

The Complexity of Subgame Perfect Equilibria in Quantitative Reachability Games

Thomas Brihaye

Université de Mons (UMONS), Belgium

Véronique Bruyère

Université de Mons (UMONS), Belgium

Aline Goeminne

Université de Mons (UMONS) and Université libre de Bruxelles (ULB), Belgium

Jean-François Raskin

Université libre de Bruxelles (ULB), Belgium

Marie van den Bogaard

Université libre de Bruxelles (ULB), Belgium

Abstract

We study multiplayer quantitative reachability games played on a finite directed graph, where the objective of each player is to reach his target set of vertices as quickly as possible. 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. It is known that there always exists an SPE in quantitative reachability games and that the constrained existence problem is decidable. We here prove that this problem is PSPACE-complete. To obtain this result, we propose a new algorithm that iteratively builds a set of constraints characterizing the set of SPE outcomes in quantitative reachability games. This set of constraints is obtained by iterating an operator that reinforces the constraints up to obtaining a fixpoint. With this fixpoint, the set of SPE outcomes can be represented by a finite graph of size at most exponential. A careful inspection of the computation allows us to establish PSPACE membership.

2012 ACM Subject Classification Software and its engineering → Formal methods; Theory of computation → Logic and verification; Theory of computation → Solution concepts in game theory

Keywords and phrases multiplayer non-zero-sum games played on graphs ; quantitative reachability objectives ; subgame perfect equilibria ; constrained existence problem

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2019.9

Related Version A full version of the paper is available at <https://arxiv.org/abs/1905.00784>.

Funding Work partially supported by the PDR project *Subgame perfection in graph games* (F.R.S.-FNRS), the ARC project *Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond* (Fédération Wallonie-Bruxelles), the EOS project *Verifying Learning Artificial Intelligence Systems* (F.R.S.-FNRS & FWO), and the COST Action 16228 *GAMENET* (European Cooperation in Science and Technology).

1 Introduction

While two-player zero-sum games played on graphs are the most studied model to formalize and solve the reactive synthesis problem [26], recent work has considered non-zero-sum extensions of this mathematical framework, see e.g. [15, 19, 7, 22, 6, 5, 17, 4, 1], see also the surveys [21, 3, 13]. In the zero-sum game approach, the system and the environment are considered as *monolithic* and fully *adversarial* entities. Unfortunately, both assumptions may turn to be too strong. First, the reactive system may be composed of several components that execute concurrently and have their own purpose. So, it is natural to model such



© Thomas Brihaye, Véronique Bruyère, Aline Goeminne, Jean-François Raskin and Marie van den Bogaard;

licensed under Creative Commons License CC-BY

30th International Conference on Concurrency Theory (CONCUR 2019).

Editors: Wan Fokkink and Rob van Glabbeek; Article No. 9; pp. 9:1–9:16



Leibniz International Proceedings in Informatics

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

systems with *multiplayer* games with each player having his own objective. Second, the environment usually has its own objective too, and this objective is usually not the negation of the objective of the reactive system as postulated in the zero-sum case. Therefore, there are instances of the reactive synthesis problem for which no solution exists in the zero-sum setting, i.e. no winning strategy for the system against a completely antagonistic environment, while there exists a strategy for the system which enforces the desired properties against all *rational* behaviors of the environment pursuing its own objective.

While the central solution concept in zero-sum games is the notion of *winning strategy*, it is well known that this concept is not sufficient to reason about non-zero-sum games. In non-zero-sum games, notions of equilibria are used to reason about the rational behavior of players. The celebrated notion of *Nash equilibrium* (NE) [24] is one of the most studied. A strategy profile is an NE if no player has an incentive to deviate, i.e. change his strategy and obtain a better reward, when this player knows that the other players will be playing their respective strategies in the profile. A well-known weakness of NE in sequential games, which include infinite duration games played on graphs, is that they are subject to *non-credible threats*: decisions in subgames that are irrational and used to threaten the other players and oblige them to follow a given behavior. To avoid this problem, the concept of *subgame perfect equilibria* (SPE) has been proposed, see e.g. [25]. SPEs are NEs with the additional property that they are also NEs in all subgames of the original game. While it is now quite well understood how to handle NEs algorithmically in games played on graphs [28, 29, 12, 17], this is not the case for SPEs.

Contributions In this paper, we provide an algorithm to decide in *polynomial space* the constrained existence problem for SPEs in *quantitative reachability games*. A quantitative reachability game is played by n players on a finite graph in which each player has his own reachability objective. The objective of each player is to reach his target set of vertices as quickly as possible. In a series of papers devoted to quantitative reachability games, it has been shown that SPEs always exist [8], and that the set of SPE outcomes is a regular language which is effectively constructible [11]. As a consequence of the latter result, the constrained existence problem for SPEs is decidable.

Unfortunately, the proof that establishes the regularity of the set of possible SPE outcomes in [11] exploits a *well-quasi order* for proving termination that cannot be used to obtain a good upper bound on the complexity for the algorithm. Here, we propose a new algorithm and we show that this set of outcomes can be represented using an automaton of size at most exponential. It follows that the constrained existence problem for SPEs is in PSPACE. We also provide a matching lower-bound showing that this problem is PSPACE-complete.

Our new algorithm iteratively builds a set of constraints that exactly characterize the set of SPEs in quantitative reachability games. This set of constraints is obtained by iterating an operator that reinforces the constraints up to obtaining a fixpoint. A careful inspection of the computation allows us to establish PSPACE membership.

Related work Algorithms to reason on NEs in graph games are studied in [28] for ω -regular objectives and in [29, 12] for quantitative objectives. Algorithms to reason on SPEs are given in [27] for ω -regular objectives. Quantitative reachability objectives are not ω -regular objectives. Reasoning about NEs and SPEs for ω -regular specifications can also be done using strategy logics [16, 23]. Other notions of rationality used for reactive synthesis have been studied in the literature: rational synthesis in cooperative [19] and adversarial [22] setting, and their algorithmic complexity in [17]. Extensions with imperfect information have

been investigated in [18]. Synthesis rules based on the notion of admissible strategies have been studied in [2, 7, 6, 5, 4, 1]. Weak SPEs have been studied in [11, 14, 9] and shown to be equivalent to SPEs for quantitative reachability objectives. Fixpoint techniques are used in [20, 11, 14] to establish the existence of (weak) SPEs in some classes of games, however they cannot be used in our context to get complexity results.

2 Preliminaries

In this section, we recall the notions of quantitative reachability game and subgame perfect equilibrium. We also state the problem studied in this paper and our main result.

Quantitative reachability games An *arena* is a tuple $G = (\Pi, V, (V_i)_{i \in \Pi}, E)$ where $\Pi = \{1, 2, \dots, n\}$ is a finite set of n players, V is a finite set of vertices with $|V| \geq 2$, $(V_i)_{i \in \Pi}$ is a partition of V between the players, and $E \subseteq V \times V$ is a set of edges such that for all $v \in V$ there exists $v' \in V$ such that $(v, v') \in E$. Without loss of generality, we suppose that $|\Pi| \leq |V|$. We denote by $\text{Succ}(v) = \{v' \mid (v, v') \in E\}$ the set of *successors* of v , for $v \in V$, and by Succ^* the transitive closure of Succ .

A *play* in G is an infinite sequence of vertices $\rho = \rho_0 \rho_1 \dots$ such that for all $k \in \mathbb{N}$, $(\rho_k, \rho_{k+1}) \in E$. A *history* is a finite sequence $h = h_0 h_1 \dots h_k$ with $k \in \mathbb{N}$ defined similarly. The *length* $|h|$ of h is the number k of its edges. We denote the set of plays by Plays and the set of histories by Hist (when it is necessary, we use notation Plays_G and Hist_G to recall the underlying arena G). Moreover, the set Hist_i is the set of histories such that their last vertex v is a vertex of player i , i.e. $v \in V_i$.

Given a play $\rho \in \text{Plays}$ and $k \in \mathbb{N}$, the prefix $\rho_0 \rho_1 \dots \rho_k$ of ρ is denoted by $\rho_{\leq k}$ and its suffix $\rho_k \rho_{k+1} \dots$ by $\rho_{\geq k}$. A play ρ is called a *lasso* if it is of the form $\rho = h\ell^\omega$ with $h\ell \in \text{Hist}$. Notice that ℓ is not necessary a simple cycle. The *length* of a lasso $h\ell^\omega$ is the length of $h\ell$.


