

We thank the two reviewers for their valuable comments. In this document we explain how we addressed their comments. The modifications such as corrections of typos are done without any comment. We provide in the sequel of this document additional details to respond to comments from reviewers for those we feel this is useful.

1 Comments of reviewer 1

- *“page 14. You could consider using p^0 instead of p_0 so it does not get confused with the p_i ”*

We hope that the footnote given in page 15 avoids this confusion.

- *“Proposition 29. I do not think the proof works exactly the same way. You probably would like to consider here ψ to be in DNF instead of CNF for this to work. You should rephrase Remark 28 from page 27 at least in this context here.”*

We provide the proof (with a little bit less formalism) of the PSPACE-hardness with Safety objectives in Section 3 of this document. It proves that the reduction works if ψ is assumed to be in CNF. We give in bold the biggest differences between the proof with reachability objectives and the proof with safety objectives.

2 Comments of reviewer 2

- *“First, while a “weak subgame equilibria” may be an interesting object to consider, it is not something that can meaningfully represent the interactions of rational agents with different objectives: a player may find themselves in a losing equilibrium while they have a winning strategy; they can even fail to achieve a feasible objective in a 1-player game. They can be depended upon to perform complex operations in order to lose: in the game of Figure 1, there is an equilibrium from v_2 in which Player 1 loses simply because some unspecified mechanism forces them to eventually go to v_3 . In fact, the players never try to achieve their (prefix-independent) objectives: they just make one-step adjustments*

if a better outcome for them is immediately available. This does not make the problem simple, but it calls into question why we should be interested in that particular case.”

We have highlighted this particular phenomenon and given some justifications by adding this paragraph: “As shown by this example, a weak Nash equilibrium has to be seen as a contract where no one has an incentive to finitely deviate alone. In such a contract a player may have a profitable deviation which needs infinitely many changes. Notice that the profitable deviation σ'_1 in Example 5 is a winning strategy¹ for player 1 from v_1 . Although playing σ'_1 is optimal for player 1 this cannot be achieved from σ_1 by finitely deviating. As in the classical setting of game theory, the study of equilibria concerns stability rather than optimality. In this spirit, weak Nash equilibria can be used to model stability under finitely deviating strategies.” (p.9)

- *Why is there ” $P_k^*(v) \neq \emptyset$ for all $v \in Succ^*(v_0)$ ” in Theorem 14 and other places. This condition does not depend on p_0 , and [18] showed that there is always a weak SPE, so how could it ever be false?*

You are right. In order to make it clearer, we have rephrased the characterization (Theorem 14) by removing the condition $P_k^*(v) \neq \emptyset$ for all $v \in Succ^*(v_0)$ and the end of the proof of Lemma 30 has been changed accordingly. However, we keep this condition in the statement of Proposition 19 (since it is useful to build the good symbolic witness in $2 \Rightarrow 3$) and we explain after Example 20 how to recover Theorem 14 from Proposition 19.

- *“Be more concise. The proof of Proposition 27 is quite simple, it should not take 4 pages. For example, I don’t think that anyone will make it that far without knowing the semantics of a fully quantified Boolean formula.”*

First reviewer asked some precisions about how the clauses C_i look like.

¹A winning strategy for player 1 from a vertex v is a strategy such that whatever the strategy of player 2, player 1 is sure that he will obtain a gain of 1 if he follows this strategy from v .

Thus, we assume that all these details may be benefic. Additionally, we need all this formalism if we want to give a proper proof of the reduction. Indeed, to make clear where the notion of subgame enters in consideration, we need to write carefully the details and so to provide a long proof. Otherwise, if we make the proof too concise we may have the feeling that the constrained existence problem with reachability objectives for Nash equilibrium is PSPACE-hard. But it is proved to be NP-complete.

- *“Is there a difference between a problem being ”fixed-parameter tractable” and it becoming polynomial when [the parameter] is fixed?”*

There is a difference between a problem being FPT and a problem becoming polynomial when the parameter is fixed.

The definition of an FPT algorithm is given in our paper just before Theorem 34. Intuitively it tries to separate the complexity into two parts: one part that is polynomial and only depends on the size of the input and another part that is some exponential or worse $f(t)$ and only depends on the parameter t . This FPT notion leads to a classification of NP-hard problems on a finer scale than in the classical complexity (see reference [33] for more details).

If an algorithm is FPT, than it becomes polynomial when the parameter is fixed. But in our Theorem 34, even if the constraint problem can be solved in polynomial time when the number of players is fixed, we were not able to prove that it is FPT when the parameter is the number of players for Rabin, Streett, and Muller objectives (additional parameters were necessary to get FPT algorithms).

- *“a Boolean combination of Buchi objective is a Muller objective.”*

It is right that a Boolean combination of Büchi objective is a Muller objective. However to get our FPT results we had to use the third result of Proposition 32 proved in [30].

3 PSPACE-hardness for Safety objectives

Proposition 1. *The constraint problem for Boolean games with Safety objectives is PSPACE-hard.*

To prove this proposition, we give a polynomial reduction from the QBF problem that is PSPACE-complete. This problem is to decide whether a fully quantified Boolean formula ψ is true. The formula ψ can be assumed to be in prenex Conjunctive Normal Form (CNF) $Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ such that the quantifiers are alternating existential and universal quantifiers ($Q_1 = \exists, Q_2 = \forall, Q_3 = \exists, \dots$), $X = \{x_1, x_2, \dots, x_m\}$ is the set of quantified variables, and $\phi(X) = C_1 \wedge \dots \wedge C_n$ is an unquantified Boolean formula over X equal to the conjunction of the clauses C_1, \dots, C_n .

Such a formula ψ is true if there exists a value of x_1 such that for all values of x_2 , there exists a value of $x_3 \dots$, such that the resulting valuation ν of all variables of X evaluates $\phi(X)$ to true. Formally, for each odd (resp. even) k , $1 \leq k \leq m$, let us denote by $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$ (resp. $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$) a valuation of variable x_k given a valuation of previous variables x_1, \dots, x_{k-1} ². Given these sequences $f = f_1, f_3, \dots$ and $g = g_2, g_4, \dots$, let us denote by $\nu = \nu_{(f,g)}$ the valuation of all variables of X such that $\nu(x_1) = f_1, \nu(x_2) = g_2(\nu(x_1)), \nu(x_3) = f_3(\nu(x_1)\nu(x_2)), \dots$. Then

$$\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X) \text{ is true}$$

if and only if

there exist $f = f_1, f_3, \dots$ such that for all $g = g_2, g_4, \dots$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

Proof. The reduction from the QBF problem to the constraint existence problem for Boolean games with safety objectives is the same as for Boolean games with reachability objectives given in the proof of Proposition 27 in the paper except that each player i , $1 \leq i \leq n$, aims at avoiding the set F'_i (instead of visiting the set F_i) defined as follows:

- for all i , $1 \leq i \leq n$, $F'_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_{n+1}\}$;
- $F'_{n+1} = \{t_1, t_2, \dots, t_n\}$;

²Notice that $f_1 : \emptyset \rightarrow \{0, 1\}$.

- $F'_{n+2} = \{t_{n+1}\}$.

Let $\psi = Q_1x_1Q_2x_2 \dots Q_mx_m \phi(X)$ with $\phi(X) = C_1 \wedge \dots \wedge C_n$ be a fully quantified Boolean formula in prenex Conjunctive Normal Form. We build the following Boolean game $\mathcal{G}_\psi = (\Pi, V, (V_i)_{i \in \Pi}, E, (\text{Gain}_i)_{i \in \Pi})$ (see Figure 1):

- the set V of vertices:
 - for each variable $x_k \in X$ under quantifier Q_k , there exist vertices x_k , $\neg x_k$ and q_k ;
 - for each clause C_k , there exist vertices c_k and t_k ;
 - there exists an additional vertex t_{n+1} ;
- the set E of edges:
 - from each vertex q_k there exist an edge to x_k and an edge to $\neg x_k$;
 - from each vertex x_k and $\neg x_k$, there exists an edge to q_{k+1} , except for $k = m$ where this edge is to c_1 ;
 - from each vertex c_k , there exist an edge to t_k and an edge to c_{k+1} , except for $k = n$ where there exist an edge to t_n and an edge to t_{n+1} ;
 - there exists a loop on each t_k ;
- the set Π of $n + 2$ players:
 - each player i , $1 \leq i \leq n$, owns vertex c_i ;
 - player $n + 1$ (resp. $n + 2$) is the player who owns the vertices q_i for each existential (resp. universal) quantifier Q_i ;
 - as all other vertices have only one outgoing edge, it does not matter which player owns them;
- each function Gain_i is associated with the objective of avoiding the set F'_i defined as follows:
 - for all i , $1 \leq i \leq n$, $F'_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_{n+1}\}$;
 - $F'_{n+1} = \{t_1, t_2, \dots, t_n\}$;
 - $F'_{n+2} = \{t_{n+1}\}$.

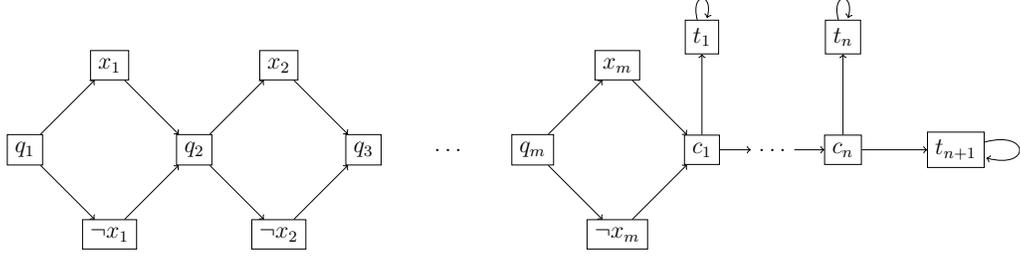


Figure 1: Reduction from the formula ψ to the Boolean game \mathcal{G}_ψ

Remark 2. (1) Notice that a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, as presented above, can be translated into a strategy σ_{n+1} of player $n+1$ in the initialized game (\mathcal{G}_ψ, q_1) , and conversely. Similarly, a sequence g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ is nothing else than a strategy σ_{n+2} of player $n+2$. (2) Notice also that if ρ is a play in (\mathcal{G}_ψ, q_1) , then $\text{Gain}_{n+1}(\rho) = 1$ if and only if $\text{Gain}_{n+2}(\rho) = 0$.

The game \mathcal{G}_ψ can be constructed from ψ in polynomial time. Let us now show that ψ is true if and only if there exists a weak SPE in (\mathcal{G}_ψ, q_1) with a payoff $p \geq (0, \dots, 0, 1, 0)$ (that is, such that $p_{n+1} = 1$).

(\Rightarrow) Suppose that ψ is true. Then there exists a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, such that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. We define a strategy profile σ as follows:

- for player $n+1$, his strategy σ_{n+1} is the strategy corresponding to the sequence f (by Remark 2);
- for player $n+2$, his strategy is an arbitrary strategy σ_{n+2} ; we denote by g the corresponding sequence $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ (by Remark 2);
- for each player i , $1 \leq i \leq n$,
 - if $hv \in \text{Hist}_i(q_1)$ with $v = c_i$, is consistent with σ_{n+1} , then $\sigma_i(hv) = c_{i+1}$ if $i \neq n$ and t_{n+1} otherwise
 - else $\sigma_i(hv) = t_i$.

First, we prove that the gain profile of the play $\rho = \langle \sigma \rangle_{q_1}$ is greater than or equal to $(0, \dots, 0, 1, 0)$. By hypothesis, the valuation $\nu_{f,g}$ evaluates $\phi(X)$

to true. That is, it evaluates all clauses C_i (with $1 \leq i \leq n$) to true. It means that ρ visits a vertex of F'_i for each $1 \leq i \leq n$ and eventually loops in t_{n+1} . It follows that: $\text{Gain}_i(\rho) = 0$ for each $1 \leq i \leq n$, $\text{Gain}_{n+1}(\rho) = 1$ and $\text{Gain}_{n+2}(\rho) = 0$.

Let us prove that σ is a weak SPE, that is, for each history $hv \in \text{Hist}(q_1)$, there is no one-shot deviating strategy in the subgame $(\mathcal{G}_{\psi|h}, v)$ that is profitable to the player who owns vertex v . This is clearly true for all $v = t_i$, $1 \leq i \leq n + 1$, since t_i has only one outgoing edge. For the other vertices v , we study two cases:

- hv is consistent with σ_{n+1} : Notice that hv may not be consistent with σ_{n+2} but with another strategy σ'_{n+2} . Let g' be the sequence corresponding with σ'_{n+2} , by hypothesis the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to true. Hence, if $\rho = h\langle\sigma_{|h}\rangle_v$, ρ visits F'_i for each $1 \leq i \leq n$ and ends in t_{n+1} . It follows that: $\text{Gain}_i(\rho) = 0$ for each $1 \leq i \leq n$, $\text{Gain}_{n+1}(\rho) = 1$ and $\text{Gain}_{n+2}(\rho) = 0$.

Assume that $v \in V_i$ and let us prove that player i has no one-shot profitable deviating strategy in $(\mathcal{G}_{|h}, v)$.

