

Subgame Perfection for Equilibria in Quantitative Reachability Games

Thomas Brihaye¹, Véronique Bruyère¹, Julie De Pril¹, and Hugo Gimbert²

¹ University of Mons - UMONS

Place du Parc 20, 7000 Mons, Belgium

{thomas.brihaye, veronique.bruyere, julie.depril}@umons.ac.be

² LaBRI & CNRS, Bordeaux, France

hugo.gimbert@labri.fr

Abstract. We study turn-based quantitative multiplayer non zero-sum games played on finite graphs with reachability objectives. In such games, each player aims at reaching his own goal set of states as soon as possible. A previous work on this model showed that Nash equilibria (resp. secure equilibria) are guaranteed to exist in the multiplayer (resp. two-player) case. The existence of secure equilibria in the multiplayer case remained, and is still an open problem. In this paper, we focus our study on the concept of subgame perfect equilibrium, a refinement of Nash equilibrium well-suited in the framework of games played on graphs. We also introduce the new concept of subgame perfect secure equilibrium. We prove the existence of subgame perfect equilibria (resp. subgame perfect secure equilibria) in multiplayer (resp. two-player) quantitative reachability games. Moreover, we provide an algorithm deciding the existence of secure equilibria in the multiplayer case.

1 Introduction

General framework. The construction of correct and efficient computer systems (hardware or software) is recognized as an extremely difficult task. To support the design and verification of such systems, mathematical logic, automata theory [16] and more recently model-checking [13] have been intensively studied. The efficiency of the model-checking approach is widely recognized when applied to systems that can be accurately modeled as a finite-state automaton. In contrast, the application of these techniques to more complex systems like embedded systems or distributed systems has been less successful. This could be partly explained by the following reasons: classical automata-based models do not faithfully capture the complex behavior of modern computational systems that are usually composed of several interacting components, also interacting with an environment that is only partially under control. One recent trend to improve the automata models used in the classical approach of verification is to generalize these models with the more flexible and mathematically deeper game-theoretic framework [22,23].

The first steps to extend computational models with concepts from game theory were done with the so-called two-player zero-sum games played on graphs [14].

Those games are adequate to model controller-environment interactions problems [26,27]. Moves of player 1 model actions of the controller whereas moves of player 2 model the uncontrollable actions of the environment, and a winning strategy for player 1 is an abstract form of a control program that enforces the control objective. However, only purely antagonist interactions between a controller and a hostile environment can be modeled in this framework. In order to study more complex systems with more than two components and objectives that are not necessarily antagonist, we need multiplayer non zero-sum games. Moreover, we do not look for winning strategies, but rather try to find relevant notions of equilibria, like the famous notion of *Nash equilibrium* [22]. We also consider the more recent concept of *secure equilibrium* [9] which is especially well-suited for assume-guarantee synthesis [11,12]. On the other hand, only qualitative objectives have been considered so far to specify, for example, that a player must be able to reach a target set of states in the underlying game graph. But, in line with the previous point, we also want to express and solve games for quantitative objectives where each player wants to force the play to reach a particular set of states within a given time bound, or within a given energy consumption limit. In summary, we need to study *equilibria* for *multiplayer non zero-sum* games played on graphs with *quantitative* objectives. This article provides some new results in this research direction, in particular it is another step in the quest for solution concepts well-suited for the computer-aided synthesis and verification of multi-agent systems.

Our contribution. We study turn-based multiplayer non zero-sum games played on finite graphs with quantitative reachability objectives, continuing work initiated in [6]. In this framework each player aims at reaching his own goal as soon as possible. In [6], among other results, it has been proved that a finite-memory Nash (resp. secure) equilibria always exists in multiplayer (resp. 2-player) games.

In this paper we consider alternative solution concepts to the classical notion of Nash equilibria. In particular, in the present framework of games on graphs, it is very natural to consider the notion of *subgame perfect equilibrium* [25]: a choice of strategies is not only required to be optimal for the initial vertex but also for every possible initial history of the game. Indeed if the initial state or the initial history of the system is not known, then a robust controller should be subgame perfect. We introduce a new and even stronger solution concept with the notion of subgame perfect *secure equilibrium*, which gathers both the sequential nature of subgame perfect equilibria and the verification-oriented aspects of secure equilibria. These different notions of equilibria are precisely defined in Section 2.

In this paper, we address the following problems:

Problem 1. *Given a multiplayer quantitative reachability game \mathcal{G} , does there exist a Nash (resp. secure, subgame perfect, subgame perfect secure) equilibrium in \mathcal{G} ?*

Problem 2. *Given a Nash (resp. secure, subgame perfect, subgame perfect secure) equilibrium in a multiplayer quantitative reachability game \mathcal{G} , does there exist such an equilibrium with finite memory?*

These questions have been positively solved by some of the authors in [6] for Nash equilibria in multiplayer games, and for secure equilibria in two-player games. Notice that these problems and related ones have been investigated a lot in the qualitative framework (see [15]).

Here we go a step further and establish the following results about subgame perfect and secure equilibria:

- in every multiplayer quantitative reachability game, there exists a subgame perfect equilibrium (Theorem 10),
- in every two-player quantitative reachability game, there exists a subgame perfect secure equilibrium (Theorem 13),
- in every multiplayer quantitative reachability game, one can decide whether there exists a secure equilibrium in ExpSpace (Theorem 14),
- if there exists a secure equilibrium in a multiplayer quantitative reachability game, then there exists one that is finite-memory (Theorem 15).

Related work. Several recent papers have considered *two-player zero-sum* games played on finite graphs with regular objectives enriched by some *quantitative* aspects. Let us mention some of them: games with finitary objectives [10], games with prioritized requirements [1], request-response games where the waiting times between the requests and the responses are minimized [17,28], and games whose winning conditions are expressed via quantitative languages [2].