A *quantitative game* $\mathcal{G} = (G, (\text{Cost}_i)_{i \in \Pi})$ is an arena equipped with a cost function profile $\text{Cost} = (\text{Cost}_i)_{i \in \Pi}$ such that each function $\text{Cost}_i : \text{Plays} \rightarrow \mathbb{R} \cup \{+\infty\}$ assigns a cost to each play. In a quantitative game \mathcal{G} , an initial vertex $v_0 \in V$ is often fixed, and we call (\mathcal{G}, v_0) an *initialized game*. A play (resp. a history) of (\mathcal{G}, v_0) is then a play (resp. a history) of \mathcal{G} starting in v_0 . The set of such plays (resp. histories) is denoted by $\text{Plays}(v_0)$ (resp. $\text{Hist}(v_0)$). We also use notation $\text{Hist}_i(v_0)$ when these histories end in a vertex $v \in V_i$.

In this article we are interested in *quantitative reachability games* such that each player has a target set of vertices that he wants to reach. The cost to pay is equal to the number of edges to reach the target set, and each player aims at minimizing his cost.

► **Definition 1** (Quantitative reachability game). *A quantitative reachability game (or simply a reachability game) is a tuple $\mathcal{G} = (G, (F_i)_{i \in \Pi}, (\text{Cost}_i)_{i \in \Pi})$ such that (i) G is an arena, (ii) for each $i \in \Pi$, $F_i \subseteq V$ is the target set of player i , and (iii) for each $i \in \Pi$ and each $\rho = \rho_0 \rho_1 \dots \in \text{Plays}$, $\text{Cost}_i(\rho)$ is equal to the least index k such that $\rho_k \in F_i$, and to $+\infty$ if no such index exists.*

Given a quantitative game \mathcal{G} , a *strategy* for player i is a function $\sigma_i : \text{Hist}_i \rightarrow V$. It assigns to each history hv , with $v \in V_i$, a vertex v' such that $(v, v') \in E$. In an initialized game (\mathcal{G}, v_0) , σ_i needs only to be defined for histories starting in v_0 . A play $\rho = \rho_0 \rho_1 \dots$ is *consistent* with σ_i if for all $\rho_k \in V_i$, $\sigma_i(\rho_0 \dots \rho_k) = \rho_{k+1}$. A strategy σ_i is *positional* if it only depends on the last vertex of the history, i.e., $\sigma_i(hv) = \sigma_i(v)$ for all $hv \in \text{Hist}_i$.

A *strategy profile* is a tuple $\sigma = (\sigma_i)_{i \in \Pi}$ of strategies, one for each player. Given an initialized game (\mathcal{G}, v_0) and a strategy profile σ , there exists a unique play from v_0 consistent



■ **Figure 1** A quantitative reachability game: player 1 (resp. player 2) owns circle (resp. square) vertices and $F_1 = \{v_2\}$ and $F_2 = \{v_2, v_5\}$.

with each strategy σ_i . We call this play the *outcome* of σ and denote it by $\langle \sigma \rangle_{v_0}$. Let $c = (c_i)_{i \in \Pi} \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, we say that σ is a strategy profile *with cost* c or that $\langle \sigma \rangle_{v_0}$ has *cost* c if $c_i = \text{Cost}_i(\langle \sigma \rangle_{v_0})$ for all $i \in \Pi$.

Solution concepts and constraint problem In the multiplayer game setting, the solution concepts usually studied are *equilibria* (see [21]). We here recall the concepts of Nash equilibrium and subgame perfect equilibrium.

Let $\sigma = (\sigma_i)_{i \in \Pi}$ be a strategy profile in a game (\mathcal{G}, v_0) . When we highlight the role of player i , we denote σ by (σ_i, σ_{-i}) where σ_{-i} is the profile $(\sigma_j)_{j \in \Pi \setminus \{i\}}$. A strategy $\sigma'_i \neq \sigma_i$ is a *deviating* strategy of player i , and it is a *profitable deviation* for him if $\text{Cost}_i(\langle \sigma \rangle_{v_0}) > \text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$.

The notion of Nash equilibrium is classical: a strategy profile σ in an initialized game (\mathcal{G}, v_0) is a *Nash equilibrium* (NE) if no player has an incentive to deviate unilaterally from his strategy, i.e. no player has a profitable deviation. Formally, σ is an NE if for each $i \in \Pi$ and each deviating strategy σ'_i of player i , we have $\text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq \text{Cost}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$.

When considering games played on graphs, a useful refinement of NE is the concept of *subgame perfect equilibrium* (SPE) which is a strategy profile being an NE in each subgame. It is well-known that contrarily to NEs, SPEs avoid non-credible threats [21]. Formally, given a quantitative game $\mathcal{G} = (G, \text{Cost})$, an initial vertex v_0 , and a history $hv \in \text{Hist}(v_0)$, the initialized game $(\mathcal{G}_{\uparrow h}, v)$ is called a *subgame* of (\mathcal{G}, v_0) such that $\mathcal{G}_{\uparrow h} = (G, \text{Cost}_{\uparrow h})$ and $\text{Cost}_{i \uparrow h}(\rho) = \text{Cost}_i(h\rho)$ for all $i \in \Pi$ and $\rho \in V^\omega$. Notice that (\mathcal{G}, v_0) is subgame of itself. Moreover if σ_i is a strategy for player i in (\mathcal{G}, v_0) , then $\sigma_{i \uparrow h}$ denotes the strategy in $(\mathcal{G}_{\uparrow h}, v)$ such that for all histories $h' \in \text{Hist}_i(v)$, $\sigma_{i \uparrow h}(h') = \sigma_i(hh')$. Similarly, from a strategy profile σ in (\mathcal{G}, v_0) , we derive the strategy profile $\sigma_{\uparrow h}$ in $(\mathcal{G}_{\uparrow h}, v)$.

► **Definition 2** (Subgame perfect equilibrium). *A strategy profile σ is a subgame perfect equilibrium in an initialized game (\mathcal{G}, v_0) if for all $hv \in \text{Hist}(v_0)$, $\sigma_{\uparrow h}$ is an NE in $(\mathcal{G}_{\uparrow h}, v)$.*

It is proved in [8] that there always exists an SPE in reachability games. In this paper, we are interested in solving the following *constraint problem*.

► **Definition 3** (Constraint problem). *Given (\mathcal{G}, v_0) an initialized reachability game and two threshold vectors $x, y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, the constraint problem is to decide whether there exists an SPE in (\mathcal{G}, v_0) with cost c such that $x \leq c \leq y$, that is, $x_i \leq c_i \leq y_i$ for all $i \in \Pi$.*

Our main result is the following one. The paper is devoted to its proof.

► **Theorem 4.** *The constraint problem for initialized reachability games is PSPACE-complete.*

► **Example 5.** A reachability game \mathcal{G} with two players is depicted in Figure 1. The circle (resp. square) vertices are owned by player 1 (resp. player 2). The target sets of both players are respectively $F_1 = \{v_2\}$ (grey vertex) and $F_2 = \{v_2, v_5\}$ (double circled vertices).

The positional strategy profile $\sigma = (\sigma_1, \sigma_2)$ is depicted by double arrows, its outcome in (\mathcal{G}, v_0) is equal to $\langle \sigma \rangle_{v_0} = (v_0 v_1 v_6 v_7 v_2)^\omega$ with cost $(4, 4)$. Let us explain that σ is an

NE. Player 1 reaches his target set as soon as possible and has thus no incentive to deviate. Player 2 has no profitable deviation that allows him to reach v_5 . For instance if he uses a deviating positional strategy σ'_2 such that $\sigma'_2(v_0) = v_4$, then the outcome of (σ_1, σ'_2) is equal to $(v_0v_4)^\omega$ with cost $(+\infty, +\infty)$ which is not profitable for player 2. One can verify that the strategy profile σ is also an SPE. For instance in the subgame $(\mathcal{G}_{\upharpoonright h}, v_5)$ with $h = v_0v_4$, we have $\rho = \langle \sigma_{\upharpoonright h} \rangle_{v_5} = v_5v_4(v_0v_1v_6v_7v_2)^\omega$ such that $\text{Cost}_{\upharpoonright h}(\rho) = \text{Cost}(h\rho) = (8, 2)$. In this subgame, with ρ , both players reach their target set as soon as possible. ◀

Extended game We here present the *extended game* of a reachability game, such that the vertices are enriched with the set of players that have already visited their target sets along a history (see e.g. [9]). Working with this extended game is essential to prove our main result.