- If $i = n + 1$: since $\text{Gain}_{n+1}(\rho) = 1$, player $n + 1$ has no incentive to deviate.
- If $i = n + 2$: assume that σ''_{n+2} is a one-shot deviating strategy from σ_{n+2} in $(\mathcal{G}_{|h}, v)$. Since hv is consistent with σ_{n+1} , the resulting play if player $n + 2$ follows strategy σ''_{n+2} also loops in t_{n+1} and so the gain of player $n + 2$ is still 0. It means that there is no one-shot profitable deviation for player $n + 2$ in $(\mathcal{G}_{|h}, v)$.
- If $1 \leq i \leq n$: as $v \in V_i$, $h'c_1 \leq hv$ for some $h' \in \text{Hist}(q_1)$. Since hv is consistent with σ_{n+1} , hv visits each F'_j ($1 \leq j \leq n$) as explained previously. It means that player i has already lost in the subgame $(\mathcal{G}_{|h}, v)$ and he can't do anything to change that.

- hv is not consistent with σ_{n+1} : In this case, let $\rho = h\langle\sigma_{|h}\rangle_v$ ends in t_1 if v is a vertex “between” q_1 and c_1 and in t_j if $v = c_j$ for some $1 \leq j \leq n$.

Assume that $v \in V_i$ and let us prove that player i has no one-shot profitable deviation in $(\mathcal{G}_{|h}, v)$.

- If $i = n + 1$: whatever player $n + 1$ tries to do in the subgame, since hv is already not consistent with σ_{n+1} , player 1 will punish

player $n + 1$ and loops in t_1 . So player $n + 1$ receives a gain of 0. He has no incentive to deviate.

- If $i = n + 2$: for the same reasons as for player $n + 1$, player $n + 2$ receives a gain of 1. He has no incentive to deviate.
- If $1 \leq i \leq n$: then ρ ends in t_i . If the clause C_i is not visited along hv then $\text{Gain}_i(\rho) = 1$ and player i has no incentive to deviate. Otherwise, whatever he chooses to do in v , this choice will lead to a gain of 0 since a state of F'_j is already visited along hv . It means that, in these two cases, player i has no incentive to deviate.

This concludes the proof of this implication.

Notice that for player $n+1$ and player $n+2$, the proof for safety objectives is the same as for reachability objectives. It is explained by the fact that wanting to avoid $F'_{n+1} = \{t_1, t_2, \dots, t_n\}$ (resp. $F'_{n+2} = \{t_{n+1}\}$) is equivalent to reach $F_{n+1} = \{t_{n+1}\}$ (resp. $F_{n+2} = \{t_1, \dots, t_n\}$). It is a little bit different for a player i , $1 \leq i \leq n$. Let σ be a strategy profile, assume that we want to know if this player has an incentive to deviate in the subgame $(\mathcal{G}_{|h}, v)$ for some $hv \in \text{Hist}(q_1)$ and $v = c_i$ (as we only check one-shot deviations, it is the only case where player i has a choice to make). Let $\rho = h\langle\sigma_{|h}\rangle_v$ be the outcome of σ in this subgame.

- **if the clause of player i is visited along hv : then for reachability objectives, player i has already won and $\text{Gain}_i(\rho) = 1$, so player i has no incentive to deviate, for safety objectives, player i has already lost so whatever he does in v he will obtain a gain of 0. He has no incentive to deviate.**
- **if the clause of player i is not visited along hv : then for reachability objectives, if player i follows $\sigma_{i|h}$ he has to go in t_i and so he visits his target set and obtains a gain of 1. Thus he has no incentive to deviate. For safety objectives, the set F'_i is not visited along hv and so player i can force to win if he chooses to go to t_i (in this way, he avoids both F'_i and t_{n+1}). This is the action prescribed by $\sigma_{i|h}$. He has no incentive to deviate from this strategy.**

(\Leftarrow) We assume that there exists a weak SPE σ in (\mathcal{G}, q_1) such that $\text{Gain}(\langle\sigma\rangle_{q_1}) \geq$

$(0, \dots, 0, 1, 0)$. Since $\text{Gain}_{n+1}(\langle \sigma \rangle_{q_1}) = 1$, the play $\langle \sigma \rangle_{q_1}$ ends in t_{n+1} and it follows that $\text{Gain}_{n+2}(\langle \sigma \rangle_{q_1}) = 0$.

We have to prove that ψ is true. To this end, consider the sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, that corresponds to strategy σ_{n+1} of player $n+1$ by Remark 2. Let us show that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

By contradiction assume that it is not the case for some sequence g' and consider the related strategy σ'_{n+2} of player $n+2$ by Remark 2. Notice that σ'_{n+2} is a finitely deviating strategy. Let us consider the outcome ρ' of the strategy profile $(\sigma'_{n+2}, \sigma_{-(n+2)})$ from q_1 . As $\text{Gain}_{n+2}(\rho) = 0$, we must have $\text{Gain}_{n+2}(\rho') = 0$, otherwise σ'_{n+2} is a profitable deviation for player $n+2$ whereas σ is a weak SPE. It follows that $\text{Gain}_{n+1}(\rho') = 1$ by Remark 2, that is, ρ' eventually loops on t_{n+1} .

Now recall that the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to false, which means that it evaluates some clause C_k of $\phi(X)$ to false. Consider the history $hc_k < \rho'$. As strategy σ'_{n+2} only acts on the left part of the underlying graph of \mathcal{G}_ψ , we have $\rho' = \langle \sigma'_{n+2}, \sigma_{-(n+2)} \rangle_{q_1} = h \langle \sigma_{|h} \rangle_{c_k}$. In the subgame $(\mathcal{G}_{\psi|h}, c_k)$, the outcome of $\sigma_{|h}$ gives a gain of 0 to player k because $\rho' = h \langle \sigma_{|h} \rangle_{c_k}$ visits t_{n+1} . In this subgame, player k has thus a profitable one-shot deviation: he has to move to t_k . It follows that σ is not a weak SPE which is impossible. Then ψ is true.

Notice that this proof is quite the same as for reachability objectives. The only difference is that, for reachability objectives, in the subgame $\text{Gain}_k(\rho') = 0$ because: (i) since C_k is evaluated to false it means that F_k is not visited along hc_k and (ii) ρ' ends eventually in t_{n+1} . So if player k chooses to go to t_k he obtains a gain of 1 (as $t_k \in F_k$).

□

Constrained existence problem for weak subgame perfect equilibria with ω -regular Boolean objectives^{*,**}.

Thomas Brihaye^a, Véronique Bruyère^{b,1}, Aline Goeminne^{a,c}, Jean-François Raskin^{c,1,2}

^a*Département de Mathématique, Université de Mons (UMONS), Belgium*

^b*Département d'Informatique, Université de Mons (UMONS), Belgium*

^c*Département d'Informatique, Université libre de Bruxelles (ULB), Belgium*

Abstract

We study multiplayer turn-based games played on a finite directed graph such that each player aims at satisfying an ω -regular Boolean objective. Instead of the well-known notions of Nash equilibrium (NE) and subgame perfect equilibrium (SPE), we focus on the recent notion of weak subgame perfect equilibrium (weak SPE), a refinement of SPE. In this setting, players who deviate can only use the subclass of strategies that differ from the original one on a finite number of histories. We are interested in the constrained existence problem for weak SPEs. We provide a complete characterization of the computational complexity of this problem: it is P-complete for Explicit Muller objectives, NP-complete for Co-Büchi, Parity, Muller, Rabin, and Streett objectives, and PSPACE-complete for Reachability and Safety objec-

*This article is based upon work from COST Action GAMENET CA 16228 supported by COST (European Cooperation in Science and Technology).

**The submission is an extended and revised version of a paper that appeared in the Proceedings of GandALF 2018 (EPTCS 277, 2018, pp. 16-29). In the submitted article, all results are proved while most of those proofs were omitted or very briefly sketched in the GandALF Proceedings, and some proofs have been divided into several lemmas in order to make the arguments more structured and easier to follow.

¹Véronique Bruyère and Jean-François Raskin are supported by the PDR project “Subgame perfection in graph games” (FNRS-F.R.S.).

²Jean-François Raskin is supported by the ERC Starting Grant inVEST (279499), by the ARC project “Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond” (Fédération Wallonie- Bruxelles), and by the EOS project “Verifying Learning Artificial Intelligence Systems” (FNRS-FWO), and he is Professeur Francqui de Recherche funded by the Francqui foundation.

tives (we only prove NP-membership for Büchi objectives). We also show that the constrained existence problem is fixed parameter tractable and is polynomial when the number of players is fixed. All these results are based on a fine-grained analysis of a fixpoint algorithm that computes the set of possible payoff profiles underlying weak SPEs.

Keywords:

Multiplayer non zero-sum graph game, weak subgame perfect equilibrium, decidability of the constrained existence of an equilibrium, omega-regular objectives

1. Introduction

Two-player zero-sum graph games with ω -regular objectives are the classical mathematical model to formalize the reactive synthesis problem [1, 2]. More recently, generalization from zero-sum to non zero-sum, and from two players to n players have been considered in the literature, see e.g. [3, 4, 5, 6, 7, 8, 9, 10, 11] and the surveys [12, 13]. Those extensions are motivated by two main limitations of the classical setting. First, zero-sum games assume a fully antagonistic environment while this is often not the case in practice: the environment usually has its own goal. While the fully antagonistic assumption is simple and sound (a winning strategy against an antagonistic environment is winning against any environment that pursues its own objective), it may fail to find a winning strategy even if solutions exist when the objective of the environment is accounted. Second, modern reactive systems are often composed of several modules, and each module has its own specification and should be considered as a player on its own right. This is why we consider *n-player graph games*.

For n -player graph games, solution concepts like *Nash equilibria* (NEs) [14] are natural to consider. A strategy profile is an NE if no player has an incentive to deviate unilaterally from his strategy, *i.e.* no player can strictly improve on the outcome of the strategy profile by changing his strategy only. In the context of sequential games (such as games played on graphs), NEs allow for non-credible threats that rational players should not carry out. To avoid non-credible threats, refinements such as *subgame perfect equilibria* (SPEs) [15] have been advocated. A strategy profile is an SPE if it is an NE in all the subgames of the original game. So players need to play rationally in all subgames, and this ensures that non-credible threats cannot exist. For

applications of this concept to n -player graph games, we refer the reader to [16, 8, 11].

In [17], the notion of *weak* subgame perfect equilibrium (weak SPE) is introduced, and it is shown how it can be used to study the existence of SPEs (possibly with constraints) in quantitative reachability games. While an SPE must be resistant to any unilateral deviation of one player, a weak SPE must be resistant to deviations restricted to deviating strategies that differ from the original one on a *finite number* of histories only. In [18] the authors study general conditions on the structure of the game graph and on the preference relations of the players that guarantee the existence of a weak SPE for quantitative games. Weak SPEs retain most of the important properties of SPEs and they coincide with them when the payoff function of each player is continuous. Weak SPEs are also easier to characterize and to manipulate algorithmically. We refer the interested reader to [17, 18] for further justifications of their interest, as well as for related work on NEs and SPEs. In particular, the understanding of weak SPEs was the key to obtain the exact computational complexity of the constrained existence problem for subgame perfect equilibria in quantitative reachability games (see [19]).

Main contributions. In this paper, we concentrate on graph games with ω -regular Boolean objectives. While SPEs, and thus weak SPEs, are always guaranteed to exist in such games, we here study the computational complexity of the *constrained existence problem* for weak SPEs, *i.e.* equilibria in which some designated players have to win and some other ones have to loose. More precisely, our main results are as follows:

- We study the constrained existence problem for games with Reachability, Safety, Büchi, Co-Büchi, Parity, Explicit Muller, Muller, Rabin, and Streett objectives. We provide a *complete characterization* of the computational complexity of this problem for all the classes of objectives with one exception: Büchi objectives. The problem is P-complete for Explicit Muller objectives, it is NP-complete for Co-Büchi, Parity, Muller, Rabin, and Streett objectives, and it is PSPACE-complete for Reachability and Safety objectives. In case of Büchi objectives, we show membership to NP but we fail to prove hardness. As a byproduct we obtain that the complexity of the constrained existence problem for Reachability and Safety objectives is also PSPACE-complete for SPEs.
- Our complexity results rely on the identification of a *symbolic witness*

for the constrained existence of a weak SPE, the size of which allows us to prove NP/PSPACE-membership. As the constrained existence problem is PSPACE-complete for Reachability and Safety objectives, symbolic witnesses as compact as those for the other objectives cannot exist unless $\text{NP} = \text{PSPACE}$. The identification of symbolic witnesses is obtained thanks to a *fixpoint algorithm* that computes the set of all possible payoff profiles underlying weak SPEs.

- When the number of players is fixed, we show that the constrained existence problem can be solved in *polynomial* time for all ω -regular objectives. We also prove that it is *fixed parameter tractable* where the parameter is the number of players, for Reachability, Safety, Büchi, Co-Büchi, and Parity objectives. For Rabin, Streett, and Muller objectives, we still establish fixed parameter tractability but we need to consider some additional parameters depending on the objectives. These tractability results are obtained by a fine-grained analysis of the complexity of the fixpoint algorithm mentioned previously.

Related work and additional contributions. In [13, 11], a tree automata-based algorithm is given to decide the constrained existence problem for SPEs on graph games with ω -regular objectives defined by parity conditions. A complexity gap is left open: this algorithm executes in EXPTIME and NP-hardness of the decision problem is proved. In this paper, we focus on weak SPEs for which we provide precise complexity results for the constrained existence problem. We also observe that our results on Reachability and Safety objectives transfer from weak SPEs to SPEs: the constrained existence problem for SPEs is PSPACE-complete for those objectives. *Quantitative* Reachability objectives are investigated in [17] where it is proved that the constrained existence problem for weak SPEs and SPEs is decidable, but its exact complexity is left open.