Other works concern *qualitative non zero-sum* games. In [9] where the notion of secure equilibrium has been introduced, it is proved that a unique maximal payoff profile of secure equilibria always exists for two-player non zero-sum games with regular objectives. In [15], general criteria ensuring existence of Nash equilibria and subgame perfect equilibria (resp. secure equilibria) are provided for multiplayer (resp. 2-player) games, as well as complexity results. In [4], the existence of Nash equilibria is studied for timed games with qualitative reachability objectives. Complexity issues are discussed in [5] about Nash equilibria in multiplayer concurrent games with Büchi objectives.

Finally, let us mention works that combine both *quantitative* and *non zero-sum* aspects. In [3], the authors study games played on graphs with terminal vertices where quantitative payoffs are assigned to the players. These games may have cycles but all the infinite plays form a single outcome (like in chess where every infinite play is a draw). That paper gives criteria that ensure the existence of Nash (and subgame perfect) equilibria in pure and memoryless strategies. In [19], the studied games are played on priced graphs similar to the ones considered in this article, however in a concurrent way. In this concurrent framework, Nash equilibria are not guaranteed to exist anymore. The authors provide an algorithm to decide existence of Nash equilibria, thanks to a Büchi automaton accepting all Nash equilibria outcomes. The complexity of some related decision problems is also studied. In [24], the authors study Muller games on finite graphs where players have a preference ordering on the sets of the Muller table. They show that Nash equilibria always exist for such games, and that it is decidable whether there exists a subgame perfect equilibrium. In both cases they give a procedure to compute an equilibrium strategy profile (when it exists).

2 Preliminaries

2.1 Games, Strategy Profiles and Equilibria

We consider here *quantitative* games played on a graph where all the players have *reachability objectives*. It means that, given a certain set of vertices Goal_i , each player i wants to reach one of these vertices as soon as possible. We recall the basic notions about these games and we introduce different kinds of equilibria, like Nash equilibria. This section is inspired from [6].

Definition 1. An infinite turn-based multiplayer quantitative reachability game is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, v_0, E, (\text{Goal}_i)_{i \in \Pi})$ where

- Π is a finite set of players,
- $G = (V, (V_i)_{i \in \Pi}, v_0, E)$ is a finite directed graph where V is the set of vertices, $(V_i)_{i \in \Pi}$ is a partition of V into the state sets of each player, $v_0 \in V$ is the initial vertex, and $E \subseteq V \times V$ is the set of edges, and
- $\text{Goal}_i \subseteq V$ is the non-empty goal set of player i .

From now on we often use the term *game* to denote a multiplayer quantitative reachability game according to Definition 1.

We assume that each vertex has at least one outgoing edge. The game is played as follows. A token is first placed on the vertex v_0 . Player i , such that $v_0 \in V_i$, has to choose one of the outgoing edges of v_0 and put the token on the vertex v_1 reached when following this edge. Then, it is the turn of the player who owns v_1 . And so on.

A *play* $\rho \in V^\omega$ (resp. a *history* $h \in V^+$) of \mathcal{G} is an *infinite* (resp. a *finite*) path through the graph G starting from vertex v_0 . Note that a history is always non-empty because it starts with v_0 . The set $H \subseteq V^+$ is made up of all the histories of \mathcal{G} , and for $i \in \Pi$, the set H_i is the set of all histories $h \in H$ whose last vertex belongs to V_i .

For any play $\rho = \rho_0\rho_1 \dots$ of \mathcal{G} , we define $\text{Cost}_i(\rho)$ the *cost* of player i as:

$$\text{Cost}_i(\rho) = \begin{cases} l & \text{if } l \text{ is the least index such that } \rho_l \in \text{Goal}_i, \\ +\infty & \text{otherwise.} \end{cases}$$

We note $\text{Cost}(\rho) = (\text{Cost}_i(\rho))_{i \in \Pi}$ the *cost profile* for the play ρ . Each player i aims to *minimize* the cost he has to pay, i.e. reach his goal set as soon as possible. The cost profile for a history h is defined similarly.

A *prefix* (resp. *proper prefix*) α of a history $h = h_0 \dots h_k$ is a finite sequence $h_0 \dots h_l$, with $l \leq k$ (resp. $l < k$), denoted by $\alpha \leq h$ (resp. $\alpha < h$). We similarly consider a prefix α of a play ρ , denoted by $\alpha < \rho$. The function **Last** returns, given a history $h = h_0 \dots h_k$, the last vertex h_k of h , and the *length* $|h|$ of h is the number k of its *edges*¹. Given a play $\rho = \rho_0\rho_1 \dots$, we denote by $\rho_{\leq l}$ the prefix of ρ of length l , i.e. $\rho_{\leq l} = \rho_0\rho_1 \dots \rho_l$. Similarly, $\rho_{< l} = \rho_0\rho_1 \dots \rho_{l-1}$.

¹ Note that the length is not defined as the number of vertices.

We say that a play $\rho = \rho_0\rho_1\dots$ *visits* a set $S \subseteq V$ (resp. a vertex $v \in V$) if there exists $l \in \mathbb{N}$ such that ρ_l is in S (resp. $\rho_l = v$). The same terminology also stands for a history h .

A *strategy* of player i in \mathcal{G} is a function $\sigma : H_i \rightarrow V$ assigning to each history $h \in H_i$, a next vertex $\sigma(h)$ such that $(\text{Last}(h), \sigma(h))$ belongs to E . We say that a play $\rho = \rho_0\rho_1\dots$ of \mathcal{G} is *consistent* with a strategy σ of player i if $\rho_{k+1} = \sigma(\rho_0\dots\rho_k)$ for all $k \in \mathbb{N}$ such that $\rho_k \in V_i$. The same terminology is used for a history h of \mathcal{G} . A *strategy profile* of \mathcal{G} is a tuple $(\sigma_i)_{i \in \Pi}$ where σ_i is a strategy for player i . It determines a unique play of \mathcal{G} consistent with each strategy σ_i , called the *outcome* of $(\sigma_i)_{i \in \Pi}$ and denoted by $\langle (\sigma_i)_{i \in \Pi} \rangle$. We write σ_{-j} for $(\sigma_i)_{i \in \Pi \setminus \{j\}}$, the set of strategies σ_i for all the players except for player j .