► **Definition 6** (Extended game). *Let $\mathcal{G} = (G, (F_i)_{i \in \Pi}, (\text{Cost}_i)_{i \in \Pi})$ be a reachability game with an arena $G = (\Pi, V, (V_i)_{i \in \Pi}, E)$, and let v_0 be an initial vertex. The extended game of \mathcal{G} is equal to $\mathcal{X} = (X, (F_i^X)_{i \in \Pi}, (\text{Cost}_i)_{i \in \Pi})$ with the arena $X = (\Pi, V^X, (V_i^X)_{i \in \Pi}, E^X)$, such that:*

- $V^X = V \times 2^\Pi$
- $((v, I), (v', I')) \in E^X$ if and only if $(v, v') \in E$ and $I' = I \cup \{i \in \Pi \mid v' \in F_i\}$
- $(v, I) \in V_i^X$ if and only if $v \in V_i$
- $(v, I) \in F_i^X$ if and only if $i \in I$
- for each $\rho \in \text{Plays}_X$, $\text{Cost}_i(\rho)$ is equal to the least index k such that $\rho_k \in F_i^X$, and to $+\infty$ if no such index exists.

The initialized extended game (\mathcal{X}, x_0) associated with the initialized game (\mathcal{G}, v_0) is such that $x_0 = (v_0, I_0)$ with $I_0 = \{i \in \Pi \mid v_0 \in F_i\}$.

Notice the way each target set F_i^X is defined: if $v \in F_i$, then $(v, I) \in F_i^X$ but also $(v', I') \in F_i^X$ for all $(v', I') \in \text{Succ}^*(v, I)$. The extended game of the reachability game of Figure 1 is depicted in Figure 2 (until Section 3 the reader should not consider the labeling indicated close to the vertices). We will come back to this example at the end of this section.

Let us state some properties of the extended game. First, notice that for each $\rho = (v_0, I_0)(v_1, I_1) \dots \in \text{Plays}_X(x_0)$, we have the next property called *I-monotonicity*:

$$I_k \subseteq I_{k+1} \quad \text{for all } k \in \mathbb{N}. \quad (1)$$

Second, given an initialized game (\mathcal{G}, v_0) and its extended game (\mathcal{X}, x_0) , there is a one-to-one correspondence between plays in $\text{Plays}_G(v_0)$ and plays in $\text{Plays}_X(x_0)$:

- from $\rho = \rho_0\rho_1 \dots \in \text{Plays}_G(v_0)$, we derive $\rho^X = (\rho_0, I_0)(\rho_1, I_1) \dots \in \text{Plays}_X(x_0)$ such that I_k is the set of players i that have seen their target set F_i along $\rho_{\leq k}$;
- from $\rho = (v_0, I_0)(v_1, I_1) \dots \in \text{Plays}_X(x_0)$, we derive $\rho^G = v_0v_1 \dots \in \text{Plays}_G(v_0)$ such that the second components $I_k, k \in \mathbb{N}$, are omitted.

Third, given $\rho \in \text{Plays}_G(v_0)$, we have that $\text{Cost}(\rho^X) = \text{Cost}(\rho)$, and conversely given $\rho \in \text{Plays}_X(x_0)$, we have that $\text{Cost}(\rho^G) = \text{Cost}(\rho)$. It follows that outcomes of SPE can be equivalently studied in (\mathcal{G}, v_0) and in (\mathcal{X}, x_0) , as stated in the next lemma.

► **Lemma 7.** *If ρ is the outcome of an SPE in (\mathcal{G}, v_0) , then ρ^X is the outcome of an SPE in (\mathcal{X}, x_0) with the same cost. Conversely, if ρ is the outcome of an SPE in (\mathcal{X}, x_0) , then ρ^G is the outcome of an SPE in (\mathcal{G}, v_0) with the same cost.*

By construction, the arena X of the initialized extended game is divided into different regions according to the players who have already visited their target set. Let us provide

some useful notions with respect to this decomposition. We will often use them in the following sections. Let $\mathcal{I} = \{I \subseteq \Pi \mid \text{there exists } v \in V \text{ such that } (v, I) \in \text{Succ}^*(x_0)\}$ be the set of sets I accessible from the initial state x_0 , and let $N = |\mathcal{I}|$ be its size. For $I, I' \in \mathcal{I}$ with $I \neq I'$, if there exists $((v, I), (v', I')) \in E^X$, we say that I' is a *successor* of I and we write $I' \in \text{Succ}(I)$. Given $I \in \mathcal{I}$, $X^I = (V^I, E^I)$ refers to the sub-arena of X restricted to the vertices $\{(v, I) \in V^X \mid v \in V\}$. We say that X^I is the *region*¹ associated with I . Such a region X^I is called a *bottom region* whenever $\text{Succ}(I) = \emptyset$.

There exists a partial order on \mathcal{I} such that $I < I'$ if and only if $I' \in \text{Succ}^+(I)$.

We fix an arbitrary *total order* on \mathcal{I} that extends this partial order $<$ as follows:

$$J_1 < J_2 < \dots < J_N. \tag{2}$$

(with X^{J_N} a bottom region).² With respect to this total order, given $n \in \{1, \dots, N\}$, we denote by $X^{\geq J_n} = (V^{\geq J_n}, E^{\geq J_n})$ the sub-arena of X restricted to the vertices $\{(v, I) \in V^X \mid I \geq J_n\}$. This total order together with the I-monotonicity leads to the following lemma.

► **Lemma 8** (Region decomposition and section). *Let π be a (finite or infinite) path in \mathcal{X} . Then there exists a region decomposition of π as $\pi[\ell]\pi[\ell+1] \dots \pi[m]$ with $1 \leq \ell \leq m \leq N$, such that for each n , $\ell \leq n \leq m$:*

- $\pi[n]$ is a (possibly empty) path in \mathcal{X} ,
- every vertex of $\pi[n]$ is of the form (v, J_n) for some $v \in V$.

Each path $\pi[n]$ is called a section. The last section $\pi[m]$ is infinite if and only if π is infinite.

► **Example 9.** Let us come back to the game (\mathcal{G}, v_0) of Figure 1. Its extended game (\mathcal{X}, x_0) is depicted in Figure 2 (only the part reachable from the initial vertex $x_0 = (v_0, \emptyset)$ is depicted).

The extended game is divided into three different regions: one region associated to $I = \emptyset$ that contains x_0 , a second region associated to $I = \{2\}$, and a third bottom region associated to $I = \Pi$. Hence the set $\mathcal{I} = \{\emptyset, \{2\}, \Pi\}$ is totally ordered as $J_1 = \emptyset < J_2 = \{2\} < J_3 = \Pi$. For all the vertices (v, I) of the region associated with $I = \{2\}$, we have $(v, I) \notin F_1^X$ and $(v, I) \in F_2^X$, and for those of the region associated with $I = \Pi$, we have $(v, I) \in F_1^X \cap F_2^X$.

From the SPE σ given in Example 5 with outcome $\rho = (v_0 v_1 v_6 v_7 v_2)^\omega \in \text{Plays}_{\mathcal{G}}(v_0)$ and cost $(4, 4)$, we derive the SPE outcome $\rho^X \in \text{Plays}_X(x_0)$ with the same cost and equal to

$$(v_0, \emptyset)(v_1, \emptyset)(v_6, \emptyset)(v_7, \emptyset)((v_2, \Pi)(v_0, \Pi)(v_1, \Pi)(v_6, \Pi)(v_7, \Pi))^\omega \tag{3}$$

3 Characterization

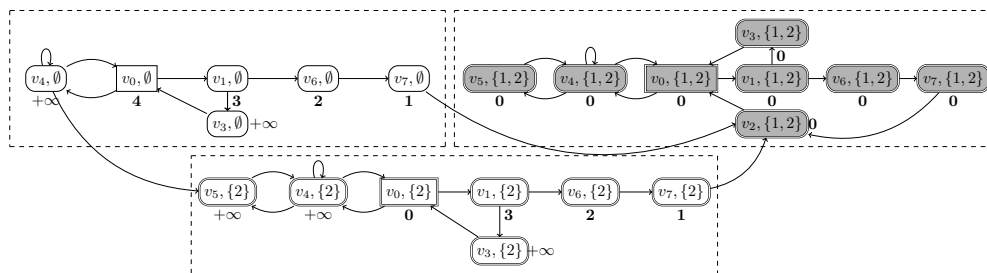
In this section, given an initialized reachability game (\mathcal{G}, v_0) , we characterize the set of plays that are outcomes of SPEs, and we provide an algorithm to construct this set. For this characterization, by Lemma 7, we can work on the extended game (\mathcal{X}, x_0) instead of (\mathcal{G}, v_0) .

All along this section, when we refer to a vertex of V^X , we use notation v (instead of (u, I)) and notation $I(v)$ means the second component I of this vertex.

