. Given a history $h = h_0 \dots h_n$ in $(\mathcal{G}, v_0)$, we denote by $I(h)$ the set of players i such that $\forall k$, $0 \leq k \leq n$, we have $h_k \notin T_i$. In other words $I(h)$ is the set of players that did not reach their target set along history h. If h is empty, then $I(h) = \Pi$. One can prove that sets $\mathbf{P}_\alpha(hv)$ only depend on v and $I(h)$, and thus not on h (we do no longer take care of players that have reached their target set along h). We can thus introduce the notations

$$\mathbf{P}^I_\alpha(v) \quad (\text{resp. } \mathbf{E}^I_\alpha(v))$$

in place of $\mathbf{P}_\alpha(hv)$ (resp. $\mathbf{E}_\alpha(hv)$). In particular, $\mathbf{P}^\Pi_\alpha(v_0) = \mathbf{P}_\alpha(v_0)$ and $\mathbf{E}^\Pi_\alpha(v_0) = \mathbf{E}_\alpha(v_0)$. Hence given α, the infinite number of sets $\mathbf{P}_\alpha(hv)$ can be replaced by the finite number of sets $\mathbf{P}^I_\alpha(v)$.

A key result in the proof of Theorem 23 is the following one: given $\mathbf{P}^I_\alpha(v)$ and $i \in I$, if for all $\rho \in \mathbf{P}^I_\alpha(v)$, we have $\lambda_i(\rho) < +\infty$, then there exists c such that for all $\rho \in \mathbf{P}^I_\alpha(v)$, we have $\lambda_i(\rho) \leq c$. The constant c only depends on α, I, v, and i. As a consequence of this result, we have that $\sup\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v)\}$ is equal to $\max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v)\}$. We use the next notation for this maximum: given $\mathbf{P}^I_\alpha(v)$, we define $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ such that

$$\Lambda_i(\mathbf{P}^I_\alpha(v)) = \begin{cases} -1 & \text{if } i \notin I, \\ \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v)\} & \text{if } i \in I. \end{cases}$$

In this definition, -1 indicates that player i has already visited his target set T_i, and the max belongs to $\mathbb{N} \cup \{+\infty\}$.

2. Fixpoint with $\alpha_* \in \mathbb{N}$. To explain why the fixpoint (when computing the sets $\mathbf{P}^I_\alpha(v)$, see Proposition 17) is reached in a finite number of steps, we need to adapt previous notation $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ to mention the maximum costs for plays in $\mathbf{P}^I_\alpha(v)$ *starting with edge* (v, v'):

$$\Lambda_i(\mathbf{P}^I_\alpha(v), v') = \begin{cases} -1 & \text{if } i \notin I, \\ \max\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v) \text{ and } \rho_0\rho_1 = vv'\} & \text{if } i \in I. \end{cases} \tag{6}$$

In this definition, the max is equal to -1 when $\{\lambda_i(\rho) \mid \rho \in \mathbf{P}^I_\alpha(v) \text{ and } \rho_0\rho_1 = vv'\} = \emptyset$. One can prove that if $\mathbf{P}^I_\alpha(v) \neq \mathbf{P}^I_{\alpha+1}(v)$, then there exist $J \subseteq \Pi$ and $(u,u') \in E$ such that $\overline{\Lambda}(\mathbf{P}^J_\alpha(u),u') \neq \overline{\Lambda}(\mathbf{P}^J_{\alpha+1}(u),u')$. Notice that there is a finite number of sequences $(\overline{\Lambda}(\mathbf{P}^I_\alpha(v),v'))_\alpha$ as they only depend on $I \subseteq \Pi$ and $(v,v') \in E$, and that they are nonincreasing for the usual component-wise ordering over $(\mathbb{N} \cup \{-1,+\infty\})^\Pi$. As this ordering is a well-quasi-ordering, there exists an integer (and not only an ordinal) α'_* such that $\overline{\Lambda}(\mathbf{P}^I_{\alpha'_*}(v),v') = \overline{\Lambda}(\mathbf{P}^I_{\alpha'_*+1}(v),v')$ for all $I \subseteq \Pi$ and $(v,v') \in E$. We get that $\alpha_* \leq \alpha'_*$, showing that $\alpha_* \in \mathbb{N}$.