A strategy σ of player i is *memoryless* if σ depends only on the current vertex, i.e. $\sigma(hv) = \sigma(v)$ for all $h \in H$ and $v \in V_i$. More generally, σ is a *finite-memory strategy* if the equivalence relation \approx_σ on H defined by $h \approx_\sigma h'$ if $\sigma(h\delta) = \sigma(h'\delta)$ for all $\delta \in H_i$ has finite index. In other words, a finite-memory strategy is a strategy that can be implemented by a finite automaton with output. A strategy profile $(\sigma_i)_{i \in \Pi}$ is called *memoryless* or *finite-memory* if each σ_i is a memoryless or a finite-memory strategy, respectively.

For a strategy profile $(\sigma_i)_{i \in \Pi}$ with outcome ρ and a strategy σ'_j of player j , we say that *player j deviates from ρ* if there exists a prefix h of ρ , consistent with σ'_j , such that $h \in H_j$ and $\sigma'_j(h) \neq \sigma_j(h)$.

We now introduce different notions of equilibria in the *quantitative* framework and give several examples to make clear the presented concepts. We first begin with the definition of *Nash equilibrium*.

Definition 2. A strategy profile $(\sigma_i)_{i \in \Pi}$ of a game \mathcal{G} is a Nash equilibrium if for all player $j \in \Pi$ and for all strategy σ'_j of player j , we have:

$$\text{Cost}_j(\rho) \leq \text{Cost}_j(\rho')$$

where $\rho = \langle (\sigma_i)_{i \in \Pi} \rangle$ and $\rho' = \langle \sigma'_j, \sigma_{-j} \rangle$.

This definition means that for all $j \in \Pi$, player j has no incentive to deviate since he can not strictly decrease his cost when using σ'_j instead of σ_j . Keeping notations of Definition 2 in mind, a strategy σ'_j such that $\text{Cost}_j(\rho) > \text{Cost}_j(\rho')$ is called a *profitable deviation* for player j w.r.t. $(\sigma_i)_{i \in \Pi}$. In this case, either player j pays an infinite cost for ρ and a finite cost for ρ' (i.e. ρ' visits Goal_j , but ρ does not), or player j pays a finite cost for ρ and a strictly lower cost for ρ' (i.e. ρ' visits Goal_j for the first time earlier than ρ does).

We now define the concept of *secure equilibrium*². We first need to associate a binary relation \prec_j on cost profiles with each player j . Given two cost profiles $(x_i)_{i \in \Pi}$ and $(y_i)_{i \in \Pi}$:

$$(x_i)_{i \in \Pi} \prec_j (y_i)_{i \in \Pi} \quad \text{iff} \quad (x_j > y_j) \vee (x_j = y_j \wedge (\forall i \in \Pi \ x_i \leq y_i) \wedge (\exists i \in \Pi \ x_i < y_i)).$$

² Our definition naturally extends the notion of *secure equilibrium* proposed in [9] to the quantitative framework.

We then say that *player j prefers $(y_i)_{i \in \Pi}$ to $(x_i)_{i \in \Pi}$* . In other words, player j prefers a cost profile to another one either if he has a strictly lower cost, or if he keeps the same cost, the other players have a greater cost, and at least one has a strictly greater cost.

Definition 3. A strategy profile $(\sigma_i)_{i \in \Pi}$ of a game \mathcal{G} is a secure equilibrium if for all player $j \in \Pi$, there does not exist any strategy σ'_j of player j such that:

$$\text{Cost}(\rho) \prec_j \text{Cost}(\rho')$$

where $\rho = \langle (\sigma_i)_{i \in \Pi} \rangle$ and $\rho' = \langle \sigma'_j, \sigma_{-j} \rangle$.

In other words, player j has no incentive to deviate w.r.t. the relation \prec_j . A strategy σ'_j such that $\text{Cost}(\rho) \prec_j \text{Cost}(\rho')$ is called a \prec_j -profitable deviation for player j w.r.t. $(\sigma_i)_{i \in \Pi}$. Clearly, any secure equilibrium is a Nash equilibrium.

We now introduce a third type of equilibrium: the *subgame perfect equilibrium*. In this case, a strategy profile is not only required to be optimal for the initial vertex, but also after every possible history of the game. Before giving the definition, we introduce the concept of *subgame* and explain some notations.

Given a game $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, v_0, E, (\text{Goal}_i)_{i \in \Pi})$ and a history hv of \mathcal{G} , with $v \in V$, the *subgame* $\mathcal{G}|_h$ of \mathcal{G} is the game $(\Pi, V, (V_i)_{i \in \Pi}, v, E, (\text{Goal}_i)_{i \in \Pi})$ with initial vertex v . Given a strategy σ_i for player i in \mathcal{G} , we define the strategy $\sigma_i|_h$ in $\mathcal{G}|_h$ by $\sigma_i|_h(h') = \sigma_i(hh')$ for all history h' of $\mathcal{G}|_h$ such that $\text{Last}(h') \in V_i$. Let σ be the strategy profile $(\sigma_i)_{i \in \Pi}$, we write $\sigma|_h$ for $(\sigma_i|_h)_{i \in \Pi}$, and $h(\sigma|_h)$ for the play in \mathcal{G} with prefix h that is consistent with $\sigma|_h$ from v .

Then, we say that $(\sigma_i|_h)_{i \in \Pi}$ is a Nash equilibrium in $\mathcal{G}|_h$ if for all player $j \in \Pi$ and for all strategy σ'_j of player j , we have that $\text{Cost}_j(\rho) \leq \text{Cost}_j(\rho')$, where $\rho = h\langle (\sigma_i|_h)_{i \in \Pi} \rangle$ and $\rho' = h\langle \sigma'_j|_h, \sigma_{-j}|_h \rangle$. Let us stress on the fact that plays ρ and ρ' both include the history h as their prefix, and that the related costs $\text{Cost}_j(\rho)$ and $\text{Cost}_j(\rho')$ thus depend on h (the goal set Goal_j could have already been visited by h). The definition of a secure equilibrium in $\mathcal{G}|_h$ is given similarly.