Our algorithm iteratively builds a set of *constraints* imposed by a *labeling function* $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ such that the plays of the extended game satisfying those constraints are exactly the SPE outcomes. Let us provide a formal definition of such a function λ with the constraints that it imposes on plays.

¹ In the sequel, we indifferently call region either X^I , or V^I , or I .

² We use notation J_n , $n \in \{1, \dots, N\}$, to avoid any confusion with the sets I_k appearing in a play $\rho = (v_0, I_0)(v_1, I_1) \dots$



■ **Figure 2** The extended game (\mathcal{X}, x_0) for the initialized game (G, v_0) of Figure 1. The values of a labeling function λ are indicated close to each vertex.

► **Definition 10** (λ -consistent play). *Let \mathcal{X} be the extended game of a reachability game \mathcal{G} , and $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ be a labeling function. Given $v \in V^X$, for all plays $\rho \in \text{Plays}_X(v)$, we say that $\rho = \rho_0\rho_1 \dots$ is λ -consistent if for all $n \in \mathbb{N}$ and $i \in \Pi$ such that $\rho_n \in V_i$:*

$$\text{Cost}_i(\rho_{\geq n}) \leq \lambda(\rho_n). \quad (4)$$

We denote by $\Lambda(v)$ the set of plays $\rho \in \text{Plays}_X(v)$ that are λ -consistent.

Thus, a play ρ is λ -consistent if for all its suffixes $\rho_{\geq n}$, if player i owns ρ_n then the number of edges to reach his target set along $\rho_{\geq n}$ is bounded by $\lambda(\rho_n)$. Before going into the details of our algorithm, let us intuitively explain on an example how a well-chosen labeling function characterizes the set of SPE outcomes.

► **Example 11.** We consider the extended game (\mathcal{X}, x_0) of Figure 2, and a labeling function λ whose values are indicated under or next to each vertex. If $v \in V_i^X$ is labeled by $\lambda(v) = c$, then if $c \in \mathbb{N}_0$, this means that player i will only accept outcomes in (\mathcal{X}, v) that reach his target set within c steps, otherwise he would have a profitable deviation. If $\lambda(v) = 0$, this means that player i has already reached his target set, and if $\lambda(v) = +\infty$, player i has no profitable deviation whatever outcome is proposed to him.

In Example 9 was given the SPE outcome equal to $\rho^X(3)$ and with cost $(4, 4)$. We have $\lambda(v_0, \emptyset) = 4$ and player 2 reaches his target set from (v_0, \emptyset) within exactly 4 steps. The constraints imposed by λ on the other vertices of ρ are respected too. On the other hand, one can prove that $\rho' = ((v_0, \emptyset)(v_4, \emptyset))^\omega$ is the outcome of no SPE. It is not λ -consistent since player 2 does not reach his target set, and so in particular not within 4 steps. ◀

Our algorithm roughly works as follows: the labeling function λ that characterizes the set of SPE outcomes is obtained (i) from an initial labeling function that imposes no constraints, (ii) by iterating an operator that reinforces the constraints step after step, (iii) up to obtaining a fixpoint which is the required function λ . Thus, if λ^k is the labeling function computed at step k and $\Lambda^k(v)$, $v \in V^X$, the related sets of λ^k -consistent plays, initially we have $\Lambda^0(v) = \text{Plays}_X(v)$, and step by step, the constraints imposed by λ^k become stronger and the sets $\Lambda^k(v)$ become smaller, until a fixpoint is reached.

Initially, we want a labeling function λ^0 that imposes no constraint in a way to have $\Lambda^0(v) = \text{Plays}_X(v)$. We define $\lambda^0(v) = +\infty$ except when $i \in I(v)$ and $v \in V_i^X$ where $\lambda^0(v) = 0$. Indeed by Definition 6, we have $v \in F_i^X$ if and only if $i \in I(v)$. Hence, given $\rho = \rho_0\rho_1 \dots$, once $\rho_k \in F_i^X$ for some $k \in \mathbb{N}$ then $\rho_n \in F_i^X$ for all $n \geq k$. It follows that for all $n \geq k$, $\text{Cost}_i(\rho_{\geq n}) = 0$ and the inequality (4) is trivially true. (See also Example 11.)

► **Definition 12** (Initial labeling). *For all $v \in V^X$, let $i \in \Pi$ be such that $v \in V_i^X$,*

$$\lambda^0(v) = 0 \text{ if } i \in I(v) \quad \text{and} \quad \lambda^0(v) = +\infty \text{ otherwise.}$$

Let us now explain how our algorithm computes the labeling functions λ^k , $k \geq 1$, and the related sets $\Lambda^k(v)$, $v \in V^X$. It works in a *bottom-up* manner, according to the total order $J_1 < J_2 < \dots < J_N$ of \mathcal{I} given in (2). It first iteratively updates the labeling function for all vertices v of the arena X^{J_N} . At some point, the values of λ^k do not change anymore in X^{J_N} and $(\lambda^k)_{k \in \mathbb{N}}$ reaches locally (on X^{J_N}) a fixpoint. Then it treats the arena $X^{\geq J_{N-1}}$ and in the same way, it updates locally the values of λ^k in $X^{\geq J_{N-1}}$ until reaching a (local) fixpoint. It then repeats this procedure in $X^{\geq J_{N-2}}, \dots, X^{\geq J_1} = X$.

Hence suppose we currently treat the arena $X^{\geq J_n}$ and we want to compute λ^{k+1} from λ^k . We define the updated function λ^{k+1} as follows (with the convention that $1 + (+\infty) = +\infty$).

► **Definition 13** (Labeling update). *Let $k \geq 0$ and suppose that we treat the arena $X^{\geq J_n}$, with $n \in \{1, \dots, N\}$. For all $v \in V^X$,*

- *if $v \in V^{\geq J_n}$, let $i \in \Pi$ be such that $v \in V_i^X$, then*

$$\lambda^{k+1}(v) = 0 \text{ if } i \in I(v) \quad \text{and} \quad \lambda^{k+1}(v) = 1 + \min_{(v,v') \in E^X} \sup\{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v')\} \text{ otherwise.}$$
- *if $v \notin V^{\geq J_n}$, then $\lambda^{k+1}(v) = \lambda^k(v)$.*

► **Remark 14.** If the sup in Definition 13 is equal to $+\infty$ then there exists $\rho \in \Lambda^k(v')$ such that $\text{Cost}_i(\rho) = +\infty$. Thus the sup can be replaced by a max which belongs to $\mathbb{N} \cup \{+\infty\}$.

Let us provide some explanations. As this update concerns the arena $X^{\geq J_n}$, we keep $\lambda^{k+1} = \lambda^k$ outside of this arena. Suppose now that v belongs to the arena $X^{\geq J_n}$ and $v \in V_i^X$. We define $\lambda^{k+1}(v) = 0$ whenever $i \in I(v)$ (as already explained for the definition of λ^0). When it is updated, the value $\lambda^{k+1}(v)$ represents what is the best cost that player i can ensure for himself with a “one-shot” choice by only taking into account plays of $\Lambda^k(v')$ with $v' \in \text{Succ}(v)$. Notice that it makes sense to run the algorithm in a bottom-up fashion according to the total ordering $J_1 < \dots < J_N$ since given a play $\rho = \rho_0 \rho_1 \dots$, if ρ_0 is a vertex of $V^{\geq J_n}$, then for all $k \in \mathbb{N}$, ρ_k is a vertex of $V^{\geq J_n}$ (by I -monotonicity). Moreover running the algorithm in this way is essential to prove that the constraint problem is in PSPACE.

We can now provide our algorithm that computes the sequence $(\lambda^k(v))_{k \in \mathbb{N}}$. From the last computed λ^k , we derive the sets $\Lambda^k(v)$, $v \in V^X$, that we need for the characterization of outcomes of SPEs (see Theorem 17 below). Such a characterization already appears in [11] however with a different algorithm that cannot be used to obtain good complexity upper bounds for the constraint problem as done in this paper. Nevertheless, the proof of our characterization and that of [11] are similar.

Algorithm 1: Fixpoint

```

 $k \leftarrow 0; n \leftarrow N$ ; compute  $\lambda^0$  (see Definition 12)
while  $n \neq 0$  do
  repeat
     $k \leftarrow k + 1$ ; compute  $\lambda^k$  from  $\lambda^{k-1}$  with respect to  $X^{\geq J_n}$  (see Definition 13)
  until  $\lambda^k = \lambda^{k-1}$ 
   $n \leftarrow n - 1$ 
end
return  $\lambda^k$ .