3. The sets $\mathbf{P}^I_\alpha(v)$ are ω-regular. We prefer to show that each set $\mathbf{P}^I_\alpha(v)$ is MSO-definable, instead of providing the (technical) construction of a Büchi automaton, It is well-known that a set of ω-words is ω-regular iff it is MSO-definable, by Büchi theorem [18]. Moreover from the Büchi automaton, one can construct an equivalent MSO-sentence, and conversely. One can also decide whether an MSO-sentence is satisfiable [18]. We recall that MSO-logic uses, in addition to variables $x, y, \ldots$ ($X, Y, \ldots$ resp.) that describe a position (a set of positions resp.) in an ω-word ρ, the relations $x < y$, $Succ(x,y)$, $X(x)$ to mention that $x \in X$, and $Q_u(x)$ to mention that vertex u is at position x of ρ.

As a first step, we show that if $\mathbf{P}^I_\alpha(v)$ is MSO-definable, say by sentence ϕ, then $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$ is computable. By definition $\Lambda_i(\mathbf{P}^I_\alpha(v))$ equals -1 if $i \notin I$, and is thus computable in this case. Given $i \in I$, one can decide whether $\Lambda_i(\mathbf{P}^I_\alpha(v)) = +\infty$ by checking whether $\phi \wedge \varphi$ is satisfiable, with $\varphi = \forall x \cdot \neg(\vee_{u \in T_i} Q_u(x))$. In case of a positive answer, $\Lambda_i(\mathbf{P}^I_\alpha(v))$ is thus computable. If not, one can prove that $\Lambda_i(\mathbf{P}^I_\alpha(v)) < n$ where n is the number of states of a Büchi automaton accepting $\mathbf{P}^I_\alpha(v)$. We can then similarly test whether $\Lambda_i(\mathbf{P}^I_\alpha(v)) = c$ by considering decreasing constants c from $n-1$ to 0, and thus compute $\Lambda_i(\mathbf{P}^I_\alpha(v))$.

As a second step, we show that each $\mathbf{P}^I_\alpha(v)$ is MSO-definable by induction on α. For $\alpha = 0$, as $\mathbf{P}^I_0(v)$ is the set of plays starting with v, the required sentence is $Q_v(0) \wedge \forall x \cdot \vee_{(u,u') \in E}(Q_u(x) \wedge Q_{u'}(x+1))$. Let $\alpha \in \mathbb{N}$, and suppose that $\mathbf{P}^I_\alpha(v)$ is a fixed MSO-definable set. To show that $\mathbf{P}^I_{\alpha+1}(v)$ is also MSO-definable, it is enough to show that $\mathbf{E}^I_\alpha(v)$ is MSO-definable. Thanks to $\overline{\Lambda}(\mathbf{P}^I_\alpha(v))$, the definition of $\rho \in \mathbf{E}^I_\alpha(v)$ (see Definition 16) can be rephrased as follows: there exist $n \in \mathbb{N}$, $i \in I$, and $u, u', v' \in V$ with $u' \neq v'$, $(u,v') \in E$, such that $\rho_n = u \in V_i$, $\rho_{n+1} = u'$, and $\lambda_i(\rho) > \Lambda_i(\mathbf{P}^{J'}_\alpha(v')) + (n+1)$, where J' is the subset of players of I that did not visit their target set along $\rho_{\leq n}$. Notice that this inequality implies that $i \in J'$ and $\Lambda_i(\mathbf{P}^{J'}_\alpha(v')) < +\infty$. The sentence ψ defining $\mathbf{E}^I_\alpha(v)$ is the following one:

$$\exists n \cdot \bigvee_{\substack{u,u' \neq v' \in V \\ (u,v') \in E}} \bigvee_{J' \subseteq I} \bigvee_{\substack{i \in J', u \in V_i \\ \Lambda_i(\mathbf{P}^{J'}_\alpha(v')) < +\infty}} (Q_u(n) \wedge Q_{u'}(n+1) \wedge \phi_{J',n} \wedge \varphi_{J',n,v',i}) \,.$$