A subgame perfect equilibrium is a strategy profile that is a Nash equilibrium after every possible history of the game, i.e. in every subgame. In particular, a subgame perfect equilibrium is also a Nash equilibrium.

Definition 4. A strategy profile $(\sigma_i)_{i \in \Pi}$ of a game \mathcal{G} is a subgame perfect equilibrium if for all history h of \mathcal{G} , $(\sigma_i|_h)_{i \in \Pi}$ is a Nash equilibrium in $\mathcal{G}|_h$.

We now introduce the last kind of equilibrium that we study. It is a new notion that combines both concepts of subgame perfect equilibrium and secure equilibrium in the following way.

Definition 5. A strategy profile $(\sigma_i)_{i \in \Pi}$ of a game \mathcal{G} is a subgame perfect secure equilibrium if for all history h of \mathcal{G} , $(\sigma_i|_h)_{i \in \Pi}$ is a secure equilibrium in $\mathcal{G}|_h$.

Notice that a subgame perfect secure equilibrium is a secure equilibrium, as well as a subgame perfect equilibrium.

In order to understand the differences between the various notions of equilibria, we provide three simple examples of games limited to two players and to finite trees. Some examples of equilibria involving cycles in two-player games can be found in [8]. They provide further motivations for the introduction of subgame perfect secure equilibria.

Example 6. Let $\mathcal{G} = (V, V_1, V_2, A, E, \text{Goal}_1, \text{Goal}_2)$ be the two-player game depicted in Fig. 1. The vertices of player 1 (resp. 2) are represented by circles (resp. squares), that is, $V_1 = \{A, D, E, F\}$ and $V_2 = \{B, C\}$. The initial vertex v_0 is A . The vertices of Goal_1 are shaded whereas the vertices of Goal_2 are doubly circled; thus $\text{Goal}_1 = \{D, F\}$ and $\text{Goal}_2 = \{F\}$. The number 2 labeling the edge (B, D) is a shortcut to indicate that there are two consecutive edges from B to D (through one intermediate vertex).




Fig. 1. Game \mathcal{G}




Fig. 2. Game \mathcal{G}'




Fig. 3. Game \mathcal{G}''

In the games \mathcal{G} , \mathcal{G}' and \mathcal{G}'' of Fig. 1, 2 and 3 (played on the same graph), we define two strategies σ_1, σ'_1 of player 1 and two strategies σ_2, σ'_2 of player 2 in the following way: $\sigma_1(A) = B, \sigma'_1(A) = C, \sigma_2(C) = E$ and $\sigma'_2(C) = F$.

In \mathcal{G} , one can easily check that the strategy profile (σ_1, σ_2) is a secure equilibrium (and thus a Nash equilibrium) with cost profile is $(3, +\infty)$. Such a secure equilibrium exists because player 2 threatens player 1 to go to vertex E in the case where vertex C is reached. This threat is not credible in this case since by acting this way, player 2 gets an infinite cost instead of a cost of 2 (that he could obtain by reaching F). For this reason, (σ_1, σ_2) is not a subgame perfect equilibrium (and thus not a subgame perfect secure equilibrium). However, one can check that the strategy profile (σ'_1, σ'_2) is a subgame perfect secure equilibrium.

Let us now consider the game \mathcal{G}' depicted in Fig. 2 (notice that the number 2 has disappeared from the edge (B, D)). One can verify that the strategy profile (σ'_1, σ'_2) is a subgame perfect equilibrium which is not a secure equilibrium (and thus not a subgame perfect secure equilibrium). A subgame perfect secure equilibrium for \mathcal{G}' is given by the strategy profile (σ_1, σ'_2) .

Finally, for the game \mathcal{G}'' depicted in Fig. 3, one can check that the strategy profile (σ_1, σ'_2) is both a subgame perfect equilibrium and a secure equilibrium. However it is not a subgame perfect secure equilibrium. In particular, this shows that being a subgame perfect secure equilibrium is *not* equivalent to be a subgame perfect equilibrium and a secure equilibrium. On the other hand, (σ_1, σ_2) is a subgame perfect secure equilibrium in \mathcal{G}'' .

2.2 Unraveling

In the proofs of this article, it will be often useful to *unravel* the graph $G = (V, (V_i)_{i \in \Pi}, v_0, E)$ from the initial vertex v_0 , which ends up in an *infinite tree*, denoted by T . This tree can be seen as a new graph where the set of vertices is the set H of histories of \mathcal{G} , the initial vertex is v_0 , and a pair $(h, hv) \in H \times H$ is an edge of T if $(\text{Last}(h), v) \in E$. A history h is a vertex of player i in T if $h \in H_i$, and h belongs to the goal set of player i if $\text{Last}(h) \in \text{Goal}_i$.

We denote by \mathcal{T} the related game. This game \mathcal{T} played on the unraveling T of G is equivalent to the game \mathcal{G} that is played on G in the following sense. A play $(\rho_0)(\rho_0\rho_1)(\rho_0\rho_1\rho_2) \dots$ in \mathcal{T} induces a unique play $\rho = \rho_0\rho_1\rho_2 \dots$ in \mathcal{G} , and conversely. Thus, we denote a play in \mathcal{T} by the respective play in \mathcal{G} . The bijection between plays of \mathcal{G} and plays of \mathcal{T} allows us to use the same cost function Cost , and to transform easily strategies in \mathcal{G} to strategies in \mathcal{T} (and conversely).

We also need to study the tree T limited to a certain depth $d \in \mathbb{N}$: we denote by $\text{Trunc}_d(T)$ the *truncated tree of T of depth d* and $\text{Trunc}_d(\mathcal{T})$ the *finite game played on $\text{Trunc}_d(T)$* . More precisely, the set of vertices of $\text{Trunc}_d(T)$ is the set of histories $h \in H$ of length $\leq d$; the edges of $\text{Trunc}_d(T)$ are defined in the same way as for T , except that for the histories h of length d , there exists no edge (h, hv) . A play ρ in $\text{Trunc}_d(\mathcal{T})$ corresponds to a history of \mathcal{G} of length *equal to d* . The notions of cost and strategy are defined exactly like in the game \mathcal{T} , but limited to the depth d . For instance, a player pays an infinite cost for a play ρ (of length d) if his goal set is not visited by ρ .