```

As already announced, the sequence $(\lambda^k)_{k \in \mathbb{N}}$ computed by this algorithm reaches a fixpoint - locally on each arena $X^{\geq J_n}$ and globally on X - in the following meaning:

- **Proposition 15.** *There exists a sequence $0 = k_N^* < k_{N-1}^* < \dots < k_1^* = k^*$ such that*
- **Local fixpoint:** *for all $J_n \in \mathcal{I}$, all $m \in \mathbb{N}$ and all $v \in V^{\geq J_n}$: $\lambda^{k_n^*+m}(v) = \lambda^{k_n^*}(v)$.*
 - **Global fixpoint:** *with $k^* = k_1^*$, for all $m \in \mathbb{N}$ and all $v \in V^X$: $\lambda^{k^*+m}(v) = \lambda^{k^*}(v)$.*
The global fixpoint λ^{k^} is also simply denoted by λ^* , and the set Λ^{k^*} is denoted by Λ^* .*

The fixpoints follow from the existence of a well-quasi-ordering on the non increasing sequences $(\lambda^k(v))_{k \in \mathbb{N}}$. Proposition 15 indicates that Algorithm 1 terminates. Indeed for each $J_n \in \mathcal{I}$, taking the *least index* k_n^* such that $\lambda^{k_n^*+1}(v) = \lambda^{k_n^*}(v)$ for every $v \in X^{\geq J_n}$ shows that the repeat loop is broken and the variable n decremented by 1. The value $n = 0$ is eventually reached and the algorithm stops with the global fixpoint λ^* . Notice that the first local fixpoint is reached with $k_N^* = 0$ as X^{J_N} is a bottom region.

Proposition 15 also shows that when a local fixpoint is reached in the arena $X^{\geq J_{n+1}}$ and the algorithm updates the labeling function λ^k in the arena $X^{\geq J_n}$, the values of $\lambda^k(v)$ do not change anymore for any $v \in V^{\geq J_{n+1}}$ but can still be modified for some $v \in V^{J_n}$. Recall also that outside of $X^{\geq J_n}$, the values of $\lambda^k(v)$ are still equal to the initial values $\lambda^0(v)$. These properties will be useful when we will prove that the constraint problem for reachability games is in PSPACE. They are summarized in the next lemma.

- **Lemma 16.** *Let k be a step of Algorithm 1, let J_n with $n \in \{1, \dots, N\}$. For all $v \in V^{J_n}$:*
- *if $k \leq k_{n+1}^*$, then $\lambda^{k+1}(v) = \lambda^k(v) = \lambda^0(v)$,*
 - *if $k_n^* \leq k$, then $\lambda^{k+1}(v) = \lambda^k(v) = \lambda^{k_n^*}(v)$,*
- Hence the values of $\lambda^k(v)$ and $\lambda^{k+1}(v)$ may be different only when $k_{n+1}^* < k < k_n^*$.*

The following theorem states how we characterize outcomes of SPEs in (\mathcal{X}, x_0) . It also provides a characterization of the outcomes of SPEs in (\mathcal{G}, v_0) by Lemma 7.

- **Theorem 17 (Characterization).** *Let (\mathcal{G}, v_0) be an initialized quantitative game and (\mathcal{X}, x_0) be its extended game. Let ρ be a play in $\text{Plays}_X(x_0)$. Then ρ is the outcome of an SPE in (\mathcal{X}, x_0) if and only if for all $v \in \text{Succ}^*(x_0)$, $\Lambda^*(v) \neq \emptyset$ and $\rho \in \Lambda^*(x_0)$.*

► **Example 18.** Let us come back to the running example of Figure 2. The different steps of Algorithm 1 are given in Table 1. The columns indicate the vertices according to their region, respectively Π , $\{2\}$, and \emptyset . Notice that for the region Π , we only write one column v as for all vertices (v, Π) the value of λ is equal to 0 all along the algorithm.

Recall that $J_1 = \emptyset < J_2 = \{2\} < J_3 = \Pi = \{1, 2\}$. The algorithm begins with the arena X^{J_3} . A fixpoint ($\lambda^1 = \lambda^0$) is immediately reached because all vertices belong to the target set of both players in X^{J_3} . Thus the first local fixpoint is reached with $k_3^* = 0$.

The algorithm then treats the arena $X^{\geq J_2}$. By Lemma 16, it is enough to consider the region X^{J_2} . Let us explain how to compute $\lambda^2(v)$ from $\lambda^1(v)$ on this region. For $v = (v_7, \{2\})$, we have that $\lambda^2(v) = 1 + \min_{(v, v') \in E^X} \sup \{\text{Cost}_1(\rho) \mid \rho \in \Lambda^1(v')\}$. As the unique successor of v is $(v_2, \{1, 2\})$, all λ^1 -consistent plays beginning in this successor have cost 0. So, we have that $\lambda^2(v) = 1$. For the computation of $\lambda^2(v_6, \{2\})$, the same argument holds since $(v_6, \{2\})$ has the unique successor $(v_7, \{2\})$. The vertex $(v_1, \{2\})$ has two successors: $(v_6, \{2\})$ and $(v_3, \{2\})$. Again, we know that all λ^1 -consistent plays beginning in $(v_6, \{2\})$ have cost 2. From $(v_3, \{2\})$ however, one can easily check that the play $(v_3, \{2\})(v_0, \{2\})(v_4, \{2\})^\omega$ is λ^1 -consistent and has cost $+\infty$ for player 1. Thus, we obtain that $\lambda^2(v_1, \{2\}) = 3$. For the other vertices of X^{J_2} , one can see that $\lambda^2(v) = \lambda^1(v)$.

Finally, we can check that the local fixed point is reached in the arena $X^{\geq J_2}$ (resp. $X^{\geq J_1}$) with $\lambda^3 = \lambda^2$ (resp. $\lambda^6 = \lambda^5 = \lambda^*$). Therefore the respective fixpoints are reached with $k_2^* = 2$ and $k_1^* = 5$. The labeling function indicated in Figure 2 is the one of λ^* . ◀

Region	Π	{2}							\emptyset					
		v	v_0	v_1	v_6	v_7	v_3	v_4	v_5	v_0	v_1	v_6	v_7	v_3
$\lambda^0 = \lambda^1$	0	0	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$
$\lambda^2 = \lambda^3$	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$	$+\infty$
λ^4	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	$+\infty$	3	2	1	$+\infty$	$+\infty$
$\lambda^5 = \lambda^*$	0	0	3	2	1	$+\infty$	$+\infty$	$+\infty$	4	3	2	1	$+\infty$	$+\infty$

■ **Table 1** The different steps of the algorithm computing λ^* for the extended game of Figure 2

4 Counter graph

In this section, given a labeling function λ , we introduce the concept of *counter graph* such that its infinite paths coincide with the plays that are λ -consistent. We then show that the counter graph associated with the fixpoint function λ^* computed by Algorithm 1 has an exponential size, an essential step to prove PSPACE membership of the constraint problem.

For the entire section, we fix a reachability game $\mathcal{G} = (G, (F_i)_{i \in \Pi}, (\text{Cost}_i)_{i \in \Pi})$ and $\mathcal{X} = (X, (F_i^X)_{i \in \Pi}, (\text{Cost}_i)_{i \in \Pi})$ its associated extended game.

A labeling function λ give constraints on costs of plays from each vertex in X , albeit only for the owner of this vertex. However, by the property of λ -consistence, constraints for a player carry over all the successive vertices, whether they belong to him or not. In order to check efficiently this property, we introduce the counter graph to keep track *explicitly* of the accumulation of constraints for all players at each step of a play. We first fix some notation.

- **Definition 19** (Maximal finite range). *Let $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ be a labeling function.*
- *We consider restrictions of λ to sub-arenas of V^X as follows. Let $n \in \{1, \dots, N\}$, we denote by $\lambda_n : V^{J_n} \rightarrow \mathbb{N} \cup \{+\infty\}$ the restriction of λ to V^{J_n} . Similarly we denote by $\lambda_{\geq n}$ (resp. $\lambda_{>n}$) the restriction of λ to $V^{\geq J_n}$ (resp. $V^{>J_n}$).*
 - *The maximal finite range of λ , denoted by $\text{mR}(\lambda)$, is equal to $\text{mR}(\lambda) = \max\{c \in \mathbb{N} \mid \lambda(v) = c \text{ for some } v \in V^X\}$ with the convention that $\text{mR}(\lambda) = 0$ if λ is the constant function $+\infty$. We extend this notion to restrictions of λ with the convention that $\text{mR}(\lambda_{>n}) = 0$ if J_n is a bottom region.*

Note that in the definition of maximal finite range, we only consider the *finite* values of λ .