In sentence ψ, we use $\phi_{J',n}$ expressing the definition of J', and $\varphi_{J',n,v',i}$ expressing that if player i visits its target set along ρ, it is after $\Lambda_i(\mathbf{P}^{J'}_\alpha(v')) + n + 1$ edges from ρ_0. Notice that the (computable) constant $\Lambda_i(\mathbf{P}^{J'}_\alpha(v'))$ can be considered as fixed, since $\mathbf{P}^I_\alpha(v)$ is fixed.

As a third step, as each $\mathbf{P}^I_\alpha(v)$ is ω-regular, then for all $i \in I$, $\mathbf{P}^I_{\alpha_*}(v)$ has a computable lasso play $\rho = h_{i,I,v}(g_{i,I,v})^\omega$ (depending on i, I, and v) with maximal cost $\lambda_i(\rho) = \Lambda_i(\mathbf{P}^I_{\alpha_*}(v))$.

4. Construction of a finite-memory SPE. We have all the ingredients to prove that each quantitative reachability game has a computable finite-memory SPE. The procedure is the same as the one developed in the proof of Lemma 20, except that it uses the lasso plays $h_{i,I,v}(g_{i,I,v})^\omega \in \mathbf{P}^I_{\alpha_*}(v)$. For the initial history v_0, we use any play $h_{i,\Pi,v_0}(g_{i,\Pi,v_0})^\omega \in \mathbf{P}^\Pi_{\alpha_*}(v_0)$, $i \in \Pi$. At the following steps, given a not yet labeled history $h'v'$, the proof of Lemma 20 requires to choose a play $\rho' \in \mathbf{P}^{J'}_{\alpha_*}(v')$ (for a certain $J' \subseteq I$) with a cost $\lambda_i(h'\rho')$ sufficiently large. We simply choose $\rho' = h_{i,J',v'} \cdot (g_{i,J',v'})^\omega$ that has maximal cost $\lambda_i(\rho') = \Lambda_i(\mathbf{P}^{J'}_{\alpha_*}(v'))$. This strategy profile $\bar{\sigma}$ is an SPE that is finite-memory since it depends on the finite number of lasso plays $h_{j,I,v} \cdot (g_{j,I,v})^\omega$.

It remains to prove the decidability of the constrained existence of an SPE for quantitative reachability games, as announced in Corollary 24. Let $\bar{c} \in \mathbb{N}^{|\Pi|}$ be a constant vector. We know that the set $\mathbf{P}^{\Pi}_{\alpha_*}(v_0)$ of outcomes of SPEs in $(\mathcal{G}, v_0)$ is MSO-definable. It is easy to see that the set $\{\rho \mid \rho \in \mathbf{P}^{\Pi}_{\alpha_*}(v_0) \text{ and } \lambda_i(\rho) \leq c_i, \forall i\}$ is also MSO-definable. We can thus decide whether this set is non-empty, which means that the constrained existence of an SPE is decidable. In case of positive answer, this set contains a lasso play $h \cdot g^\omega$. Exactly as done above, one can construct a finite-memory SPE $\bar{\sigma}$ such that $\langle \bar{\sigma} \rangle_{v_0} = h \cdot g^\omega$.

To conclude this section, we present another large class of games for which one can decide whether there exists a weak SPE.[8] The hypotheses are general conditions on the cost functions λ_i, $i \in \Pi$:

▶ **Theorem 25.** *Let $(\mathcal{G}, v_0)$ be an initialized game such that:*
- *each cost function λ_i is prefix-independent, and with finite range $C_i \subset \mathbb{Q}$,*
- *for all $i \in \Pi$, $c_i \in C_i$, and $v \in V$, the set of plays ρ in $(\mathcal{G}, v)$ with $\lambda_i(\rho) = c_i$ is an ω-regular set.*

Then one can decide whether $(\mathcal{G}, v_0)$ has a weak SPE $\bar{\sigma}$ (resp. such that $\lambda_i(\langle \bar{\sigma} \rangle_{v_0}) \leq c_i$ forall i for given $c_i \in C_i$, $i \in \Pi$). In case of positive answer, one can construct such a finite-memory weak SPE.