In [17, 18, 20], the existence of (weak) SPEs in graph games is established using a construction based on a fixpoint. Our fixpoint algorithm is mainly inspired by the fixpoint technique of [18]. However, we provide complexity results based on this fixpoint while transfinite induction is used in [18]. Furthermore, we have modified the technique of [18] in a way to get a fixpoint that contains exactly all the possible payoff profiles of weak SPEs. This is necessary to get a decision algorithm for the constrained existence problem.

Profiles of strategies with finite-memory are more appealing from a practical point of view. It was shown in [11] that when there exists an SPE in a

graph game with ω -regular objectives, then there exists one that uses finite-memory strategies and has the same payoff profile. Thanks to the symbolic witnesses, we have refined those results for weak SPEs.

Structure of the paper. In Section 2, we recall the notions of n -player graph games and of (weak) SPE, and we state the studied constrained existence problem. In Section 3, we provide a fixpoint algorithm that computes all the possible payoff profiles for weak SPEs on a given graph game. From this fixpoint, we derive symbolic witnesses of weak SPEs. In Section 4, we study the complexity classes of the constrained existence problem for all objectives except Explicit Muller objectives. In Section 5, we prove the fixed parameter tractability of the constrained existence problem and we show that is in polynomial time when the number of players is fixed. We also show that this problem is P-complete for Explicit Muller objectives. In Section 6, we give a conclusion and propose future work.

2. Preliminaries

In this section, we introduce multiplayer graph games in which each player aims to achieve his Boolean objective. We focus on classical ω -regular objectives, like Reachability, Büchi, and so on. We recall two classical concepts of equilibria: Nash equilibrium and subgame perfect equilibrium (see [13]). We also recall weak variants of these equilibria as proposed in [17, 18]. We conclude the section by the constrained existence problem that is studied in this paper.

2.1. Multiplayer Boolean games

Definition 1 (Boolean game). A *multiplayer Boolean game* is a tuple $\mathcal{G} = (\Pi, V, (V_i)_{i \in \Pi}, E, (\text{Gain}_i)_{i \in \Pi})$ where

- $\Pi = \{1, 2, \dots, n\}$ is a finite set of n players;
- $G = (V, E)$ is a finite directed graph and for all $v \in V$ there exists $v' \in V$ such that $(v, v') \in E$;
- $(V_i)_{i \in \Pi}$ is a partition of V between the players;
- $\text{Gain} = (\text{Gain}_i)_{i \in \Pi}$ is a tuple of functions $\text{Gain}_i : V^\omega \rightarrow \{0, 1\}$ that assigns a Boolean value to each infinite path of G for player i .

A *play* in \mathcal{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 $h = h_0h_1\dots h_n$ ($n \in \mathbb{N}$) defined similarly. We denote the set of plays by Plays and the set of histories by Hist . Moreover, the set Hist_i is the set of histories such that the last vertex v is a vertex of player i , i.e. $v \in V_i$. The *length* $|h|$ of h is the number n of its edges. 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 $h\ell$. For all $h \in \text{Hist}$, we denote by $\text{First}(h)$ the first vertex h_0 of h . We use notation $h < \rho$ when a history h is prefix of a play (or a history) ρ . Given a play $\rho = \rho_0\rho_1\dots$, the set $\text{Occ}(\rho) = \{v \in V \mid \exists k, \rho_k = v\}$ is the set of vertices *visited* by ρ , and $\text{Inf}(\rho) = \{v \in V \mid \forall k, \exists j \geq k, \rho_j = v\}$ is the set of vertices *infinitely often visited* by ρ . Given a vertex $v \in V$, $\text{Succ}(v) = \{v' \mid (v, v') \in E\}$ is the set of successors of v , and $\text{Succ}^*(v)$ is the set of vertices reachable from v in G .

When an *initial* vertex $v_0 \in V$ is fixed, we call (\mathcal{G}, v_0) an *initialized game*. A play (resp. a history) of (\mathcal{G}, v_0) is 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 $v \in V_i$.

The goal of each player i is to achieve his objective, *i.e.*, to maximize his gain.

Definition 2 (Objective). For each player $i \in \Pi$, let $\text{Obj}_i \subseteq V^\omega$ be his *objective*. In the setting of multiplayer Boolean game, the gain function Gain_i is defined such that $\text{Gain}_i(\rho) = 1$ (resp. $\text{Gain}_i(\rho) = 0$) if and only if $\rho \in \text{Obj}_i$ (resp. $\rho \notin \text{Obj}_i$).

An objective Obj_i (or the related gain function Gain_i) is *prefix-independent* if for all $h \in V^*$ and $\rho \in V^\omega$, we have $\rho \in \text{Obj}_i$ if and only if $h\rho \in \text{Obj}_i$. In this paper, we focus on classical ω -regular objectives: Reachability, Safety, Büchi, Co-Büchi, Parity, Explicit Muller, Muller, Rabin, and Streett and we suppose that each player has the *same type* of objective. For instance, we say that \mathcal{G} is a *Boolean game with Büchi objectives* to express that all players have a Büchi objective.

Definition 3 (Classical ω -regular objectives). The set Obj_i is a *Reachability, Safety, Büchi, Co-Büchi, Parity, Explicit Muller, Muller, Rabin, or Streett* objective for player i if and only if Obj_i is composed of the plays ρ satisfying:

- *Reachability*: given $F \subseteq V$, $\text{Occ}(\rho) \cap F \neq \emptyset$;

- *Safety*: given $F \subseteq V$, $\text{Occ}(\rho) \cap F = \emptyset$;
- *Büchi*: given $F \subseteq V$, $\text{Inf}(\rho) \cap F \neq \emptyset$;
- *Co-Büchi*: given $F \subseteq V$, $\text{Inf}(\rho) \cap F = \emptyset$;
- *Parity*: given a coloring function $\Omega : V \rightarrow \{1, \dots, d\}$, $\max(\text{Inf}(\Omega(\rho)))^3$ is even;
- *Explicit Muller*: given $\mathcal{F} \subseteq 2^V$, $\text{Inf}(\rho) \in \mathcal{F}$;
- *Muller*: given a coloring function $\Omega : V \rightarrow \{1, \dots, d\}$, and $\mathcal{F} \subseteq 2^{\Omega(V)}$, $\text{Inf}(\Omega(\rho)) \in \mathcal{F}$;
- *Rabin*: given $(G_j, R_j)_{1 \leq j \leq k}$ a family of pair of sets $G_j, R_j \subseteq V$, there exists $j \in 1, \dots, k$ such that $\text{Inf}(\rho) \cap G_j \neq \emptyset$ and $\text{Inf}(\rho) \cap R_j = \emptyset$;
- *Streett*: given $(G_j, R_j)_{1 \leq j \leq k}$ a family of pair of sets $G_j, R_j \subseteq V$, for all $j \in 1, \dots, k$, $\text{Inf}(\rho) \cap G_j = \emptyset$ or $\text{Inf}(\rho) \cap R_j \neq \emptyset$.

All these objectives are prefix-independent except Reachability and Safety objectives.

A *strategy* of a player $i \in \Pi$ is a function $\sigma_i : \text{Hist}_i \rightarrow V$. This function 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 $\rho_k \in V_i$ we have that $\sigma_i(\rho_0 \dots \rho_k) = \rho_{k+1}$. 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 $hv \in \text{Hist}_i$. It is *finite-memory* if it can be encoded by a deterministic *Moore machine* $\mathcal{M} = (M, m_0, \alpha_u, \alpha_n)$ where M is a finite set of states (the memory of the strategy), $m_0 \in M$ is the initial memory state, $\alpha_u : M \times V \rightarrow M$ is the update function, and $\alpha_n : M \times V_i \rightarrow V$ is the next-action function. The Moore machine \mathcal{M} defines a strategy σ_i such that $\sigma_i(hv) = \alpha_n(\widehat{\alpha}_u(m_0, h), v)$ for all histories $hv \in \text{Hist}_i$, where $\widehat{\alpha}_u$ extends α_u to histories as expected. The *size* of the strategy σ_i is the size $|M|$ of its machine \mathcal{M} . Note that σ_i is positional when $|M| = 1$.

A *strategy profile* is a tuple $\sigma = (\sigma_i)_{i \in \Pi}$ of strategies, one for each player. It is called positional (resp. finite-memory) if for all $i \in \Pi$, σ_i is positional (resp. finite-memory). Given an initialized game (\mathcal{G}, v_0) and a strategy

³Where $\Omega(\rho) = \Omega(\rho_0)\Omega(\rho_1) \dots \Omega(\rho_n) \dots$

profile σ , there exists an unique play from v_0 consistent with each strategy σ_i . We call this play the *outcome* of σ and it is denoted by $\langle \sigma \rangle_{v_0}$. Let $p = (p_i)_{i \in \Pi} \in \{0, 1\}^{|\Pi|}$, we say that σ is a strategy profile *with payoff* p or that $\langle \sigma \rangle_{v_0}$ *has payoff* p if $p_i = \text{Gain}_i(\langle \sigma \rangle_{v_0})$ for all $i \in \Pi$.

2.2. Solution concepts

In the multiplayer game setting, the solution concepts usually studied are *equilibria* (see [13]). We here recall the concepts of Nash equilibrium and subgame perfect equilibrium, as well as some variants. We begin with the notion of deviating strategy.

Let $\sigma = (\sigma_i)_{i \in \Pi}$ be a strategy profile in an initialized Boolean game (\mathcal{G}, v_0) . Given $i \in \Pi$, a strategy $\sigma'_i \neq \sigma_i$ is a *deviating* strategy of player i , and (σ'_i, σ_{-i}) denotes the strategy profile σ where σ'_i replaces σ_i . Such a strategy is a *profitable deviation* for player i if $\text{Gain}_i(\langle \sigma \rangle_{v_0}) < \text{Gain}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$. We say that σ'_i is *finitely deviating* from σ_i if σ'_i and σ_i only differ on a finite number of histories, and that σ'_i is *one-shot deviating* from σ_i if σ'_i and σ_i only differ on v_0 [17, 18].

The notion of Nash equilibrium (NE) is classical: a strategy profile σ in an initialized game (\mathcal{G}, v_0) is a *Nash equilibrium* if no player has an incentive to deviate unilaterally from his strategy since he has no profitable deviation, *i.e.*, for each $i \in \Pi$ and each deviating strategy σ'_i of player i from σ_i , the following inequality holds: $\text{Gain}_i(\langle \sigma \rangle_{v_0}) \geq \text{Gain}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$. In this paper we focus on two variants of NE: weak/very weak NE [17, 18].

Definition 4 (Weak/very weak Nash equilibrium). A strategy profile σ is a *weak NE* (resp. *very weak NE*) in (\mathcal{G}, v_0) if, for each player $i \in \Pi$, for each finitely (resp. one-shot) deviating strategy σ'_i of player i from σ_i , we have $\text{Gain}_i(\langle \sigma \rangle_{v_0}) \geq \text{Gain}_i(\langle \sigma'_i, \sigma_{-i} \rangle_{v_0})$.

Example 5. Figure 1 illustrates an initialized Boolean game (\mathcal{G}, v_0) with Büchi objectives in which there exists a weak NE that is not an NE. In this game, player 1 (resp. player 2) owns round (resp. square) vertices and wants to visits v_1 (resp. v_3 or v_5) infinitely often. The positional strategy profile $\sigma = (\sigma_1, \sigma_2)$ is depicted by dashed arrows, its outcome is equal to $\langle \sigma \rangle_{v_0} = v_0 v_1 v_2 v_3^\omega$, and σ has payoff $(0, 1)$. Notice that player 1 has an incentive to deviate from his strategy σ_1 with a strategy σ'_1 that goes to v_1 for all histories ending in v_2 . This is indeed a profitable deviation for him since $\text{Gain}(\langle (\sigma'_1, \sigma_2) \rangle_{v_0}) = (1, 0)$. So, σ is not an NE. Nevertheless, it is a weak NE

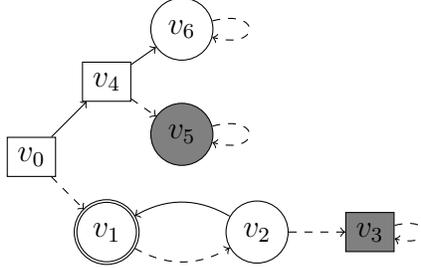


Figure 1: Example of a Boolean game with Büchi objectives. Player 1 (resp. player 2) owns round (resp. square) vertices and wants to visits v_1 (resp. v_3 or v_5) infinitely often.

because σ'_1 is the only profitable deviation and it is not finitely deviating (it differs from σ_1 on all histories of the form $v_0(v_1v_2)^k$ for $k \in \mathbb{N} \setminus \{0\}$).

As shown by this example, a weak Nash equilibrium has to be seen as a contract where no one has an incentive to finitely deviate alone. In such a contract a player may have a profitable deviation which needs infinitely many changes. Notice that the profitable deviation σ'_1 in Example 5 is a winning strategy⁴ for player 1 from v_1 . Although playing σ'_1 is optimal for player 1 this cannot be achieved from σ_1 by finitely deviating. As in the classical setting of game theory, the study of equilibria concerns stability rather than optimality. In this spirit, weak Nash equilibria can be used to model stability under finitely deviating strategies.