- **Definition 20** (Counter Graph). *Let $\lambda : V^X \rightarrow \mathbb{N} \cup \{+\infty\}$ be a labeling function. Let $\mathcal{K} := \{0, \dots, K\} \cup \{+\infty\}$ with $K = \text{mR}(\lambda)$. The counter graph $\mathbb{C}(\lambda)$ for \mathcal{G} and λ is equal to $\mathbb{C}(\lambda) = (\Pi, V^C, (V_i^C)_{i \in \Pi}, E^C)$, such that:*

- $V^C = V^X \times \mathcal{K}^{|\Pi|}$
- $(v, (c_i)_{i \in \Pi}) \in V_j^C$ if and only if $v \in V_j^X$
- $((v, (c_i)_{i \in \Pi}), (v', (c'_i)_{i \in \Pi})) \in E^C$ if and only if:
 - $(v, v') \in E^X$, and

$$\text{■ for every } i \in \Pi, c'_i = \begin{cases} 0 & \text{if } i \in I(v') \\ c_i - 1 & \text{if } i \notin I(v'), v' \notin V_i^X \text{ and } c_i > 1 \\ \min(c_i - 1, \lambda(v')) & \text{if } i \notin I(v'), v' \in V_i^X \text{ and } c_i > 1. \end{cases}$$

Intuitively, the counter graph is constructed such that once a value $\lambda(v)$ is finite for a vertex $v \in V_i^X$ along a play in \mathcal{X} , the corresponding path in $\mathbb{C}(\lambda)$ keeps track of the induced constraint by (i) decrementing the counter value c_i for the concerned player i by 1 at every step, (ii) updating this counter if a stronger constraint for player i is encountered by visiting a vertex v' with a smaller value $\lambda(v')$, and (iii) setting the counter c_i to 0 if player i has indeed reached his target set.

Note that there may be vertices with *no outgoing edges*. Indeed, consider a vertex $(v, (c_i)_{i \in \Pi}) \in V^C$ such that $c_j = 1$ for some player j . By construction of $\mathbb{C}(\lambda)$, the only outgoing edges from $(v, (c_i)_{i \in \Pi})$ must link to vertices $(v', (c'_i)_{i \in \Pi})$ such that $(v, v') \in E^X$, $c'_j = 0$ and $j \in I(v')$. However, it may be that no successor v' of v in X is such that $j \in I(v')$.

Note as well that for each vertex $v \in V^X$, there exist many different vertices $(v, (c_i)_{i \in \Pi})$ in $\mathbb{C}(\lambda)$, one for each counter values profile. However, the intended goal of $\mathbb{C}(\lambda)$ is to monitor explicitly the constraints accumulated by each player along a play in \mathcal{X} regarding λ . Thus, we will only consider paths in $\mathbb{C}(\lambda)$ that start in vertices $(v, (c_i)_{i \in \Pi})$ such that the counter values correspond indeed to the constraint at the *beginning of a play* in \mathcal{X} regarding λ :

► **Definition 21** (Starting vertex in $\mathbb{C}(\lambda)$). *Let $v \in V^X$. We distinguish one vertex $v^C = (v, (c_i)_{i \in \Pi})$ in V^C , such that for every $i \in \Pi$, the counter value c_i is equal to 0 if $i \in I(v)$, to $\lambda(v)$ if $i \notin I(v)$ and $v \in V_i^X$, and to $+\infty$ otherwise. We call v^C the starting vertex associated with v , and denote by $SV(\lambda)$ the set of all starting vertices in $\mathbb{C}(\lambda)$.*

There exists a correspondence between λ -consistent plays in \mathcal{X} and *infinite* paths from starting vertices in $\mathbb{C}(\lambda)$, called *valid paths*, in the following way. On one hand, every play ρ in \mathcal{X} that is not λ -consistent does not appear in the counter graph: the first constraint regarding λ that is violated along ρ is reflected by a vertex in $\mathbb{C}(\lambda)$ with a counter value getting to 1 and no outgoing edges. On the other hand, λ -consistent plays in \mathcal{X} have a corresponding infinite path in the counter graph $\mathbb{C}(\lambda)$. This is formalized in the next lemma:

► **Lemma 22.** *There exists a λ -consistent play $\rho = \rho_0 \rho_1 \dots$ in $\text{Plays}_X(v)$ with $v \in V^X$ if, and only if there exists an associated infinite path $\pi = \pi_0 \pi_1 \dots$ in $\mathbb{C}(\lambda)$ such that $\pi_0 = v^C \in SV(\lambda)$ and ρ is the projection of π on V^X (π_n is of the form $(\rho_n, (c'_i)_{i \in \Pi})$ for all $n \in \mathbb{N}$).*

Since the edge relation E^C in $\mathbb{C}(\lambda)$ respects the edge relation E^X in \mathcal{X} , the region decomposition of a path in \mathcal{X} given in Lemma 8 can also be applied to a path in $\mathbb{C}(\lambda)$.

In order to prove the PSPACE membership for the constraint problem, we need to show that the counter graph $\mathbb{C}(\lambda^*)$, with λ^* the fixpoint function computed by Algorithm 1, has an exponential size. To this end it is enough to show an exponential upper bound on the maximal finite range $\text{mR}(\lambda^*) (= \text{mR}(\lambda_{\geq 1}^{k_1^*}))$ of λ^* . To do so, we prove with Theorem 23, by induction on the number of computation steps k of Algorithm 1, an exponential upper bound on $\text{mR}(\lambda_\ell^k)$ and $\text{mR}(\lambda_{\geq \ell}^k)$ for every region X^{J_ℓ} and every step k .

► **Theorem 23.** *For every $k \in \mathbb{N}$ and region X^{J_ℓ} , we have*

$$\text{mR}(\lambda_\ell^k) \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi|+2)}) \quad \text{and} \quad \text{mR}(\lambda_{\geq \ell}^k) \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi|+2)}).$$

The next corollary will be useful in Section 5 where not only the maximal finite range must be exponentially bounded, but also the cost of any play satisfying the current constraints.

► **Corollary 24.** *Let $v \in V^X$ with $I(v) = J_\ell$ with $\ell < N$. Let $k \in \mathbb{N}$ such that $k > k_{\ell+1}^*$. Suppose there exists $c \in \mathbb{N}$ such that $\sup \{\text{Cost}_i(\rho) \mid \rho \in \Lambda^k(v)\} = c$. Then, we have:*


$$c \leq \mathcal{O}(|V|^{(|V|+3) \cdot (|\Pi|+2)}).$$

Let us give a few ingredients of the induction of Theorem 23. For each region X^{J_ℓ} , the value $\text{mR}(\lambda_\ell^0)$ is always equal to 0. For the general case, the value $\text{mR}(\lambda_\ell^{k+1})$ depends on the values of λ_ℓ^{k+1} , which, when finite, are in turn determined by the maximal cost of the λ^k -consistent plays starting in region X^{J_ℓ} (see Definition 13). A crucial point to obtain these bounds is that this maximal cost corresponds to the length of the longest cycle-free prefix of an infinite path starting in region X^{J_ℓ} in $\mathbb{C}(\lambda^k)$. Furthermore, we can evaluate this maximal length in terms of the values of $\text{mR}(\lambda_\ell^k)$ and $\text{mR}(\lambda_{\geq \ell}^k)$, which are bounded by the previous step of the induction. The following key technical lemma formalizes this idea.

► **Lemma 25.** *Let v^C be a starting vertex in $SV(\lambda)$ associated with $v \in V^X$ such that $I(v) = J_\ell$. Let π be a finite prefix of a valid path in $\mathbb{C}(\lambda)$ such that $\pi_0 = v^C$ and π does not contain any cycle. Then*

$$|\pi| \leq |V| + 2 \cdot \text{mR}(\lambda_\ell) + \sum_{r=|J_\ell|+1}^{|\Pi|} |V| + 2 \cdot \max_{\substack{J_j > J_\ell \\ |J_j|=r}} \text{mR}(\lambda_j).$$

Proof sketch. Let π be a finite prefix of a valid path in $\mathbb{C}(\lambda)$ as in the statement. Let $\pi[\ell] \dots \pi[m]$ be its region decomposition according to Lemma 8, graphically represented in Figure 3. Let ρ be the corresponding path in \mathcal{X} and $\rho[\ell] \dots \rho[m]$ be its region decomposition. Let us consider a fixed non-empty section $\pi[n]$.