2.3 Kuhn’s Theorem

This section is devoted to the classical Kuhn’s theorem [20]. It claims the existence of a subgame perfect equilibrium (resp. subgame perfect secure equilibrium) in multiplayer games played on *finite trees*.

A *preference relation* is a total, reflexive and transitive binary relation.

Theorem 7 (Kuhn’s theorem). *Let Γ be a finite tree and \mathcal{G} a game played on Γ . For each player $i \in \Pi$, let \succsim_i be a preference relation on cost profiles. Then there exists a strategy profile $(\sigma_i)_{i \in \Pi}$ such that for all history h of \mathcal{G} , all player $j \in \Pi$, and all strategy σ'_j of player j in \mathcal{G} , we have*

$$\text{Cost}(\rho') \succsim_j \text{Cost}(\rho)$$

where $\rho = h \langle (\sigma_i|_h)_{i \in \Pi} \rangle$ and $\rho' = h \langle \sigma'_j|_h, \sigma_{-j}|_h \rangle$.

One can easily be convinced that the binary relation on cost profiles used to define the notion of Nash equilibrium (see Definition 2) is total, reflexive and transitive. We thus have the following corollary.

Corollary 8. *Let \mathcal{G} be a game and T be the unraveling of G . Let $\text{Trunc}_d(\mathcal{T})$ be the game played on the truncated tree of T of depth $d \in \mathbb{N}$. Then there exists a subgame perfect equilibrium in $\text{Trunc}_d(\mathcal{T})$.*

Let \preceq_j be the relation defined by $x \preceq_j y$ iff $x \prec_j y$ or $x = y$, where \prec_j is the relation used in Definition 3. We notice that in the two-player case, this relation is total, reflexive and transitive. However when there are more than two players, \preceq_j is no longer total. Nevertheless, it is proved in [21] that Kuhn’s theorem remains true when \preceq_j is only transitive. So, the next corollary holds.

Corollary 9. *Let \mathcal{G} be a game and T be the unraveling of G . Let $\text{Trunc}_d(\mathcal{T})$ be the game played on the truncated tree of T of depth $d \in \mathbb{N}$. Then there exists a subgame perfect secure equilibrium in $\text{Trunc}_d(\mathcal{T})$.*

3 Subgame Perfection

In this section, we positively solve Problem 1 for subgame perfect equilibria, and for subgame perfect secure equilibria in the two-player case.

Theorem 10. *In every multiplayer quantitative reachability game, there exists a subgame perfect equilibrium.*

The proof uses techniques completely different from the ones given in [6,7] for the existence of Nash equilibria, and secure equilibria in two-player games.

Let \mathcal{G} be a game and \mathcal{T} be the infinite game played on the unraveling T of \mathcal{G} . Kuhn’s theorem (and in particular Corollary 8) guarantees the existence of a subgame perfect equilibrium in each finite game $\text{Trunc}_n(\mathcal{T})$ for all depth $n \in \mathbb{N}$. Given a sequence of such equilibria, the keypoint is to derive the existence of a subgame perfect equilibrium in the infinite game \mathcal{T} . This is possible by the following lemma.

Lemma 11. *Let $(\sigma^n)_{n \in \mathbb{N}}$ be a sequence of strategy profiles such that for every $n \in \mathbb{N}$, σ^n is a strategy profile in the truncated game $\text{Trunc}_n(\mathcal{T})$. Then there exists a strategy profile σ^* in the game \mathcal{T} with the property:*

$$\forall d \in \mathbb{N}, \exists n \geq d, \sigma^* \text{ and } \sigma^n \text{ coincide on histories of length up to } d. \quad (1)$$

Proof. This result is a direct consequence of the compactness of the set of infinite trees with bounded outdegree [18]. An alternative proof is as follows. We give a tree structure, denoted by Γ , to the set of all strategy profiles in the games $\text{Trunc}_n(\mathcal{T})$, $n \in \mathbb{N}$: the nodes of Γ are the strategy profiles, and we draw an edge from a strategy profile σ in $\text{Trunc}_n(\mathcal{T})$ to a strategy profile σ' in $\text{Trunc}_{n+1}(\mathcal{T})$ if and only if σ is the restriction of σ' to histories of length less than n . It means that the nodes at depth d correspond to strategy profiles of $\text{Trunc}_d(\mathcal{T})$. We then consider the tree Γ' derived from Γ where we only keep the nodes σ^n , $n \in \mathbb{N}$, and their ancestors. Since Γ' has finite outdegree, it has an infinite path by König’s lemma. This path goes through infinitely many nodes that are ancestors of nodes in the set $\{\sigma^n, n \in \mathbb{N}\}$. Therefore there exists a strategy profile σ^* in the infinite game \mathcal{T} (given by the previous infinite path in Γ') with property (1). \square

Proof (of Theorem 10). Let $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, v_0, E, (\text{Goal}_i)_{i \in \Pi})$ be a multi-player quantitative reachability game and \mathcal{T} the game played on the unraveling of \mathcal{G} . For all $n \in \mathbb{N}$, we consider the finite game $\text{Trunc}_n(\mathcal{T})$ and get a subgame perfect equilibrium $\sigma^n = (\sigma_i^n)_{i \in \Pi}$ in this game by Corollary 8. According to Lemma 11, there exists a strategy profile σ^* in the game \mathcal{T} with property (1).