When considering games played on graphs, a well-known refinement of NE is the concept of *subgame perfect equilibrium* (SPE) which a strategy profile being an NE in each subgame. Variants of weak/very weak SPE can also be studied as done with NEs. Formally, given an initialized Boolean game (\mathcal{G}, v_0) and a history $hv \in \text{Hist}(v_0)$, the initialized game $(\mathcal{G}_{|h}, v)$ is called a *subgame*⁵ of (\mathcal{G}, v_0) such that $\mathcal{G}_{|h} = (\Pi, V, (V_i)_{i \in \Pi}, E, \text{Gain}_{|h})$ and $\text{Gain}_{i|h}(\rho) = \text{Gain}_i(h\rho)$ for all $i \in \Pi$ and $\rho \in V^\omega$. Moreover if σ_i is a strategy for player i in (\mathcal{G}, v_0) , then $\sigma_{i|h}$ denotes the strategy in $(\mathcal{G}_{|h}, v)$ such that for all histories $h' \in \text{Hist}_i(v)$, $\sigma_{i|h}(h') = \sigma_i(hh')$. Similarly, from a strategy profile σ in (\mathcal{G}, v_0) , we derive the strategy profile $\sigma_{|h}$ in $(\mathcal{G}_{|h}, v)$. The play

⁴A winning strategy for player 1 from a vertex v is a strategy such that whatever the strategy of player 2, player 1 is sure that he will obtain a gain of 1 if he follows this strategy from v .

⁵Notice that (\mathcal{G}, v_0) is subgame of itself.

$\langle \sigma_{\upharpoonright h} \rangle_v$ is called a *subgame outcome* of σ .

Definition 6 (Subgame perfect equilibrium and weak/very weak subgame perfect equilibrium). A strategy profile σ is a (*resp. weak, very weak*) *subgame perfect equilibrium* in (\mathcal{G}, v_0) if for all $hv \in \text{Hist}(v_0)$, $\sigma_{\upharpoonright h}$ is a (*resp. weak, very weak*) NE in $(\mathcal{G}_{\upharpoonright h}, v)$.

When one needs to show that a strategy profile is a weak SPE, the next proposition is very useful because it states that it is enough to consider one-shot deviating strategies.

Proposition 7 ([17]). *A strategy profile σ is a weak SPE if and only if σ is a very weak SPE.*

Example 8. In Example 5 is given a weak NE σ in the game (\mathcal{G}, v_0) depicted in Figure 1. This strategy profile is also a very weak SPE (and thus a weak SPE by Proposition 7). For instance, in the subgame $(\mathcal{G}_{\upharpoonright h}, v)$ with $h = v_0v_1$ and $v = v_2$, the only one-shot deviating strategy σ'_1 is such that σ'_1 coincides with $\sigma_{\upharpoonright h}$ except that $\sigma'_1(v_2) = v_1$. This is not a profitable deviation for player 1 in $(\mathcal{G}_{\upharpoonright h}, v)$. Notice that σ is not an SPE since it is not an NE as explained in Example 5.

In general, the notions of SPE and weak SPE are not equivalent (see Example 8). Nevertheless they coincide for the class of Boolean games with Reachability objectives.

Proposition 9. *Let σ be a strategy profile in an initialized Boolean game (\mathcal{G}, v_0) with Reachability objectives. Then σ is an SPE if and only if σ is a weak SPE.*

Proof. Each player i has a Reachability objective, let F_i be the set of vertices he aims to visit.

(\Rightarrow) This implication is a consequence of the definitions of SPE and weak SPE.

(\Leftarrow) Let σ be a weak SPE in (\mathcal{G}, v_0) . Assume that σ is not an SPE, *i.e.*, there exists $hv \in \text{Hist}(v_0)$ such that $\sigma_{\upharpoonright h}$ is not an NE in $(\mathcal{G}_{\upharpoonright h}, v)$. Then some player i has a profitable deviation σ'_i in the subgame $(\mathcal{G}_{\upharpoonright h}, v)$. As Gain_i takes its values in $\{0, 1\}$, this means that

$$0 = \text{Gain}_i(h\rho) < \text{Gain}_i(h\rho') = 1$$

with $\rho = \langle \sigma_{\uparrow h} \rangle_v$ and $\rho' = \langle \sigma'_i, \sigma_{-i \uparrow h} \rangle_v$. We consider the first occurrence of a vertex of F_i along $h\rho'$ (which appears in ρ' and not in h as $\text{Gain}_i(h\rho) = 0$): let g' of minimal length such that $hg' < h\rho'$ and g' ends in some $v' \in F_i$. Let us define a strategy τ_i that is finitely deviating from $\sigma_{i \uparrow h}$ and profitable for player i in $(\mathcal{G}_{\uparrow h}, v)$. This will be in contradiction with our hypothesis. For all $g \in \text{Hist}_i(v)$, let

$$\tau_i(g) = \begin{cases} \sigma'_i(g) & \text{if } g \leq g' \\ \sigma_{i \uparrow h}(g) & \text{otherwise} \end{cases}.$$

By definition of τ_i , we have that $\text{Gain}_i(h\langle \tau_i, \sigma_{-i \uparrow h} \rangle_v) = \text{Gain}_i(h\rho') = 1$ and τ_i is finitely deviating from $\sigma_{i \uparrow h}$ since $|g'|$ is finite. \square

2.3. Constraint problem

It is proved in [18] that there always exists a weak SPE in Boolean games. In this paper, we are interested in solving the following *constraint problem*:

Definition 10 (Constraint problem). Given an initialized Boolean game (\mathcal{G}, v_0) and thresholds $x, y \in \{0, 1\}^{|\Pi|}$, decide whether there exists a weak SPE in (\mathcal{G}, v_0) with payoff p such that $x \leq p \leq y$.⁶

In the next sections, we solve the constraint problem for the classical ω -regular objectives. The complexity classes that we obtain are shown in Table 1. They are detailed in Section 4 with the case of Explicit Muller objectives postponed to Section 5.3. The case of Büchi objectives remains open, since we only propose a non-deterministic algorithm in polynomial time but no matching lower bound. In Section 5, we prove that the constraint problem for weak SPEs is fixed parameter tractable and becomes polynomial when the number of players is fixed. All these results are based on a characterization of the set of possible payoffs of a weak SPE, that is described in Section 3.

Table 1: Complexity classes of the constraint problem for ω -regular objectives

	Expl. Muller	Co-Büchi	Parity	Muller	Rabin	Streett	Reach.	Safety
P-complete	×							
NP-complete		×	×	×	×	×		
PSPACE-complete							×	×

⁶The order \leq is the componentwise order, that is, $x_i \leq p_i \leq y_i$, for all $i \in \Pi$.

3. Characterization

In this section our aim is twofold: first, we characterize the set of possible payoffs of weak SPEs and second, we show how it is possible to build a weak SPE given a set of lassoes with some “good properties”. Those characterizations work for Boolean games with *prefix-independent gain functions*. We make this hypothesis all along Section 3.

3.1. Remove-Adjust procedure

Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. The computation of the set of all the payoffs of weak SPEs in (\mathcal{G}, v_0) is inspired by a fixpoint procedure explained in [18]. Each vertex v is *labeled* by a set of payoffs $p \in \{0, 1\}^{|\mathbb{I}|}$. Initially, these payoffs are those for which there exists a play in $\text{Plays}(v)$ with payoff p . Then step by step, some payoffs are removed for the labeling of v as soon as we are sure they cannot be the payoff of $\sigma_{|h}$ in a subgame $(\mathcal{G}_{|h}, v)$ for some weak SPE σ .⁷ When a fixpoint is reached, the labeling of the initial vertex v_0 exactly contains all the payoffs of weak SPEs in (\mathcal{G}, v_0) . Hence, at each step of this procedure, the payoffs labeling a vertex v are payoffs of *potential* subgame outcomes of a weak SPE. Their number decreases until reaching a fixpoint.

We formally proceed as follows. For all $v \in V$, we define the initial labeling of v as:

$$\mathbf{P}_0(v) = \{p \in \{0, 1\}^{|\mathbb{I}|} \mid \text{there exists } \rho \in \text{Plays}(v) \text{ such that } \text{Gain}(\rho) = p\}.$$

Then for each step $k \in \mathbb{N} \setminus \{0\}$, we compute the set $\mathbf{P}_k(v)$ by alternating between two operations: *Remove* and *Adjust*. To this end, we need to introduce the notion of (p, k) -labeled play. Let p be a payoff and k be a step, a play $\rho = \rho_0\rho_1\rho_2\dots$ is (p, k) -labeled if for all $j \in \mathbb{N}$ we have $p \in \mathbf{P}_k(\rho_j)$, that is, ρ visits only vertices that are labeled by p at step k . We first give some intuition about the Remove-Adjust procedure and then give the definition.

We start with the Remove operation. Let p that labels vertex v . This means that it is the payoff of a potential subgame outcome of a weak SPE that starts in v . Suppose that v is a vertex of player i and v has a successor

⁷The value of h is not important since the gain functions are prefix independent. This is why we only focus on v and not on hv .

v' such that $p_i < p'_i$ for all p' labeling v' . Then p cannot be the payoff of $\sigma_{|h}$ in the subgame $(\mathcal{G}_{|h}, v)$ for some weak SPE σ and some history h , otherwise player i would have a profitable (one-shot) deviation by moving from v to v' in this subgame.

Let us now explain the Adjust operation. It may happen that for another vertex u having p in its labeling, all potential subgame outcomes of a weak SPE from u with payoff p necessarily visit vertex v . As p has been removed from the labeling of v , these potential plays do no longer survive and p is also removed from the labeling of u by the Adjust operation.

Let us now formally define the Remove-Adjust procedure.

Definition 11 (Remove-Adjust procedure). Let $k \in \mathbb{N} \setminus \{0\}$.

- If k is odd, process the *Remove* operation:
 - If for some $v \in V_i$ there exists $p \in \mathbf{P}_{k-1}(v)$ and $v' \in \text{Succ}(v)$ such that $p_i < p'_i$ for all $p' \in \mathbf{P}_{k-1}(v')$, then $\mathbf{P}_k(v) = \mathbf{P}_{k-1}(v) \setminus \{p\}$ and for all $u \neq v$, $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$.
 - If such a vertex v does not exist, then $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$ for all $u \in V$.
- If k is even, process the *Adjust* operation:
 - If some payoff p was removed from $\mathbf{P}_{k-2}(v)$ (that is, $\mathbf{P}_{k-1}(v) = \mathbf{P}_{k-2}(v) \setminus \{p\}$), then
 - * For all $u \in V$ such that $p \in \mathbf{P}_{k-1}(u)$, check whether there still exists a $(p, k-1)$ -labeled play with payoff p from u . If it is the case, then $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$, otherwise $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u) \setminus \{p\}$.
 - * For all $u \in V$ such that $p \notin \mathbf{P}_{k-1}(u)$: $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$.
 - Otherwise $\mathbf{P}_k(u) = \mathbf{P}_{k-1}(u)$ for all $u \in V$.

We can state the existence of a fixpoint of the sequences $(\mathbf{P}_k(v))_{k \in \mathbb{N}}$, $v \in V$, in the following meaning:

Proposition 12 (Existence of a fixpoint). *There exists an even natural number $k^* \in \mathbb{N}$ such that for all $v \in V$, $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v) = \mathbf{P}_{k^*+2}(v)$.*

Proof. For all $v \in V$, the sequence $(\mathbf{P}_k(v))_{k \in \mathbb{N}}$ is nonincreasing because the Remove and Adjust operations never add a new payoff. As each $\mathbf{P}_0(v)$ is finite (it contains at most $2^{|\Pi|}$ payoffs), there exists a natural odd number $k^* + 1$ such that for all $v \in V$, $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v)$ during the Remove operation, and thus for all $v \in V$, $\mathbf{P}_{k^*+1}(v) = \mathbf{P}_{k^*+2}(v)$ during the Adjust operation. \square

Example 13. We illustrate the different steps of the Remove-Adjust procedure on the example depicted in Figure 1, and we display the result of this computation in Table 2. Initially, the sets $\mathbf{P}_0(v)$, $v \in V$, contains all payoffs p such that there exists a play $\rho \in \text{Plays}(v)$ with $\text{Gain}(\rho) = p$. At step $k = 1$, we apply a Remove operation to $v = v_4$ (this is the only possible v): v is a vertex of player $i = 2$ and v has a successor $v' = v_5$ such that $(0, 1) \in \mathbf{P}_0(v_5)$. Therefore $(0, 0)$ is removed from $\mathbf{P}_0(v_4)$ to get $\mathbf{P}_1(v_4)$. By definition of the Remove operation, the other sets $\mathbf{P}_0(u)$ are not modified and are thus equal to $\mathbf{P}_1(u)$. At step $k = 2$, we apply an Adjust operation. The only way to obtain the payoff $(0, 0)$ from v_0 is by visiting v_4 with the play $v_0 v_4 v_6^\omega$. As there does not exist a $((0, 0), 1)$ -labeled play with payoff $(0, 0)$ anymore, we have to remove $(0, 0)$ from $\mathbf{P}_1(v_0)$. The other sets $\mathbf{P}_1(v)$ remain unchanged. At step $k = 3$, the Remove operation removes payoff $(1, 0)$ from $\mathbf{P}_2(v_0)$ due to the unique payoff $(0, 1)$ in $\mathbf{P}_2(v_4)$. At step $k = 4$, the Adjust operation leaves all sets $\mathbf{P}_3(v)$ unchanged. Finally at step $k = 5$, the Remove operation also leaves all sets $\mathbf{P}_4(v)$ unchanged, and the fixpoint is reached. Therefore, we have $k^* = 4$.