■ **Figure 3** Region decomposition of π

Suppose first that the counter values at $\pi[n]_0$ are either 0 or $+\infty$. Let us prove that along $\pi[n]$, there can be at most $|V|$ steps before reaching a vertex with a finite positive value of λ :

- assume there is a cycle in the corresponding section $\rho[n]$ in \mathcal{X} such that from $\rho[n]_0$ and along the cycle, all the values of λ are either 0 or $+\infty$,
- by construction of $\mathbb{C}(\lambda)$, the counter values in the corresponding prefix of $\pi[n]$ remain fixed for each vertex of this prefix: as no value of λ is positive and finite, no counter value can be decremented,
- thus, the cycle in $\rho[n]$ is also a cycle in $\pi[n]$ which is impossible by hypothesis,
- thus there is no such cycle in $\rho[n]$, and as there are at most $|V|$ vertices in region X^{J_n} , $\rho[n]$ can have a prefix of length at most $|V|$ with only values 0 or $+\infty$ for λ , implying that this is also the case for $\pi[n]$.

Therefore, we can decompose $\pi[n]$ into a (possibly empty) prefix of length at most $|V|$, and a (possibly empty) suffix where at least one counter value c'_i , for some i , is a positive finite value in its first vertex v' . This frontier between prefix and suffix of $\pi[n]$ is represented by a vertical double bar \parallel with caption $\infty \rightarrow c$ in Figure 3. This value c'_i is bounded by $\text{mR}(\lambda_n)$, the maximal finite range of λ_n . From there, as the corresponding ρ is λ -consistent, player i reaches his target set in at most c'_i steps, and ρ enters a new region, which means that the section $\pi[n]$ is over. So, in that case, the length of $\pi[n]$ can be bounded by $|V| + \text{mR}(\lambda_n)$.

Suppose now that at vertex $\pi[n]_0$, there exists a counter value c_i for some player i that is neither 0 nor $+\infty$. This means that there was a constraint for player i initialized in a previous section $\pi[n']$, with $n' < n$, that has carried over to $\pi[n]_0$, via decrements of at least 1 per step. We know that the initial finite counter value is bounded by $\text{mR}(\lambda_{n'})$, and appeared before the end of section $\pi[n']$. Thus the length from the end of section $\pi[n']$ to the end of section $\pi[n]$ is bounded by $\text{mR}(\lambda_{n'})$, as again, once the counter value attains 0 for player i , the path π has entered the next section.

Therefore, considering the possible cases for each section, we can bound the total length of π as follows: $|\pi| \leq \sum_{j=\ell}^m |V| + 2 \cdot \text{mR}(\lambda_j)$.

Finally, remark that by I -monotonicity, it is actually the case that only (and at most) $|\Pi|$ different non-empty sections can appear in the decomposition of π . Furthermore, for each

$n \in \{\ell + 1, \dots, N\}$, we have $\text{mR}(\lambda_n) \leq \max\{\text{mR}(\lambda_j) \mid J_j > J_\ell, |J_j| = |J_n|\}$ by Definition 19. Thus, we obtain the bound stated in Lemma 25. ◀

5 PSPACE completeness

In this section, we prove that the constraint problem is PSPACE-complete (Theorem 4). Given a reachability game (\mathcal{G}, v_0) and two thresholds $x, y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$, this problem is to decide whether there exists an SPE σ in this game such that for all $i \in \Pi$, $x_i \leq \text{Cost}_i(\langle \sigma \rangle_{v_0}) \leq y_i$. As the proof of the PSPACE-hardness (based on a reduction from the QBF problem) is nearly the same as for qualitative reachability games [10], we only prove the PSPACE-easiness.

► **Proposition 26.** *The constraint problem for quantitative reachability games is in PSPACE.*

Let us provide a high level sketch of the proof of our PSPACE procedure. Thanks to Theorem 17, solving the constraint problem reduces in finding a λ^* -consistent play ρ in (\mathcal{X}, x_0) satisfying the constraints. By Lemma 22, the latter problem reduces in finding a corresponding valid path π in the counter graph $\mathbb{C}(\lambda^*)$, satisfying the constraints. Roughly speaking, in order to solve our initial problem, it suffices to decide the existence of a valid path that is a lasso and satisfies the constraints in the counter graph, which is exponential in the size of our original input. Classical arguments, using Savitch's Theorem can thus be used to prove the PSPACE membership. Nevertheless, the detailed proof is more intricate for two reasons. The first reason is that the counter graph is based on the labeling function λ^* . We thus also have to prove that this function λ^* can be computed in PSPACE. The second reason is that, a priori, although we know that the counter graph is of exponential size, we do not know explicitly its size. This is problematic when using classical NPSpace algorithms that guess some path in a graph of exponential size, where a counter bounded by the size of the graph is needed to guarantee the termination of the procedure. In order to overcome this, we also need a PSPACE procedure to obtain the actual size of the counter graph. Recall that, as $|\mathbb{C}(\lambda^*)| = |V| \cdot 2^{|\Pi|} \cdot (K^* + 1)^{|\Pi|}$ with $K^* = \text{mR}(\lambda^*)$, we only have to determine the value of K^* . This is possible thanks to Proposition 27.

► **Proposition 27.** *Given a quantitative game (\mathcal{G}, v_0) , for all $k \in \mathbb{N}$, for all J_ℓ with $\ell \in \{1, \dots, N\}$, the set $\{\lambda^k(v) \mid v \in V^{J_\ell}\}$ and the value $\text{mR}(\lambda_{\geq \ell}^k)$ can be computed in PSPACE.*

Proof sketch. The whole procedure works by induction on k , the steps in the computation of the labeling function λ^* . Moreover, it exploits the structural evolution of the local fixpoints formalized in Proposition 15: once a step is fixed, we proceed region by region, beginning with the bottom region J_N and then proceeding bottom-up by following the total order on \mathcal{I} .

Let X^{J_ℓ} be a region, we aim at proving that $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ and $\text{mR}(\lambda_{\geq \ell}^{k+1})$ are computable in PSPACE, assuming that (i) we know $\{\lambda^k(v) \mid v \in V^{J_\ell}\}$, and (ii) we can compute $\{\lambda^k(v) \mid v \in V^{J'}\}$ for all $J' > J_\ell$ and $\text{mR}(\lambda_{\geq \ell}^k)$ in PSPACE.

We first focus on the computation of λ^{k+1} from λ^k . Let k_ℓ^* (resp. $k_{\ell+1}^*$) be the step where the fixpoint is reached for region X^{J_ℓ} (resp. $X^{J_{\ell+1}}$). Recall that $k_{\ell+1}^* < k_\ell^*$. When $k \leq k_{\ell+1}^*$ (resp. $k > k_\ell^*$), we have that $\lambda^{k+1}(v) = \lambda^k(v)$, for each $v \in V^{J_\ell}$, by Lemma 16. The tricky case is when $k_{\ell+1}^* < k \leq k_\ell^*$, let us focus on it.

In this case, the computation of $\lambda^{k+1}(v)$ from $\lambda^k(v)$, for all $v \in V^{J_\ell}$, relies on the maximum of the cost of the λ^k -consistent plays (see Definition 13). Computing this maximum

³ Notice the little difference with the inequalities given in Lemma 16. When $k = k_\ell^*$, we still need to compute λ^{k+1} to realize that the fixpoint is effectively reached.

is equivalent to guess the existence of lassoes with a certain cost in the counter graph $\mathbb{C}(\lambda_{\geq \ell}^k)$. The size of these lassoes depends on the size of $\mathbb{C}(\lambda_{\geq \ell}^k)$ equal to $|V| \cdot 2^{|\Pi|} \cdot (K+1)^{|\Pi|}$ where $K = \text{mR}(\lambda_{\geq \ell}^k)$. Moreover, each lasso is guessed region by region. As we can compute in PSPACE the value K and the sets $\{\lambda^k(v) \mid v \in V^{J'}\}$ for all $J' \geq J_\ell$ by induction hypothesis, we can also compute in PSPACE the maximum of the cost of the λ^k -consistent plays and thus the set $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$. The correctness of this approach relies on Proposition 15 which ensures that the values of λ^* have already reached a fixpoint for all the $v \in V^{J'}$ such that $J' > J_\ell$. Once we have computed $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$ in PSPACE, it is clear that $\text{mR}(\lambda_{\geq \ell}^{k+1})$ can also be computed in PSPACE because $\text{mR}(\lambda_{\geq \ell}^{k+1}) = \max\{\text{mR}(\lambda_\ell^{k+1}), \text{mR}(\lambda_{\geq \ell+1}^{k+1})\}$ (notice that $\text{mR}(\lambda_{\geq \ell+1}^{k+1}) = \text{mR}(\lambda_{\geq \ell+1}^k)$ because $k > k_{\ell+1}^*$ and thus $X^{J_{\ell+1}}$ has already reached its local fixpoint and so the induction hypothesis can be used).