It remains to show that σ^* is a subgame perfect equilibrium in \mathcal{T} , and thus in \mathcal{G} . Let $h \in H$ be a history of the game. We have to prove that $\sigma^*|_h$ is a Nash equilibrium in $\mathcal{T}|_h$. As a contradiction, suppose that there exists a profitable deviation σ'_j for some player $j \in \Pi$ w.r.t. $\sigma^*|_h$ in $\mathcal{T}|_h$. This means that $\text{Cost}_j(\rho) > \text{Cost}_j(\rho')$ for $\rho = h\langle \sigma^*|_h \rangle$ and $\rho' = h\langle \sigma'_j|_h, \sigma^*_{-j}|_h \rangle$, that is ρ' visits Goal_j for the first time at a certain depth d , such that $|h| < d < +\infty$, and ρ visits Goal_j at a depth strictly greater than d (see Figure 4). Thus:

$$\text{Cost}_j(\rho) > \text{Cost}_j(\rho') = d.$$

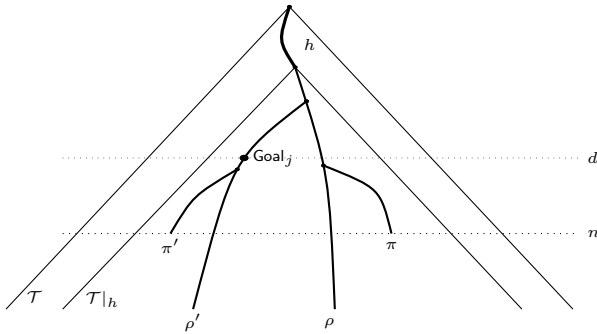


Fig. 4. The game \mathcal{T} with its subgame $\mathcal{T}|_h$

According to property (1), there exists $n \in \mathbb{N}$ such that σ^* coincide with σ^n on histories of length up to d . It follows that for $\pi = h\langle \sigma^n|_h \rangle$ and $\pi' = h\langle \sigma'_j|_h, \sigma^n_{-j}|_h \rangle$, we have that (see Figure 4)

$$\text{Cost}(\pi') = \text{Cost}(\rho') = d \quad \text{and} \quad \text{Cost}(\pi) > d.$$

And so, σ'_j is a profitable deviation for player j w.r.t. $\sigma^n|_h$ in $\text{Trunc}_n(\mathcal{T})|_h$, which leads to a contradiction with the fact that σ^n is a subgame perfect equilibrium in $\text{Trunc}_n(\mathcal{T})$ by hypothesis. □

As an extension, we consider *multiplayer quantitative reachability games with tuples of costs on edges* (see [7, Definition 31]). In these games, we assume that edges are labelled with tuples of strictly positive costs (one cost for each player). Here we do not only count the number of edges to reach the goal of a player, but we sum up his costs along the path until his goal is reached. His aim is still to minimize his global cost for a play. In this framework, we can also prove

the existence of a subgame perfect equilibrium. The proof is similar to the one of Theorem 10, the only difference lies in the choice of the different considered depths.

Theorem 12. *In every multiplayer quantitative reachability game with tuples of costs on edges, there exists a subgame perfect equilibrium.*

Regarding subgame perfect *secure* equilibria, we positively solve Problem 1 but only in the case of two-player games.

Theorem 13. *In every two-player quantitative reachability game, there exists a subgame perfect secure equilibrium.*

The main ideas of the proof are similar to the ones for Theorem 10 (see [8]). Unfortunately the proof does not seem to extend to the multiplayer case. Indeed we face the same kind of problems encountered in [6,7], where the existence of secure equilibria is proved for two-player games, and left open for multiplayer games. (See [7] for further discussion on these problems).

4 Decidability of the Existence of a Secure Equilibrium

In this section, we study Problems 1 and 2 in the context of secure equilibria. Both problems have been positively solved in [6] for two-player games only. To the best of our knowledge, the existence of secure equilibria in the multiplayer framework is still an open problem. We here provide an algorithm that *decides* the existence of a secure equilibrium. We also show that if there exists a secure equilibrium, then there exists one that is finite-memory.

Theorem 14. *In every multiplayer quantitative reachability game, one can decide whether there exists a secure equilibrium in ExpSpace .*

Theorem 15. *If there exists a secure equilibrium in a multiplayer quantitative reachability game, then there exists one that is finite-memory.*

The proof of Theorem 14 is inspired from ideas developed in [6,7]. It is rather technical and can be found in [8]. Nevertheless, let us give some flavor of the proof. The keypoint is to show that the existence of a secure equilibrium in a game \mathcal{G} is equivalent to the existence of a secure equilibrium (with two additional properties) in the finite game $\text{Trunc}_d(\mathcal{T})$ for a well-chosen depth d . The existence of the latter equilibrium is decidable. Notice that by Corollary 9 a secure equilibrium always exists in $\text{Trunc}_d(\mathcal{T})$; however we do not know if a secure equilibrium with the two required additional properties always exists in $\text{Trunc}_d(\mathcal{T})$. Let us now give some details about these two properties.

The first property requires that the secure equilibrium is *goal-optimized*, meaning that all the visited goal sets are visited for the first time *before a certain given depth*. Let $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, v_0, E, (\text{Goal}_i)_{i \in \Pi})$ be a game. We fix the following constant: $d_{\text{goal}} := 2 \cdot |\Pi| \cdot |V|$.

Definition 16. Given a strategy profile $(\sigma_i)_{i \in \Pi}$ in a game \mathcal{G} , with outcome ρ , we say that $(\sigma_i)_{i \in \Pi}$ is goal-optimized if and only if for all $i \in \Pi$ such that $\text{Cost}_i(\rho) < +\infty$, we have that $\text{Cost}_i(\rho) < d_{goal}$.

The second property asks for a secure equilibrium that is *deviation-optimized*, meaning that whenever a player deviates, he realizes *within a certain given number of steps* that his deviation is not profitable for him.

Definition 17. Given a secure equilibrium $(\sigma_i)_{i \in \Pi}$ in a game \mathcal{G} , we say that $(\sigma_i)_{i \in \Pi}$ is deviation-optimized if and only if for all player $j \in \Pi$, for all strategy σ'_j of player j ,

$$\text{Cost}(\rho_{<d_{dev}}) \not\leq_j \text{Cost}(\rho'_{<d_{dev}}),$$

where $\rho = \langle (\sigma_i)_{i \in \Pi} \rangle$, $\rho' = \langle \sigma'_j, \sigma_{-j} \rangle$, and $d_{dev} = \max\{\text{Cost}_i(\rho) \mid \text{Cost}_i(\rho) < +\infty\} + |V|$.