For example, the hypotheses of this theorem are satisfied by the liminf games and the limsup games; they are also satisfied by the game of Example 12. The proof of this theorem shares similar points with the proof given for quantitative reachability games. Again, it uses our Folk Theorem for weak SPEs.

5 Conclusion and Future Work

In this article, we have studied the existence of (weak) SPEs in quantitative games. We have proposed a Folk Theorem that characterizes all the outcomes of weak SPEs. To illustrate the potential of this theorem, we have given two applications. The first one is concerned with quantitative reachability games for which we have provided an algorithm to compute a finite-memory SPE, and a second algorithm for deciding the constrained existence of a (finite-memory) SPE. The second application is concerned with another large class of games for which we have proved that the (constrained) existence of a (finite-memory) weak SPE is decidable.

Future possible directions of research are the following ones. We would like to study the complexities of the problems studied for the two classes of games. We also want to investigate the application of our Folk Theorem to other classes of games. The example of Figure 3 is a game with a weak SPE but no SPE (see Example 12). Recall that for this game, the cost $\lambda_i(\rho)$ can be seen as either the mean-payoff, or the liminf, or the limsup, of the weights of ρ. We do not know if games with this kind of payoff functions always have a weak SPE or not. Notice that using a variant of the techniques developed for weak SPEs, we were able to obtain a Folk theorem for SPEs, nevertheless with a weaker characterization since one implication requires that the cost functions are upper-semicontinuous (see [2]).

[8] Contrarily to quantitative reachability games, we do not know if a weak SPE always exists for games in this class.

References

1 Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9, 2012.

2 Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Weak subgame perfect equilibria and their application to quantitative reachability. *CoRR*, abs/1504.01557, 2015.

3 Thomas Brihaye, Julie De Pril, and Sven Schewe. Multiplayer cost games with simple Nash equilibria. In *LFCS*, volume 7734 of *Lecture Notes in Comput. Sci.*, pages 59–73. Springer, 2013.

4 Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Secure equilibria in weighted games. In *CSL-LICS*, pages 26:1–26:26. ACM, 2014.

5 Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. Quantitative languages. *ACM Trans. Comput. Log.*, 11, 2010.

6 János Flesch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Math. Oper. Res.*, 35:742–755, 2010.

7 Drew Fudenberg and David Levine. Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory*, 31:251–268, 1983.

8 Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

9 Miroslav Klimos, Kim G. Larsen, Filip Stefanak, and Jeppe Thaarup. Nash equilibria in concurrent priced games. In *LATA*, volume 7183 of *Lecture Notes in Comput. Sci.*, pages 363–376. Springer, 2012.

10 H.W. Kuhn. Extensive games and the problem of information. *Classics in Game Theory*, pages 46–68, 1953.

11 Jean-François Mertens. Repeated games. In *Internat. Congress Mathematicians*, American Mathematical Society, pages 1528–1577. Providence, RI, 1987.

12 J. F. Nash. Equilibrium points in n-person games. In *PNAS*, volume 36, pages 48–49. National Academy of Sciences, 1950.

13 Martin J. Osborne and Ariel Rubinstein. *A course in Game Theory*. MIT Press, Cambridge, MA, 1994.

14 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.

15 Stéphane Le Roux and Arno Pauly. Infinite sequential games with real-valued payoffs. In *CSL-LICS*, pages 62:1–62:10. ACM, 2014.

16 Ariel Rubinstein. Comments on the interpretation of game theory. *Econometrica*, 59:909–924, 1991.

17 Eilon Solan and Nicolas Vieille. Deterministic multi-player dynkin games. *Journal of Mathematical Economics*, 39:911–929, 2003.

18 Wolfgang Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B)*, pages 133–192. Elsevier Science Publishers, Amsterdam, 1990.

19 Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *FSTTCS*, volume 4337 of *Lecture Notes in Comput. Sci.*, pages 212–223. Springer, 2006.