Table 2: Computation of the fixpoint on the example of Figure 1

	v_0	v_1	v_2	v_3	v_4	v_5	v_6
$\mathbf{P}_0(v)$	$\{(0, 0), (1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(\mathbf{0}, \mathbf{0}), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_1(v)$	$\{(\mathbf{0}, \mathbf{0}), (1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_2(v)$	$\{(\mathbf{1}, \mathbf{0}), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_3(v)$	$\{(0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$
$\mathbf{P}_4(v)$	$\{(0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(1, 0), (0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 1)\}$	$\{(0, 0)\}$

3.2. Characterization and good symbolic witness

The fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$, provides a characterization of the payoffs of all weak SPEs as described in the following theorem. This result is in the spirit of the classical *Folk Theorem* which characterizes the payoffs of all NEs in infinitely repeated games (see for instance [21, Chapter 8]).

Theorem 14 (Characterization). *Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. Then there exists a weak SPE σ with payoff p_0 in (\mathcal{G}, v_0) if and only if $p_0 \in \mathbf{P}_{k^*}(v_0)$.⁸*

The rest of this section is devoted to the proof of Theorem 14. We begin with a lemma that states that if a given payoff p survives at step k (where k is even) in the labeling of v , this means that there exists a play with payoff p from v that only visits vertices also labeled by p .

Lemma 15. *For all even k and in particular for $k = k^*$, p belongs to $\mathbf{P}_k(v)$ if and only if there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$.*

Proof. (\Leftarrow) Suppose that there exists a (p, k) -labeled play $\rho = \rho_0\rho_1 \dots \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$. By definition of a (p, k) -labeled play, we have $p \in \mathbf{P}_k(\rho_j)$ for all j , and so in particular for $j = 0$.

(\Rightarrow) Let us prove that if p belongs to $\mathbf{P}_k(v)$, then there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$. We proceed by induction on k . For $k = 0$, the assertion is satisfied by definition of $\mathbf{P}_0(v)$ and because Gain_i is prefix-independent for all $i \in \Pi$.

Suppose that the assertion is true for an even k and let us prove that it remains true for $k + 2$. Let $p \in \mathbf{P}_{k+2}(v)$. As $\mathbf{P}_{k+2}(v) \subseteq \mathbf{P}_{k+1}(v) \subseteq \mathbf{P}_k(v)$, we have $p \in \mathbf{P}_k(v)$ and there exists a (p, k) -labeled play $\rho \in \text{Plays}(v)$ such that $\text{Gain}(\rho) = p$ by induction hypothesis. In other words $p \in \mathbf{P}_k(\rho_j)$ for all j .

We suppose that there exists v' such that $\mathbf{P}_{k+2}(v') \neq \mathbf{P}_k(v')$ (the fixpoint is not reached), otherwise $p \in \mathbf{P}_{k+2}(\rho_j)$ for all j and ρ is also a $(p, k + 2)$ -labeled play. Therefore the Remove operation has removed some payoff p' from one $\mathbf{P}_k(v')$ and the Adjust operation has possibly removed p' from some other $\mathbf{P}_k(u)$. If $p' \neq p$, then clearly p still belongs to each $\mathbf{P}_{k+2}(\rho_j)$ and ρ is again a $(p, k + 2)$ -labeled play. If $p' = p$, then $v' \neq v$ since $p \in \mathbf{P}_{k+2}(v)$ by hypothesis. Moreover, by the Adjust operation, this means that there exists a $(p, k + 1)$ -labeled play $\pi = \pi_0\pi_1 \dots$ with payoff p from v which never visits v' . Let us show that π is also a $(p, k + 2)$ -labeled play, that is, $p \in \mathbf{P}_{k+2}(\pi_j)$ for all j . Each suffix $\pi_j\pi_{j+1} \dots$ of π is a $(p, k + 1)$ -labeled play with payoff p thanks to prefix-independence of Gain. By the Adjust operation, it follows that $\mathbf{P}_{k+2}(\pi_j) = \mathbf{P}_{k+1}(\pi_j)$ for all j . This concludes the proof. \square

⁸We use notation $p_0 \in \{0, 1\}^{|\Pi|}$ to highlight that this is the payoff of σ from vertex v_0 . It should not be confused with any component p_i , $i \in \Pi$, of a payoff p .

The proof of Theorem 14 uses the concept of (good) symbolic witness defined hereafter but we begin with some intuition about it.

A symbolic witness \mathcal{P} is a *compact* representation of some finite-memory strategy profile σ in (\mathcal{G}, v_0) . It is a finite set of lassoes that represent some subgame outcomes of σ : the lasso ρ_{0,v_0} of \mathcal{P} represents the outcome $\langle \sigma \rangle_{v_0}$, and each other lasso $\rho_{i,v'}$ represents the subgame outcome $\langle \sigma_{|h} \rangle_{v'}$ for some particular histories $hv' \in \text{Hist}(v_0)$. The index i records that player i can move from v (the last vertex of h) to v' (with the convention that $i = 0$ for the outcome $\langle \sigma \rangle_{v_0}$). When σ is a weak SPE, the related symbolic witness \mathcal{P} is good, that is, its lassoes avoid profitable one-shot deviations between them.

Let us now define it properly.

Definition 16 (Symbolic witness). Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. Let $I \subseteq (\Pi \cup \{0\}) \times V$ be the set

$$I = \{(0, v_0)\} \cup \{(i, v') \mid \text{there exists } (v, v') \in E \text{ such that } v, v' \in \text{Succ}^*(v_0) \text{ and } v \in V_i\}.$$

A *symbolic witness* is a set $\mathcal{P} = \{\rho_{i,v} \mid (i, v) \in I\}$ such that each $\rho_{i,v} \in \mathcal{P}$ is a lasso of G with $\text{First}(\rho_{i,v}) = v$ and with length bounded by $2 \cdot |V|^2$.

A symbolic witness has thus at most $|V| \cdot |\Pi| + 1$ lassoes (by definition of I) with polynomial length.

Definition 17 (Good symbolic witness). A symbolic witness \mathcal{P} is *good* if for all $\rho_{j,u}, \rho_{i,v'} \in \mathcal{P}$, for all vertices $v \in \rho_{j,u}$ such that $v \in V_i$ and $v' \in \text{Succ}(v)$, we have $\text{Gain}_i(\rho_{j,u}) \geq \text{Gain}_i(\rho_{i,v'})$.

The condition of Definition 17 is depicted in Figure 2.

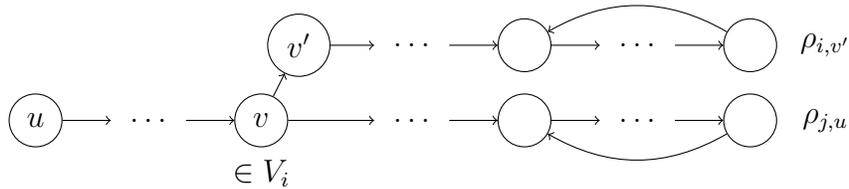


Figure 2: The condition of Definition 17

Example 18. We come back to our running example. The weak SPE of Example 8 depicted in Figure 1 has payoff $p = (0, 1)$. A symbolic witness \mathcal{P} of σ is given in Table 3 which is here composed of all the subgame outcomes of σ . One can check that \mathcal{P} is a good symbolic witness. For instance, consider its lassoes $\rho_{0,v_0} = v_0v_1v_2v_3^\omega$ and $\rho_{1,v_1} = v_1v_2v_3^\omega$, the vertex $v_2 \in V_1$ of ρ_{0,v_0} and the edge (v_2, v_1) . We have $\text{Gain}_1(\rho_{0,v_0}) \geq \text{Gain}_1(\rho_{1,v_1})$. Indeed in the subgame $(\mathcal{G}_{|v_0v_1}, v_2)$, player 1 has no profitable one-shot deviation by using the edge (v_2, v_1) .

Table 3: An example of good symbolic witness

(i, v)	$(0, v_0)$	$(2, v_4)$	$(1, v_2)$	$(1, v_1)$	$(1, v_3)$	$(2, v_5)$	$(2, v_6)$	$(1, v_5)$	$(1, v_6)$
lasso	$v_0v_1v_2v_3^\omega$	$v_4v_5^\omega$	$v_2v_3^\omega$	$v_1v_2v_3^\omega$	v_3^ω	v_5^ω	v_6^ω	v_5^ω	v_6^ω
payoff	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 1)$	$(0, 0)$	$(0, 1)$	$(0, 0)$

In Proposition 19 below, we are going to prove that there exists a weak SPE if and only if there exists a good symbolic witness representing a finite-memory weak SPE with the same payoff, and that the existence of this witness is equivalent to the non-emptiness of the fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$. In this way, we will prove Theorem 14. We will see that the lassoes $\rho_{i,v}$ of a good symbolic witness can be constructed from (p, k^*) -labeled plays for well-chosen payoffs $p \in \mathbf{P}_{k^*}(v)$.

Notice that in the second assertion of Proposition 19, we not only ask that $p_0 \in \mathbf{P}_{k^*}(v_0)$ (as in Theorem 14) but also that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$. We will come back to this observation when we will prove Theorem 14 after Example 20.

Proposition 19. *Let (\mathcal{G}, v_0) be an initialized Boolean game with prefix-independent gain functions. The following assertions are equivalent:*

1. *There exists a weak SPE with payoff p_0 in (\mathcal{G}, v_0) ;*
2. *$\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$ and $p_0 \in \mathbf{P}_{k^*}(v_0)$;*
3. *There exists a good symbolic witness \mathcal{P} that contains a lasso ρ_{0,v_0} with payoff p_0 ;*
4. *There exists a finite-memory weak SPE σ with payoff p_0 in (\mathcal{G}, v_0) such that the size of each strategy σ_i is in $\mathcal{O}(|V|^3 \cdot |\Pi|)$.*

Proof. We prove that $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 4 \Rightarrow 1$.

(1 \Rightarrow 2) Suppose that there exists a weak SPE σ with payoff p_0 in (\mathcal{G}, v_0) . To show that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$, let us prove by induction on k that

$$\text{Gain}(\langle \sigma_{\uparrow h} \rangle_v) \in \mathbf{P}_k(v) \text{ for all } hv \in \text{Hist}(v_0). \quad (1)$$

For case $k = 0$, this is true by definition of $\mathbf{P}_0(v)$. Suppose that this assertion is satisfied for an even k . Let us prove that it remains true for $k+2$ by showing that payoff $p = \text{Gain}(\langle \sigma_{\uparrow h} \rangle_v) \in \mathbf{P}_k(v)$ can be removed neither from $\mathbf{P}_k(v)$ at step $k+1$, nor from $\mathbf{P}_{k+1}(v)$ at step $k+2$.

- Payoff p cannot be removed from $\mathbf{P}_k(v)$ by the Remove operation at step $k+1$. Otherwise, if $v \in V_i$, this means that there exists $v' \in \text{Succ}(v)$ such that

$$\forall p' \in \mathbf{P}_k(v'), p_i < p'_i. \quad (2)$$

By induction hypothesis,

$$\text{Gain}(\langle \sigma_{\uparrow hv} \rangle_{v'}) \in \mathbf{P}_k(v'). \quad (3)$$

To get a contradiction, we prove that in the subgame $(\mathcal{G}_{\uparrow h}, v)$ there exists a one-shot deviating strategy σ'_i from $\sigma_{i\uparrow h}$ that is a profitable deviation for player i . We define σ'_i that only differs from $\sigma_{i\uparrow h}$ on v : $\sigma'_i(v) = v'$. Therefore we get $\text{Gain}(h\langle \sigma'_i, \sigma_{-i\uparrow h} \rangle_v) = \text{Gain}(hv\langle \sigma_{\uparrow hv} \rangle_{v'})$. It follows by (2), (3), and prefix-independence of Gain_i that $\text{Gain}_i(h\langle \sigma_{\uparrow h} \rangle_v) = p_i < p'_i = \text{Gain}_i(hv\langle \sigma_{\uparrow hv} \rangle_{v'}) = \text{Gain}(h\langle \sigma'_i, \sigma_{-i\uparrow h} \rangle_v)$. This is impossible since σ is a weak SPE.

- Payoff p cannot be removed from $\mathbf{P}_{k+1}(v)$ by the Adjust operation at step $k+2$. Otherwise, this means that there exists u such that $\mathbf{P}_{k+1}(u) = \mathbf{P}_k(u) \setminus \{p\}$ (by the Remove operation at step $k+1$) and there is no $(p, k+1)$ -labeled play with payoff p from v . However by Lemma 15, as $p \in \mathbf{P}_k(v)$, there exists a (p, k) -labeled play π with payoff p from v . This means that π visit u . Let $h'u \in \text{Hist}(v)$ such that $h'u < \pi$. Then we get a contradiction with σ being a weak SPE if we repeat the argument done before in the previous item for u and the subgame $(\mathcal{G}_{\uparrow hh'}, u)$ (instead of v and $(\mathcal{G}_{\uparrow h}, v)$).

Now that we know that $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$, it remains to prove that $p_0 \in \mathbf{P}_{k^*}(v_0)$. By (1), we have $p_0 = \text{Gain}(\langle \sigma \rangle_{v_0}) \in \mathbf{P}_{k^*}(v_0)$.

(2 \Rightarrow 3) Let us show how to build a good symbolic witness \mathcal{P} from the non-empty fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$. First recall that if $p \in \mathbf{P}_{k^*}(v)$, then by Lemma 15 there exists a (p, k^*) -labeled play with payoff p from v . Notice that such a play can be supposed to be a lasso with length at most $2 \cdot |V|^2$. Indeed it is proved in [22, Proposition 3.1] that given a play ρ , one can construct a lasso ρ' of length bounded by $2 \cdot |V|^2$ such that $\text{First}(\rho) = \text{First}(\rho')$, $\text{Occ}(\rho) = \text{Occ}(\rho')$, and $\text{Inf}(\rho) = \text{Inf}(\rho')$ (this construction eliminates some cycles of ρ in a clever way). Therefore, if ρ is a (p, k^*) -labeled play with payoff p from v , the lasso ρ' is also a (p, k^*) -labeled play with payoff p from v . The required set \mathcal{P} will be composed of some of these lassoes.