We can now state the key proposition.

Proposition 18. Given a game \mathcal{G} , there exists a secure equilibrium in \mathcal{G} iff there exists a goal-optimized and deviation-optimized secure equilibrium in $\text{Trunc}_d(\mathcal{T})$, for $d = d_{goal} + 3 \cdot |V|$.

The proof of this proposition is detailed in [8]. It is here difficult to give the significance of choosing constants d_{goal} , d_{dev} and $d = d_{goal} + 3 \cdot |V|$ as is done. We just propose a sketch of proof.

Proof (of Proposition 18 - Sketch).

Suppose that there exists a secure equilibrium $(\sigma_i)_{i \in \Pi}$ in \mathcal{G} . The first step consists in transforming $(\sigma_i)_{i \in \Pi}$ into a goal-optimized and deviation-optimized secure equilibrium in \mathcal{G} . To get a goal-optimized equilibrium, the idea is to eliminate unnecessary cycles between two successively visited goal sets. Such an idea is already developed in [7, Lemma 19] for Nash equilibria. Unfortunately, this lemma cannot be applied for secure equilibria. We need to adapt it to the context of secure equilibria, by modifying the strategies of the coalition against a deviating player. In this way, we get a goal-optimized equilibrium that is also deviation-optimized due to this particular form of the coalitions strategies.

Once we have a goal-optimized and deviation-optimized secure equilibrium in \mathcal{G} , the second step consists in showing that its restriction to $\text{Trunc}_d(\mathcal{T})$ with $d = d_{goal} + 3 \cdot |V|$ is still a goal-optimized and deviation-optimized secure equilibrium in $\text{Trunc}_d(\mathcal{T})$.

Suppose now that there exists a goal-optimized and deviation-optimized secure equilibrium $(\sigma_i)_{i \in \Pi}$ in $\text{Trunc}_d(\mathcal{T})$, for $d = d_{goal} + 3 \cdot |V|$. To get from $(\sigma_i)_{i \in \Pi}$ a secure equilibrium in \mathcal{G} , we inspire from a construction proposed in [7, Proposition 25] where it is shown, in the context of two-player games, how to extend a secure equilibrium in a finite truncation of \mathcal{G} to a secure equilibrium in \mathcal{G} . The rough idea is as follows. Due to the hypotheses, the outcome π of $(\sigma_i)_{i \in \Pi}$ has a prefix $\alpha\beta$ such that all goal sets visited by π are already visited by α , and such that β is a cycle. The required secure equilibrium is specified such that its

outcome is equal to $\alpha\beta^\omega$ and any deviating player is punished by the coalition of the other players in a way that this deviation is not profitable for him.³ \square

We are now able to prove the two theorems of this section.

Proof (of Theorem 14). By Proposition 18, there exists a secure equilibrium in \mathcal{G} iff there exists a goal-optimized and deviation-optimized secure equilibrium in $\text{Trunc}_d(\mathcal{T})$, with $d = d_{\text{goal}} + 3 \cdot |V|$. The latter property is decidable in NExpSpace (in $|V|$ and $|I|$). Indeed, $\text{Trunc}_d(\mathcal{T})$ has an exponential size. Guessing a strategy profile $(\sigma_i)_{i \in I}$ in this tree also needs an exponential size. Then we can test in exponential size whether $(\sigma_i)_{i \in I}$ is a goal-optimized and deviation-optimized secure equilibrium in $\text{Trunc}_d(\mathcal{T})$. By Savitch's theorem, deciding the existence of a secure equilibria is thus in ExpSpace . \square

Proof (of Theorem 15). This theorem is a direct consequence of Proposition 18. Indeed consider a secure equilibrium in a game \mathcal{G} . We first apply Proposition 18 to this strategy profile to get a goal-optimized and deviation-optimized secure equilibrium $(\sigma_i)_{i \in I}$ in $\text{Trunc}_d(\mathcal{T})$, for $d = d_{\text{goal}} + 3 \cdot |V|$. Then we apply Proposition 18, in the other direction, to the equilibrium $(\sigma_i)_{i \in I}$, to get a secure equilibrium back in \mathcal{G} . The latter equilibrium can be supposed to be finite-memory as explained in the proof of Proposition 18 (see Footnote 5). \square

5 Conclusion and Perspectives

In this paper, we study the concept of subgame perfect equilibrium, a refinement of Nash equilibrium well-suited in the framework of games played on graphs. We also introduce the new concept of subgame perfect secure equilibrium. We prove the existence of subgame perfect equilibria in multiplayer quantitative reachability games. We also prove the existence of subgame perfect secure equilibria, but only in the two-player framework. Finally, we provide an algorithm deciding in ExpSpace the existence of secure equilibria in the multiplayer case. On one hand, the first two results have been obtained by topological techniques, that are completely different from the techniques used in [6,7]. On the other hand, proofs of the last result are strongly inspired by proofs developed in these references, but have required new ideas about the coalition strategies.

There are several interesting directions for future research. We are currently working on the model of quantitative game, enriched by allowing n -tuples of positive weights on edges (see Theorem 12). We do believe that our results remain true in this context. The case of Nash equilibria is already treated in [7]. Notice that our results trivially generalize to the particular case where the weights of the edges are of the form (c, \dots, c) with $c \in \mathbb{N}_0$. Indeed it is enough to replace each such edge by a path of length c composed of c new edges (of cost 1).

To the best of our knowledge, the existence of secure equilibria in the multiplayer framework is still an open problem. We prove that the existence of a secure

³ It should be noted that this secure equilibrium can be constructed in a way to be finite-memory.

equilibrium in an infinite game is equivalent to the existence of a goal-optimized and deviation-optimized secure equilibrium in a finite game. This open problem could be positively solved if Corollary 9 could be adapted in a way to get a goal-optimized and deviation-optimized secure equilibrium in the finite game, and then by applying Proposition 18. A deeper understanding of equilibria with unnecessary cycles could also be helpful. For the moment, we are not able to solve this problem with more than two players. The same kind of question is also open for subgame perfect secure equilibria.