Notice that all these arguments hold because: (i) the value of $\text{mR}(\lambda_{\geq \ell}^k)$ and any value of $\lambda^k(v)$ with $v \in V^{\geq \ell}$ (or of $\lambda^{k+1}(v)$ with $v \in V^\ell$) can be encoded in polynomial size memory (by Theorem 23); (ii) the value of the maximal cost of λ^k -consistent plays is at most exponential in the inputs (by Corollary 24); (iii) any set $\{\lambda^k(v) \mid v \in V^{J'}\}$, $J' \geq J_\ell$ (or the set $\{\lambda^{k+1}(v) \mid v \in V^{J_\ell}\}$) is composed of at most $|V|$ values which can be encoded with polynomial size memory by (i), and (iv) we have a polynomial number of such values to keep in memory. ◀

Proof of Proposition 26. Let (\mathcal{G}, v_0) be an initialized reachability game and let $x, y \in (\mathbb{N} \cup \{+\infty\})^{|\Pi|}$ be two thresholds. Let (\mathcal{X}, x_0) be its extended game with $x_0 = (v_0, I_0)$. Let $\mathbb{C}(\lambda_{\geq I_0}^*)$ be the counter graph restricted to the arena $X^{\geq I_0}$ and so $|\mathbb{C}(\lambda_{\geq I_0}^*)| = |V| \cdot 2^{|\Pi|} \cdot (K^0 + 1)^{|\Pi|}$ with $K^0 = \text{mR}(\lambda_{\geq I_0}^*)$. Let us recall that K^0 can be computed (resp. encoded) in PSPACE thanks to Proposition 27 (resp. Theorem 23).

We have to guess a lasso $\pi = hg^\omega$ in the counter graph $\mathbb{C}(\lambda_{\geq I_0}^*)$ such that its corresponding play ρ^X in the extended game satisfies the constraints, *i.e.*, $x_i \leq \text{Cost}_i(\rho^X) \leq y_i$ for all $i \in \Pi$. The length L of π is at most $\max(\max\{y_i \mid y_i < +\infty\}, |\mathbb{C}(\lambda_{\geq I_0}^*)|) + 2 \cdot |\mathbb{C}(\lambda_{\geq I_0}^*)|$ as the cycle of π can appear after the constraints $y_i < +\infty$ are satisfied.

In order to have a PSPACE procedure, we cannot guess π entirely and we have to proceed region by region. If $I_0 = J_\ell$ for some $\ell \in \{1, \dots, N\}$, we consider the region decomposition $\pi[\ell]\pi[\ell+1] \dots \pi[t]$ of π , with $t \in \{\ell, \dots, N\}$, where some sections $\pi[m]$ may be empty.

We first guess the first vertex of the cycle g and then we guess successively $\pi[\ell]\pi[\ell+1]$ and so on. To guess $\pi[m]$ with $\ell \in \{m, \dots, t\}$, assuming it is not empty, we guess one by one its vertices. To guess a vertex $(v', J_m, (c'_i)_{i \in \Pi})$ we only have to keep its predecessor in memory and to know the value $\lambda^*(v', J_m)$. So, we need to know $\{\lambda^*(v) \mid v \in V^{J_m}\}$. Thanks to Proposition 27, we can obtain this set in polynomial size memory and once we move to another region, we can forget this set and compute the new one.

Additionally, we have a counter C_L to count the length of π and for each player $i \in \Pi$ we have a counter C_i keeping track the current cost of player i along π . For all $i \in \Pi$, we have that $C_i, C_L \leq L$. Henceforth, all these counters can be encoded with polynomial size memory. Moreover, in addition to these counters, we currently maintain: (i) the first vertex of g , (ii) the current vertex we are guessing, (iii) its predecessor and, (iv) the set $\{\lambda^*(v) \mid v \in V^{J_m}\}$ (only $|V|$ values which can be encoded in polynomial size memory by Theorem 23) if we are guessing $\pi[m]$. As for (ii) and (iii) a vertex in the counter graph is composed of a vertex of V , a subset I of Π and $|\Pi|$ counter values which are at most exponential in the input (Theorem 23), all this procedure can be done in PSPACE. ◀

References

- 1 Nicolas Basset, Ismaël Jecker, Arno Pauly, Jean-François Raskin, and Marie van den Bogaard. Beyond admissibility: Dominance between chains of strategies. In *CSL*, volume 119 of *LIPICs*, pages 10:1–10:22. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
- 2 Dietmar Berwanger. Admissibility in infinite games. In *STACS*, volume 4393 of *LNCS*, pages 188–199. Springer, 2007.
- 3 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In *LATA*, volume 9618 of *LNCS*, pages 3–23. Springer, 2016.
- 4 Romain Brenguier, Arno Pauly, Jean-François Raskin, and Ocan Sankur. Admissibility in Games with Imperfect Information (Invited Talk). In *CONCUR*, volume 85 of *LIPICs*, pages 2:1–2:23. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- 5 Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin, and Ocan Sankur. Admissibility in Quantitative Graph Games. In *FSTTCS*, volume 65 of *LIPICs*, pages 42:1–42:14, 2016.
- 6 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. In *CONCUR*, *LIPICs* 42, pages 100–113. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015.
- 7 Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS*, pages 23:1–23:10. ACM, 2014.
- 8 Thomas Brihaye, Véronique Bruyère, Julie De Pril, and Hugo Gimbert. On subgame perfection in quantitative reachability games. *Logical Methods in Computer Science*, 9(1), 2012.
- 9 Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with ω -regular boolean objectives. In *GandALF*, pages 16–29, 2018.
- 10 Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with omega-regular boolean objectives. *CoRR*, abs/1806.05544, 2018.
- 11 Thomas Brihaye, Véronique Bruyère, Noémie Meunier, and Jean-François Raskin. Weak subgame perfect equilibria and their application to quantitative reachability. In *CSL*, volume 41 of *LIPICs*, pages 504–518. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.
- 12 Thomas Brihaye, Julie De Pril, and Sven Schewe. Multiplayer cost games with simple nash equilibria. In *LFCS*, volume 7734 of *Lecture Notes in Computer Science*, pages 59–73. Springer, 2013.
- 13 Véronique Bruyère. Computer aided synthesis: A game-theoretic approach. In *DLT*, volume 10396 of *LNCS*, pages 3–35. Springer, 2017.
- 14 Véronique Bruyère, Stéphane Le Roux, Arno Pauly, and Jean-François Raskin. On the existence of weak subgame perfect equilibria. In *FOSSACS*, volume 10203 of *LNCS*, pages 145–161, 2017.
- 15 K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with secure equilibria. *Theoretical Computer Science*, 365:67–82, 2006.
- 16 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- 17 Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. The complexity of rational synthesis. In *ICALP*, volume 55 of *LIPICs*, pages 121:1–121:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
- 18 Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. Rational synthesis under imperfect information. In *LICS*, pages 422–431. ACM, 2018.
- 19 Dana Fisman, Orna Kupferman, and Yoad Lustig. Rational synthesis. In *TACAS*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.
- 20 János Fleisch, Jeroen Kuipers, Ayala Mashiah-Yaakovi, Gijs Schoenmakers, Eilon Solan, and Koos Vrieze. Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operation Research*, 35:742–755, 2010.

- 21 Erich Grädel and Michael Ummels. Solution Concepts and Algorithms for Infinite Multiplayer Games. In *New Perspectives on Games and Interaction*, volume 4, pages 151–178. Amsterdam University Press, 2008.
- 22 Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. Synthesis with rational environments. In *EUMAS*, LNCS 8953, pages 219–235. Springer, 2014.
- 23 Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. Reasoning About Strategies. In *FSTTCS 2010*, volume 8 of *LIPICs*, pages 133–144. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2010.
- 24 J. F. Nash. Equilibrium points in n -person games. In *PNAS*, volume 36, pages 48–49. National Academy of Sciences, 1950.
- 25 Martin J. Osborne. *An introduction to game theory*. Oxford Univ. Press, 2004.
- 26 A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
- 27 Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In *FSTTCS*, volume 4337 of *LNCS*, pages 212–223. Springer, 2006.
- 28 Michael Ummels. The complexity of nash equilibria in infinite multiplayer games. In *FOSSACS*, volume 4962 of *LNCS*, pages 20–34. Springer, 2008.
- 29 Michael Ummels and Dominik Wojtczak. The complexity of Nash equilibria in limit-average games. In *CONCUR*, volume 6901 of *LNCS*, pages 482–496. Springer, 2011.