We start with $\mathcal{P} = \emptyset$. As $p_0 \in \mathbf{P}_{k^*}(v_0)$, then there exists a (p_0, k^*) -labeled lasso ρ_{0,v_0} with payoff p_0 from v_0 that we add to \mathcal{P} . For all $v, v' \in \text{Succ}^*(v_0)$ such that $v \in V_i$ and $v' \in \text{Succ}(v)$, let p' be a payoff in $\mathbf{P}_{k^*}(v')$ such that

$$p'_i = \min\{q_i \mid q \in \mathbf{P}_{k^*}(v')\}. \quad (4)$$

This payoff exists since $\mathbf{P}_{k^*}(v') \neq \emptyset$ by hypothesis. Then there exists a (p', k^*) -labeled lasso $\rho_{i,v'}$ with payoff p' from v' that we add to \mathcal{P} .

This set \mathcal{P} is a symbolic witness by construction. It remains to prove that it is good. Let $v \in \rho_{j,u}$ such that $v \in V_i$ and $\rho_{j,u} \in \mathcal{P}$. As $\rho_{j,u}$ is a (p, k^*) -labeled lasso for some payoff p , we have $p \in \mathbf{P}_{k^*}(v)$. Furthermore, as $\mathbf{P}_{k^*}(v) = \mathbf{P}_{k^*+1}(v)$ (by the fixpoint), this means that p was not removed from $\mathbf{P}_{k^*}(v)$ by the Remove operation at step k^* . In particular, by definition of the payoff p' of $\rho_{i,v'}$ (see (4)), we have $p_i \geq p'_i$, that is $\text{Gain}_i(\rho_{j,u}) \geq \text{Gain}_i(\rho_{i,v'})$. This shows that \mathcal{P} is a good symbolic witness.

(3 \Rightarrow 4) Let $\mathcal{P} = \{\rho_{i,v} \mid (i, v) \in I\}$ be a good symbolic witness that contains a lasso ρ_{0,v_0} with payoff p_0 . We define a strategy profile σ step by step by induction on the subgames of (\mathcal{G}, v_0) . We first partially build σ such that $\langle \sigma \rangle_{v_0} = \rho_{0,v_0}$. Consider next $hv v' \in \text{Hist}(v_0)$ with $v \in V_i$ such that $\langle \sigma \rangle_{|h}$ is already built but not $\langle \sigma \rangle_{|hv}$. Then we extend the definition of σ such that

$$\langle \sigma \rangle_{|hv} = \rho_{i,v'}. \quad (5)$$

Notice that $\langle \sigma \rangle_{|h}$ being already built means that there exists $h' \leq h$ and $(j, u) \in I$ such that

$$h' \langle \sigma \rangle_{|h'} = h' \rho_{j,u} = h \langle \sigma \rangle_{|h}. \quad (6)$$

Let us prove that σ is a very weak SPE (and so a weak SPE by Proposition 7). Consider the subgame $(\mathcal{G}_{\uparrow h}, v)$ and the one-shot deviating strategy σ'_i from $\sigma_{i|\uparrow h}$ such that $\sigma'_i(v) = v'$. We have to prove that

$$\text{Gain}_i(h\langle\sigma_{\uparrow h}\rangle_v) \geq \text{Gain}_i(hv\langle\sigma_{\uparrow hv}\rangle_{v'}). \quad (7)$$

By (5), (6), and prefix independence of Gain_i , we have

$$\begin{aligned} \text{Gain}_i(hv\langle\sigma_{\uparrow hv}\rangle_{v'}) &= \text{Gain}_i(\rho_{i,v'}), \\ \text{Gain}_i(h\langle\sigma_{\uparrow h}\rangle_v) &= \text{Gain}_i(\rho_{j,u}). \end{aligned}$$

Inequality (7) follows from these equalities and the fact that \mathcal{P} is a good symbolic witness (see Figure 2).

Notice that σ has payoff p_0 by construction (ρ_{0,v_0} has payoff p_0). It remains to show that σ is finite-memory. Having (j, u) in memory (the last deviating player j and the vertex u where he moved), the Moore machines \mathcal{M}_i , $i \in \Pi$, representing each strategy σ_i , have to produce together the lasso $\rho_{j,u}$ of length bounded by $2 \cdot |V|^2$. As $|\mathcal{P}| = |I| \leq |V| \cdot |\Pi| + 1$, the size of each σ_i is in $\mathcal{O}(|\Pi| \cdot |V|^3)$.

(4 \Rightarrow 1) This implication is obvious. \square

Example 20. We consider again the running example. Thanks to Theorem 14, payoff $p_0 = (0, 1)$ is the only possible payoff for a weak SPE in (\mathcal{G}, v_0) since $\mathbf{P}_{k^*}(v_0) = \{(0, 1)\}$ (see Example 13). Let us illustrate the constructions of the proof of Proposition 19. We build a good symbolic witness \mathcal{P} as follows. First, we choose a (p_0, k^*) -labeled lasso ρ_{0,v_0} with payoff p_0 from v_0 : $\rho_{0,v_0} = v_0v_1v_2v_3^\omega$ (we could also have chosen $v_0v_4v_5^\omega$). Then let $v = v_1, v' = v_2 \in \text{Succ}^*(v_0)$ such that $v \in V_1$ and $v' \in \text{Succ}(v)$. As $\mathbf{P}_{k^*}(v_2) = \{(1, 0), (0, 1)\}$, by (4), we choose a (p', k^*) -labeled lasso ρ_{1,v_2} with payoff $p' = (0, 1)$ from v_2 , for instance, $\rho_{1,v_2} = v_2v_3^\omega$. The whole set \mathcal{P} is depicted in Table 3. From \mathcal{P} , we get the weak SPE of Example 8 depicted in Figure 1.

We are now able to prove Theorem 19.

Proof of Theorem 14. If there exists a weak SPE with payoff p_0 in (\mathcal{G}, v_0) , then $p_0 \in \mathbf{P}_{k^*}(v_0)$ by Proposition 19. It is proved in [18] that each initialized game with prefix-independent gain functions always has a weak SPE. Therefore $\mathbf{P}_{k^*}(v) \neq \emptyset$ for each $v \in \text{Succ}^*(v_0)$ by Proposition 19. The latter assertion together with the assertion that $p_0 \in \mathbf{P}_{k^*}(v_0)$ implies the existence of a weak SPE with payoff p_0 in (\mathcal{G}, v_0) . \square

Proposition 19 highlights a property that will be very useful in Section 4 (the equivalence between the existence of a weak SPE and the existence of a good symbolic witness). It also shows that when the gain functions are prefix-independent, if there exists a weak SPE with a given payoff, then there always exists one with the same payoff but with strategies of polynomial size. We prove in Section 4.2.3 that for Reachability and Safety objectives which are not prefix-independent, we have the same result however with strategies of exponential size.

Corollary 21. *Let (\mathcal{G}, v_0) be an initialized Boolean game. There exists a weak SPE in (\mathcal{G}, v_0) if and only if there exists a finite-memory weak SPE σ with the same payoff. Moreover, the size of each strategy σ_i is*

- *in $\mathcal{O}(|V|^3 \cdot |\Pi|)$ for Büchi, Co-Büchi, Parity, Explicit Muller, Muller, Rabin, and Streett objectives,*
- *in $\mathcal{O}(|V|^3 \cdot |\Pi| \cdot 2^{3 \cdot |\Pi|})$ for Reachability and Safety objectives.*

4. Complexity classes of the constraint problem

In this section, we study the complexity classes of the constraint problem for Boolean games with classical ω -regular objectives, except the case of Explicit Muller objectives that is postponed to Section 5.3 (see Table 1). The concept of good symbolic witness is essential in this study.

4.1. NP-completeness

We first prove that the constraint problem for co-Büchi, Parity, Muller, Rabin, and Streett objectives is NP-complete, and that it is in NP for Büchi objectives.

Theorem 22. *The constraint problem for Boolean games with co-Büchi, Parity, Muller, Rabin, and Streett objectives is NP-complete. It is in NP for Büchi objectives.*

Proof. We begin with the NP-easiness. The objectives considered in Theorem 22 are prefix-independent. We can thus apply Proposition 19. Given thresholds $x, y \in \{0, 1\}^{|\Pi|}$, there exists a weak SPE in (\mathcal{G}, v_0) with payoff p such that $x \leq p \leq y$ if and only if there exists a good symbolic witness \mathcal{P} that contains a lasso ρ_{0, v_0} with payoff p . Hence a nondeterministic polynomial algorithm works as follows: guess a set \mathcal{P} composed of at most $|\Pi| \cdot |V| + 1$

lassoes of length at most $2 \cdot |V|^2$ and check that \mathcal{P} is a good symbolic witness that contains a lasso ρ_{0,v_0} with payoff p such that $x \leq p \leq y$. Clearly checking that \mathcal{P} is a symbolic witness can be done in polynomial time. Checking that it is good requires to compute the payoffs of its lassoes and to compare them. This can also be done in polynomial time for Büchi, co-Büchi, Parity, Muller, Rabin, and Streett objectives.

We now proceed to the NP-hardness. It is obtained thanks to a polynomial reduction from SAT. In [23] is provided a polynomial reduction from SAT to the constraint problem for NEs in Boolean games with co-Büchi objectives. Due to the structure of the game constructed in this approach, the same reduction holds for the constraint problem for weak SPEs. As co-Büchi objectives can be polynomially translated into Parity, Muller, Rabin, and Streett objectives (see [24]), the constraint problem for Boolean games with those objectives is also NP-hard. \square

4.2. PSPACE-completeness

In this section, we show that the constraint problem for Reachability and Safety objectives is PSPACE-complete.

Theorem 23. *The constraint problem for Boolean games with Reachability and Safety objectives is PSPACE-complete.*

Recall that weak SPEs and SPEs are equivalent notions for Reachability objectives (Proposition 9). It follows from Theorem 23 that the constraint problem for SPEs (instead of weak SPEs) for Boolean games with Reachability objectives is PSPACE-complete. We will see later (in Section 4.2.3, from the proof of Theorem 23) that the constraint problem for SPEs is also PSPACE-complete for Safety objectives.

Corollary 24. *The constraint problem for SPEs in Boolean games with Reachability and Safety objectives is PSPACE-complete.*

We detail the proof of Theorem 23 in the next two sections for Reachability objectives, and we also show how to adapt it for Safety objectives. To get the PSPACE-easiness, we transform the Boolean game (\mathcal{G}, v_0) with Reachability objectives (which are not prefix-independent) into a Boolean game (\mathcal{G}', v'_0) with Büchi objectives. In this way, it is possible to use the concept of good symbolic witness as done before in Section 4.1. Even if the size of the game (\mathcal{G}', v'_0) is exponential in the size of the initial game (\mathcal{G}, v_0) ,

we manage to get a PSPACE-membership thanks to the classical complexity result $\text{PSPACE} = \text{APTIME}$. The PSPACE-hardness is obtained with a polynomial reduction from QBF. The reduction is more involved than the one in Theorem 22. Indeed the reduction for NP-hardness already works for NEs whereas the reduction for PSPACE-hardness really exploits the subgame perfect aspects.

4.2.1. PSPACE-easiness

We here prove that the constraint problem is in PSPACE for Reachability objectives, and we then explain how to adapt the proof for Safety objectives.

Proposition 25. *The constraint problem for Boolean games with Reachability objectives is in PSPACE.*

Proof. First, we transform the Boolean game (\mathcal{G}, v_0) with Reachability objectives into a Boolean game (\mathcal{G}', v'_0) with Büchi objectives. This construction is classical, it stores inside the vertices the set of players who have already satisfied their objective. Suppose that in (\mathcal{G}, v_0) , each player i aims at reaching a vertex of F_i , then we build $\mathcal{G}' = (\Pi, V', E', (V'_i)_{i \in \Pi}, (\text{Gain}'_i)_{i \in \Pi})$ such that

- $V' = V \times 2^\Pi$;
- $((v, I), (u, I')) \in E'$ if and only if $(v, u) \in E$ and $I' = I \cup \{i \in \Pi \mid u \in F_i\}$;
- $(v, I) \in V'_i$ if and only if $v \in V_i$;
- Gain'_i corresponds to the Büchi objective $F'_i = \{(v, I) \mid v \in V, i \in I\}$; and
- $v'_0 = (v_0, I_0)$ is the initial vertex such that $I_0 = \{i \in \Pi \mid v_0 \in F_i\}$.

There is a one-to-one correspondence between plays $\rho = v_1 v_2 \dots v_k \dots$ in \mathcal{G} from v_1 and $\rho' = (v_1, I_1)(v_2, I_2) \dots (v_k, I_k) \dots$ in \mathcal{G}' from (v_1, I_1) , with the important property that

$$I_k \subseteq I_{k+1} \text{ for all } k \geq 1. \quad (8)$$

In particular, there exists a weak SPE with payoff p_0 in (\mathcal{G}, v_0) if and only if there is one in (\mathcal{G}', v'_0) if and only if there is a good symbolic witness \mathcal{P} containing a lasso ρ_{0, v'_0} with payoff p_0 in (\mathcal{G}', v'_0) (by Proposition 19).