Another research direction concerns a deeper study of the memory needed in the different kinds of equilibria. In the case of subgame perfect equilibria and subgame perfect secure equilibria, the topological techniques give no results on the memory needed. However, in the case of secure equilibria, we prove that we can limit to finite-memory equilibria.

Acknowledgements. This work has been partly supported by the ESF project GASICS and a grant from the National Bank of Belgium. The third author is supported by a grant from L'Oréal-UNESCO/F.R.S.-FNRS.

References

1. Alur, R., Kanade, A., Weiss, G.: Ranking Automata and Games for Prioritized Requirements. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 240–253. Springer, Heidelberg (2008)
2. Bloem, R., Chatterjee, K., Henzinger, T., Jobstmann, B.: Better Quality in Synthesis through Quantitative Objectives. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 140–156. Springer, Heidelberg (2009)
3. Boros, E., Gurvich, V.: Why chess and back gammon can be solved in pure positional uniformly optimal strategies. Rutcor Research Report 21-2009. Rutgers University (2009)
4. Bouyer, P., Brenguier, R., Markey, N.: Nash Equilibria for Reachability Objectives in Multi-Player Timed Games. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 192–206. Springer, Heidelberg (2010)
5. Bouyer, P., Brenguier, R., Markey, N., Ummels, M.: Nash equilibria in concurrent games with Büchi objectives. In: Foundations of Software Technology and Theoretical Computer Science, FSTTCS. LIPIcs, vol. 13, pp. 375–386. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2011)
6. Brihaye, T., Bruyère, V., De Pril, J.: Equilibria in Quantitative Reachability Games. In: Ablayev, F., Mayr, E.W. (eds.) CSR 2010. LNCS, vol. 6072, pp. 72–83. Springer, Heidelberg (2010)
7. Brihaye, T., Bruyère, V., De Pril, J.: On equilibria in quantitative games with reachability/safety objectives. Technical report (2011), <http://www.ulb.ac.be/di/gasics/equilibria.pdf>
8. Brihaye, T., Bruyère, V., De Pril, J., Gimbert, H.: Subgame Perfection for Equilibria in Quantitative Reachability Games. Technical report (2011), <http://www.ulb.ac.be/di/gasics/subgame-perfection.pdf>
9. Chatterjee, K., Henzinger, T., Jurdziński, M.: Games with secure equilibria. Theoretical Computer Science 365(1-2), 67–82 (2006)

10. Chatterjee, K., Henzinger, T.A.: Finitary Winning in Omega-Regular Games. In: Hermanns, H. (ed.) TACAS 2006. LNCS, vol. 3920, pp. 257–271. Springer, Heidelberg (2006)
11. Chatterjee, K., Henzinger, T.A.: Assume-Guarantee Synthesis. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 261–275. Springer, Heidelberg (2007)
12. Chatterjee, K., Raman, V.: Assume-guarantee synthesis for digital contract signing. CoRR, abs/1004.2697 (2010)
13. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (2000)
14. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games. LNCS, vol. 2500. Springer, Heidelberg (2002)
15. Grädel, E., Ummels, M.: Solution concepts and algorithms for infinite multiplayer games. In: New Perspectives on Games and Interaction. Texts in Logic and Games, vol. 4, pp. 151–178. Amsterdam University Press (2008)
16. Hopcroft, J.E., Ullman, J.D.: Introduction to automata theory, languages, and computation. Addison-Wesley Series in Computer Science. Addison-Wesley Publishing Co., Reading (1979)
17. Horn, F., Thomas, W., Wallmeier, N.: Optimal Strategy Synthesis in Request-Response Games. In: Cha, S(S.), Choi, J.-Y., Kim, M., Lee, I., Viswanathan, M. (eds.) ATVA 2008. LNCS, vol. 5311, pp. 361–373. Springer, Heidelberg (2008)
18. Kechris, A.: Classical descriptive set theory. Springer, Heidelberg (1995)
19. Klimoš, M., Larsen, K.G., Štefaňák, F., Thaarup, J.: Nash Equilibria in Concurrent Priced Games. In: Dediu, A.-H. (ed.) LATA 2012. LNCS, vol. 7183, pp. 363–376. Springer, Heidelberg (2012)
20. Kuhn, H.: Extensive games and the problem of information. *Classics in Game Theory*, 46–68 (1953)
21. Le Roux, S.: Acyclic Preferences and Existence of Sequential Nash Equilibria: A Formal and Constructive Equivalence. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 293–309. Springer, Heidelberg (2009)
22. Nash, J.: Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America* 36(1), 48–49 (1950)
23. Osborne, M., Rubinstein, A.: A course in game theory. MIT Press, Cambridge (1994)
24. Paul, S., Simon, S., Kannan, R., Kumar, K.: Nash equilibrium in generalised muller games. In: Foundations of Software Technology and Theoretical Computer Science, FSTTCS. LIPIcs, vol. 4, pp. 335–346. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2009)
25. Selten, R.: Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrage-trägheit. *Zeitschrift für die gesamte Staatswissenschaft* 121, 301–324, 667–689 (1965)
26. Thomas, W.: On the Synthesis of Strategies in Infinite Games. In: Mayr, E.W., Puech, C. (eds.) STACS 1995. LNCS, vol. 900, pp. 1–13. Springer, Heidelberg (1995)
27. Thomas, W.: Church’s Problem and a Tour through Automata Theory. In: Avron, A., Dershowitz, N., Rabinovich, A. (eds.) Pillars of Computer Science. LNCS, vol. 4800, pp. 635–655. Springer, Heidelberg (2008)
28. Zimmermann, M.: Time-Optimal Winning Strategies for Poset Games. In: Maneth, S. (ed.) CIAA 2009. LNCS, vol. 5642, pp. 217–226. Springer, Heidelberg (2009)