Second, let us study the lengths of the lassos $\rho_{i,v'}$, with $i \in \Pi \cup \{0\}$, $v' \in V'$, of \mathcal{P} . There exists a better bound than the bound $2 \cdot |V'|^2 = 2 \cdot |V|^2 \cdot 2^{2 \cdot |\Pi|}$ of Definition 16: each lasso $\rho_{i,v'}$ of \mathcal{P} has a polynomial length bounded by

$$(|\Pi| + 1) \cdot |V|. \quad (9)$$

Indeed recall that these lassoes are constructed from some (p, k^*) -labeled plays ρ with payoff p from v' (see the proof of implication $(2 \Rightarrow 3)$ of Proposition 19). We adapt this construction as follows. Consider such a play $\rho = (v_1, I_1)(v_2, I_2) \dots (v_k, I_k) \dots$ from v' . By (8), there exists $I \subseteq \Pi$ and $k \in \mathbb{N}$ such that for all $k' \geq k$, $I_{k'} = I$. Hence from ρ , we can construct a lasso ρ' of length bounded by (9) such that $\text{First}(\rho') = \text{First}(\rho)$, $\text{Occ}(\rho') \subseteq \text{Occ}(\rho)$, and $\text{Gain}(\rho') = \text{Gain}(\rho)$,

- by eliminating all cycles in the history $(v_1, I_1)(v_2, I_2) \dots (v_{k-1}, I_{k-1})$ (leading to a history of length at most $|\Pi| \cdot |V|$), and
- by detecting in the play $(v_k, I)(v_{k+1}, I) \dots$ the first repeated vertex $(v_{k'}, I) = (v_{k'+\ell+1}, I)$ and replacing this play by the lasso

$$(v_k, I)(v_{k+1}, I) \dots ((v_{k'}, I) \dots (v_{k'+\ell}, I))^\omega$$

of length at most $|V|$.

In this way, if ρ is a (p, k^*) -labeled play with payoff p from v' , then the constructed lasso ρ' is also a (p, k^*) -labeled play with payoff p from v' .

Third we prove PSPACE-membership of the constraint problem by proving that it is in APTIME. Given the game (\mathcal{G}', v'_0) and two thresholds $x, y \in \{0, 1\}^{|\Pi|}$, the alternating Turing machine works as follows. Existential and universal states (respectively controlled by player \vee and player \wedge) alternate along an execution of the machine. Player \vee proposes a lasso $\rho_{j,w'}$ of length bounded by $(|\Pi| + 1) \cdot |V|$ (in the initial state, he proposes a lasso ρ_{0,v'_0}). Then player \wedge chooses a vertex $w' \in V'_i$ of $\rho_{j,w'}$ and proposes to move to v' such that $(w', v') \in E'$. Player \vee reacts by proposing a lasso $\rho_{i,v'}$ of length bounded by $(|\Pi| + 1) \cdot |V|$, and so on. The execution stops after

$$2 \cdot |\Pi|^2 \cdot |V| + 1 \text{ turns.} \quad (10)$$

Such an execution is accepting if:

- for the payoff p_0 of the initial lasso ρ_{0,v'_0} , we have $x \leq p_0 \leq y$;

- for each lasso $\rho_{j,w'}$ proposed by player \vee , for the corresponding move $(w', v') \in E'$ with $w' \in V'_i$ made by player \wedge , and the answer $\rho_{i,v'}$ of player \vee , we have $\text{Gain}_i(\rho_{j,w'}) \geq \text{Gain}_i(\rho_{i,v'})$.

The intuition is that if there exists in (\mathcal{G}', v'_0) a good symbolic witness \mathcal{P} containing a lasso ρ_{0,v'_0} with payoff p_0 such that $x \leq p_0 \leq y$, then player \vee will play with the lassoes of \mathcal{P} according to Definition 17. Notice that along an execution of the Turing machine, player \wedge has no interest to choose twice the same pair (i, v') since player \vee will react with the same lasso $\rho_{i,v'}$. Remembering property (8), the maximum number of times that player \wedge has to play is

$$|\Pi|^2 \cdot |V|. \quad (11)$$

Indeed for a fixed $I \subseteq \Pi$, player \wedge can choose at most $|\Pi| \cdot |V|$ different pairs (i, v') with v' of the form (v, I) , and the size of I can only increase. This explains the number of turns of any execution of the machine (see (10)): an initial lasso proposed by player \vee followed by $|\Pi|^2 \cdot |V|$ alternations between moves of both players \vee and \wedge .

Checking whether an execution is accepting is done in polynomial time since player \vee proposes lassoes of polynomial size by (9), there is a polynomial numbers of turns by (10), and computing and comparing payoffs of lassoes is done in polynomial time. So the constraint problem is in $\text{APTIME} = \text{PSPACE}$. \square

The constraint problem for Safety objectives is solved similarly by transforming the given Boolean game into one with co-Büchi (instead of Büchi) objectives.

Proposition 26. *The constraint problem for Boolean games with Safety objectives is in PSPACE.*

4.2.2. PSPACE-hardness

We now prove that the constraint problem is PSPACE-hard for Reachability objectives, and we then show how to adapt the proof for Safety objectives.

Proposition 27. *The constraint problem for Boolean games with Reachability objectives is PSPACE-hard.*

To prove this proposition, we give a polynomial reduction from the QBF problem that is PSPACE-complete. This problem is to decide whether a fully quantified Boolean formula ψ is true. The formula ψ can be assumed to be in prenex Conjunctive Normal Form (CNF) $Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ such that the quantifiers are alternating existential and universal quantifiers ($Q_1 = \exists, Q_2 = \forall, Q_3 = \exists, \dots$), $X = \{x_1, x_2, \dots, x_m\}$ is the set of quantified variables, and $\phi(X) = C_1 \wedge \dots \wedge C_n$ is an unquantified Boolean formula over X equal to the conjunction of the clauses C_1, \dots, C_n .

Such a formula ψ is true if there exists a value of x_1 such that for all values of x_2 , there exists a value of $x_3 \dots$, such that the resulting valuation ν of all variables of X evaluates $\phi(X)$ to true. Formally, for each odd (resp. even) k , $1 \leq k \leq m$, let us denote by $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$ (resp. $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$) a valuation of variable x_k given a valuation of previous variables x_1, \dots, x_{k-1} ⁹. Given these sequences $f = f_1, f_3, \dots$ and $g = g_2, g_4, \dots$, let us denote by $\nu = \nu_{(f,g)}$ the valuation of all variables of X such that $\nu(x_1) = f_1, \nu(x_2) = g_2(\nu(x_1)), \nu(x_3) = f_3(\nu(x_1)\nu(x_2)), \dots$. Then

$$\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X) \text{ is true} \\ \text{if and only if}$$

there exist $f = f_1, f_3, \dots$ such that for all $g = g_2, g_4, \dots$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

Proof of Proposition 27. Let us detail a polynomial reduction from the QBF problem to the constraint problem for Boolean games with Reachability objectives. Let $\psi = Q_1x_1Q_2x_2\dots Q_mx_m\phi(X)$ with $\phi(X) = C_1 \wedge \dots \wedge C_n$ be a fully quantified Boolean formula in prenex Conjunctive Normal Form. We build the following Boolean game $\mathcal{G}_\psi = (\Pi, V, (V_i)_{i \in \Pi}, E, (\text{Gain}_i)_{i \in \Pi})$ (see Figure 3):

- the set V of vertices:
 - for each variable $x_k \in X$ under quantifier Q_k , there exist vertices $x_k, \neg x_k$ and q_k ;
 - for each clause C_k , there exist vertices c_k and t_k ;
 - there exists an additional vertex t_{n+1} ;

⁹Notice that $f_1 : \emptyset \rightarrow \{0, 1\}$.

- the set E of edges:
 - from each vertex q_k there exist an edge to x_k and an edge to $\neg x_k$;
 - from each vertex x_k and $\neg x_k$, there exists an edge to q_{k+1} , except for $k = m$ where this edge is to c_1 ;
 - from each vertex c_k , there exist an edge to t_k and an edge to c_{k+1} , except for $k = n$ where there exist an edge to t_n and an edge to t_{n+1} ;
 - there exists a loop on each t_k ;
- the set Π of $n + 2$ players:
 - each player i , $1 \leq i \leq n$, owns vertex c_i ;
 - player $n + 1$ (resp. $n + 2$) is the player who owns the vertices q_i for each existential (resp. universal) quantifier Q_i ;
 - as all other vertices have only one outgoing edge, it does not matter which player owns them;
- each function Gain_i is associated with the objective of visiting the set F_i defined as follows:
 - for all i , $1 \leq i \leq n$, $F_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_i\}$;
 - $F_{n+1} = \{t_{n+1}\}$;
 - $F_{n+2} = \{t_1, \dots, t_n\}$.

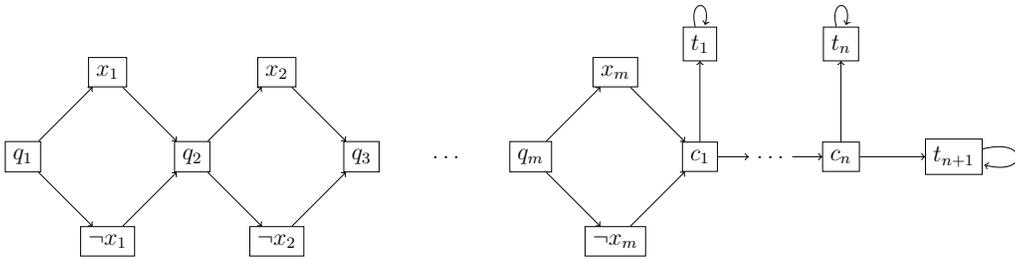


Figure 3: Reduction from the formula ψ to the Boolean game \mathcal{G}_ψ

Remark 28. (1) Notice that a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, as presented above, can be translated into a strategy σ_{n+1} of player $n+1$ in the initialized game (\mathcal{G}_ψ, q_1) , and conversely. Similarly, a sequence g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ is nothing else than a strategy σ_{n+2} of player $n+2$. (2) Notice also that if ρ is a play in (\mathcal{G}_ψ, q_1) , then $\text{Gain}_{n+1}(\rho) = 1$ if and only if $\text{Gain}_{n+2}(\rho) = 0$. Moreover, suppose that ρ visits t_{n+1} , then for all i , $1 \leq i \leq n$, $\text{Gain}_i(\rho) = 1$ if and only if for all i , $1 \leq i \leq n$, ρ visits a vertex that is a literal of C_i if and only if there is a valuation of all variables of X that evaluates $\phi(X)$ to true.

The game \mathcal{G}_ψ can be constructed from ψ in polynomial time. Let us now show that ψ is true if and only if there exists a weak SPE in (\mathcal{G}_ψ, q_1) with a payoff $p \geq (0, \dots, 0, 1, 0)$ (that is, such that $p_{n+1} = 1$).

(\Rightarrow) Suppose that ψ is true. Then there exists a sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, such that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. We define a strategy profile σ as follows:

- for player $n+1$, his strategy σ_{n+1} is the strategy corresponding to the sequence f (by Remark 28);
- for player $n+2$, his strategy is an arbitrary strategy σ_{n+2} ; we denote by g the corresponding sequence $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$ (by Remark 28);
- for each player i , $1 \leq i \leq n$,
 - if $hv \in \text{Hist}_i(q_1)$ with $v = c_i$, is consistent with σ_{n+1} , then $\sigma_i(hv) = c_{i+1}$ if $i \neq n$ and t_{n+1} otherwise
 - else $\sigma_i(hv) = t_i$.

Let us prove that σ is a weak SPE, that is, for each history $hv \in \text{Hist}(q_1)$, there is no one-shot deviating strategy in the subgame $(\mathcal{G}_{\psi|h}, v)$ that is profitable to the player who owns vertex v (by Proposition 7). This is clearly true for all $v = t_i$, $1 \leq i \leq n+1$, since t_i has only one outgoing edge. For the other vertices v , we study two cases:

- hv is consistent with σ_{n+1} : First notice that $\text{Gain}(\rho) = (1, 1, \dots, 1, 0)$ with $\rho = h\langle \sigma_{|h} \rangle_v$. Indeed by hypothesis, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true. Hence by Remark 28, the play ρ visits a vertex of F_i

for all i , $1 \leq i \leq n$, and by definition of σ , ρ eventually loops on t_{n+1} . Second, as σ_{n+2} is arbitrary in the definition of σ , using another strategy σ'_{n+2} in place of σ_{n+2} will lead to a play ρ' such that $\text{Gain}(\rho') = \text{Gain}(\rho) = (1, 1, \dots, 1, 0)$.

Now if hv is consistent with σ_{n+1} , it is maybe not consistent with σ_{n+2} , but with another arbitrary strategy σ'_{n+2} , and thus $\text{Gain}(h\langle\sigma_{\uparrow h}\rangle_v) = (1, 1, \dots, 1, 0)$ as explained previously. Thus only player $n + 2$ has an incentive to deviate in the subgame $(\mathcal{G}_{\psi \uparrow h}, v)$ to increase his gain. Nevertheless, using another strategy σ''_{n+2} will not change his gain (again by the same argument).

- hv is not consistent with σ_{n+1} : Suppose that $v = c_k$. Then by definition of σ , the play $h\langle\sigma_{\uparrow h}\rangle_v$ eventually loops on t_k leading to a gain of 1 for player k . This player has thus no incentive to deviate with a one-shot deviation in the subgame $(\mathcal{G}_{\psi \uparrow h}, v)$.

Suppose that $v = q_k$. Then by definition of σ , the play $\rho = h\langle\sigma_{\uparrow h}\rangle_v$ eventually loops on t_1 . It follows that $\text{Gain}_{n+1}(\rho) = 0$ and $\text{Gain}_{n+2}(\rho) = 1$. As we only have to consider one-shot deviating strategies, if $q_k \in V_{n+2}$, player $n + 2$ has no incentive to deviate, and if $q_k \in V_{n+1}$, player $n + 1$ could try to use a one-shot deviating strategy, however the resulting play still eventually loops on t_1 .

This proves that σ is a weak SPE. Its payoff is equal to $p = (1, 1, \dots, 1, 0)$ as explained previously. Therefore it satisfies the constraint $p \geq (0, \dots, 0, 1, 0)$.

(\Leftarrow) Suppose that there exists a weak SPE σ in (\mathcal{G}_ψ, q_1) with outcome ρ and payoff $\text{Gain}(\rho) \geq (0, \dots, 0, 1, 0)$, that is, $\text{Gain}_{n+1}(\rho) = 1$. By Remark 28, it follows that $\text{Gain}_{n+2}(\rho) = 0$. We have to prove that ψ is true. To this end, consider the sequence f of functions $f_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k odd, $1 \leq k \leq m$, that corresponds to strategy σ_{n+1} of player $n + 1$ by Remark 28. Let us show that for all sequences g of functions $g_k : \{0, 1\}^{k-1} \rightarrow \{0, 1\}$, with k even, $1 \leq k \leq m$, the valuation $\nu_{f,g}$ evaluates $\phi(X)$ to true.

By contradiction assume that it is not the case for some sequence g' and consider the related strategy σ'_{n+2} of player $n + 2$ by Remark 28. Notice that σ'_{n+2} is a finitely deviating strategy. Let us consider the outcome ρ' of the strategy profile $(\sigma'_{n+2}, \sigma_{-(n+2)})$ from q_1 . As $\text{Gain}_{n+2}(\rho) = 0$, we must have $\text{Gain}_{n+2}(\rho') = 0$, otherwise σ'_{n+2} is a profitable deviation for player $n + 2$ whereas σ is a weak SPE. It follows that $\text{Gain}_{n+1}(\rho') = 1$ by Remark 28, that is, ρ' eventually loops on t_{n+1} .

Now recall that the valuation $\nu_{f,g'}$ evaluates $\phi(X)$ to false, which means that it evaluates some clause C_k of $\phi(X)$ to false. Consider the history $hc_k < \rho'$. As strategy σ'_{n+2} only acts on the left part of the underlying graph of \mathcal{G}_ψ , we have $\rho' = \langle \sigma'_{n+2}, \sigma_{-(n+2)} \rangle_{q_1} = h \langle \sigma_{\uparrow h} \rangle_{c_k}$. In the subgame $(\mathcal{G}_{\psi \uparrow h}, c_k)$, the outcome of $\sigma_{\uparrow h}$ gives a gain of 0 to player k because $\rho' = h \langle \sigma_{\uparrow h} \rangle_{c_k}$ does not visit t_k and $\nu_{f,g'}$ evaluates C_k to false. In this subgame, player k has thus a profitable one-shot deviation: to move to t_k . It follows that σ is not a weak SPE which is impossible. Therefore ψ is true. \square

For Boolean games with Safety objectives, we can use the same reduction and the same kind of arguments as for Reachability objectives.

Proposition 29. *The constraint problem for Boolean games with Safety objectives is PSPACE-hard.*

Proof. Given a fully quantified Boolean formula ψ , we construct the same game as in the proof of Proposition 27 (see Figure 3), except that each player i , $1 \leq i \leq n+2$, aims at avoiding the set F'_i (instead of visiting the set F_i) defined as follows:

- for all i , $1 \leq i \leq n$, $F'_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_{n+1}\}$;
- $F'_{n+1} = \{t_1, t_2, \dots, t_n\}$;
- $F'_{n+2} = \{t_{n+1}\}$.

Recall how the sets F_i were defined: $F_{n+1} = \{t_{n+1}\}$, $F_{n+2} = \{t_1, t_2, \dots, t_n\}$, and for all i , $1 \leq i \leq n$, $F_i = \{\ell \in V \mid \ell \text{ is a literal of clause } C_i\} \cup \{t_i\}$. Hence we have a clear duality for players $n+1$ and $n+2$: a play ρ visits F_{n+1} (resp. F_{n+2}) if and only if ρ avoids F'_{n+1} (resp. F'_{n+2}). This is not the case for the others players, but one can check that the proof works in the same way as for Boolean games with Reachability objectives. \square

4.2.3. Proof of Corollaries 21 and 24

We conclude Section 4 with the proof of two previous corollaries. We begin with Corollary 24 stating that the constraint problem for SPEs (instead of weak SPEs) in Boolean games with Reachability and Safety objectives is PSPACE-complete.

Proof of Corollary 24. As weak SPEs and SPEs are equivalent notions for Reachability objectives (Proposition 9), by Theorem 23, the constraint problem for SPEs for Boolean games with Reachability objectives is PSPACE-complete.

We need to use other arguments for the case of Safety objectives. The reduction from QBF proposed in the proof of Proposition 29 uses the game \mathcal{G}_ψ of Figure 3. Due to the structure of the underlying graph, all weak SPEs of \mathcal{G}_ψ are SPEs since any deviating strategy from a given strategy is necessarily finitely deviating. This shows that the constraint problem for SPEs is PSPACE-hard for Safety objectives. It is proved in [25] that this problem is in PSPACE. \square

We now turn to Corollary 21 that states that if there exists a weak SPE in a Boolean game, then there exists a finite-memory weak SPE with the same payoff such that the strategy memory sizes are polynomial for all objectives, except for Reachability and Safety objectives where the sizes are exponential.

Proof of Corollary 21. For the objectives that are prefix-independent, this is an immediate consequence of Proposition 19 with strategy sizes in $\mathcal{O}(|V|^3 \cdot |\Pi|)$. For Reachability and Safety objectives, we need to transform the Boolean game with Reachability (resp. Safety) objectives into a Boolean game with Büchi (resp. Co-Büchi) objectives, as done in the proof of Proposition 25 (resp. Proposition 26). Recall that the set of vertices of the latter game is equal to $V' = V \times 2^{|\Pi|}$. We can thus again apply Proposition 19 and obtain strategy sizes in $\mathcal{O}(|V'|^3 \cdot |\Pi|) = \mathcal{O}(|V|^3 \cdot |\Pi| \cdot 2^{3|\Pi|})$. \square

5. Fixed parameter tractability

In this section, we show that the constraint problem is P -complete for Explicit Muller objectives, that it is fixed parameter tractable for the other classical ω -regular objectives, and that it becomes polynomial when the number of players is fixed.

These results do not rely on the concept of good symbolic witness (as in Section 4) but rather on the following algorithm based on Theorem 14 to solve the constraint problem. Given a Boolean game (\mathcal{G}, v_0) and thresholds $x, y \in \{0, 1\}^{|\Pi|}$,

- Compute the initial sets $\mathbf{P}_0(v)$, $v \in V$, and repeat the Remove-Adjust procedure (see Definition 11) until reaching the fixpoint $\mathbf{P}_{k^*}(v)$, $v \in V$,

- Then check whether $\mathbf{P}_{k^*}(v) \neq \emptyset$ for all $v \in \text{Succ}^*(v_0)$ and whether there exists a payoff $p \in \mathbf{P}_{k^*}(v_0)$ such that $x \leq p \leq y$.

We call this algorithm the *decision algorithm* and its first part computing the fixpoint the *fixpoint algorithm*.

5.1. Complexity of the decision algorithm

We here study the time complexity of the decision algorithm. We express it in terms of three parameters:

- *init*: the complexity of computing $\mathbf{P}_0(v)$ for some given vertex v ,
- $m = \max_{v \in V} |\mathbf{P}_0(v)|$: the maximum number of payoffs in the sets $\mathbf{P}_0(v)$, $v \in V$,
- *path*: the complexity of determining whether there exists a play with a given payoff p from a given vertex v . (This test is required in both the computation of $\mathbf{P}_0(v)$ and the Adjust operation.)

Lemma 30. *The time complexity of the decision algorithm is in $\mathcal{O}(m^3 \cdot |V| \cdot |\Pi| \cdot \text{init} \cdot \text{path} \cdot (|V| + |E|))$.*

Expressing the complexity in this way will be useful in Section 5.3 (dedicated to Boolean games with Explicit Muller objectives and to the fixed parameter tractability of the constraint problem). We do not claim that the given complexity is the tightest one but this is enough for our purpose.

The gain function Gain_i takes its values in $\{0, 1\}$ for all i , hence m is bounded by $2^{|\Pi|}$. Moreover by definition of $\mathbf{P}_0(v)$, *init* is in $\mathcal{O}(2^{|\Pi|} \cdot \text{path})$. For Boolean games with Explicit Muller objectives, we will provide a polynomial bound for both m and *init* (see Section 5.3). Next Section 5.2 is devoted to the study of *path*.

Proof of Lemma 30. We suppose that the graph G is given as $|V|$ lists of successors and the sets $\mathbf{P}_k(v)$ are given as $|V|$ lists of payoffs. The comparison between two payoffs is in $\mathcal{O}(|\Pi|)$ time.

We first study the time complexity of the fixpoint algorithm.

- The first step of the algorithm is to compute the sets $\mathbf{P}_0(v)$ for all v , that takes $\mathcal{O}(|V| \cdot \text{init})$ time.

- Then the algorithm repeats the Remove-Adjust procedure until reaching a fixpoint. There are at most $m \cdot |V|$ repetitions of this procedure since it removes at least one payoff from some $\mathbf{P}_k(v)$.
- The computation of one Remove operation is in $\mathcal{O}((|V| + |E|) \cdot m^2 \cdot |\Pi|)$ time. Indeed we potentially have to consider all the vertices v and their successors v' to check whether there exists $p \in \mathbf{P}_k(v)$ such that $p_i < p'_i$ for all $p' \in \mathbf{P}_k(v')$ (with $v \in V_i$). This takes $\mathcal{O}((|V| + |E|) \cdot m^2 \cdot |\Pi|)$ time. If the check is positive, we have to remove p from $\mathbf{P}_k(v)$ that takes $\mathcal{O}(m \cdot |\Pi|)$ time.
- The computation of one Adjust operation is in $\mathcal{O}((|V| + |E|) \cdot m \cdot |\Pi| \cdot path)$ time. Indeed if p is the payoff removed by the Remove operation just before, for all u such that $p \in \mathbf{P}_k(u)$, we have to check whether there exists a (p, k) -labeled play with payoff p from u and remove p from $\mathbf{P}_k(u)$ in case of non existence of such a play. This can be done by computing a graph G' from G by restricting V to the vertices u such that $p \in \mathbf{P}_k(u)$ (in $\mathcal{O}(|V| \cdot m \cdot |\Pi| + (|V| + |E|))$ time), and for each u by checking in G' whether there exists a play with payoff p from u and then remove p from $\mathbf{P}_k(u)$ in case of non existence (in $\mathcal{O}(|V| \cdot (path + m \cdot |\Pi|))$ time).

Therefore the total time complexity of the fixpoint algorithm is in $\mathcal{O}(|V| \cdot init + m \cdot |V| \cdot [(|V| + |E|) \cdot m^2 \cdot |\Pi| + (|V| + |E|) \cdot m \cdot |\Pi| \cdot path]) = \mathcal{O}(|V| \cdot init + m^3 \cdot |V| \cdot |\Pi| \cdot path \cdot (|V| + |E|))$. This is bounded by $\mathcal{O}(m^3 \cdot |V| \cdot |\Pi| \cdot init \cdot path \cdot (|V| + |E|))$.

To get the time complexity of the decision algorithm, it remains to add the time complexity to test if there exists a payoff $p \in \mathbf{P}_{k^*}(v_0)$ such that $x \leq p \leq y$. This is done in $\mathcal{O}(m \cdot |\Pi|)$ time. The overall complexity of the decision algorithm is thus in $\mathcal{O}(m^3 \cdot |V| \cdot |\Pi| \cdot init \cdot path \cdot (|V| + |E|))$ time as announced. \square

5.2. Existence of a play with a given payoff

The purpose of this section is to prove the next lemma stating the complexity $path$ for all kinds of ω -regular objectives.

Lemma 31. *Let \mathcal{G} be a Boolean game. Let $p \in \{0, 1\}^{|\Pi|}$ and $v \in V$.*

1. *Determining whether there exists a play with payoff p from v is*

- in polynomial time for Büchi, co-Büchi, Explicit Muller, and Parity objectives,
 - in $\mathcal{O}(2^{|\Pi|}(|V| + |E|))$ time for Reachability and Safety objectives, and
 - in $\mathcal{O}(2^L \cdot M + (L^L \cdot |V|)^5)$ time for Rabin, Streett, and Muller objectives, where $L = 2^\ell$ and
 - $\ell = \sum_{i=1}^{|\Pi|} 2 \cdot k_i$ and $M = \mathcal{O}(\sum_{i=1}^{|\Pi|} 2 \cdot k_i)$ such that for each player $i \in \Pi$, k_i is the number of his pairs $(G_j^i, R_j^i)_{1 \leq j \leq k_i}$ in the case of Rabin and Streett objectives, and
 - $\ell = \sum_{i=1}^{|\Pi|} d_i$ and $M = \mathcal{O}(\sum_{i=1}^{|\Pi|} |\mathcal{F}_i| \cdot d_i)$ such that for each player $i \in \Pi$, d_i (resp. $|\mathcal{F}_i|$), is the number of his colors (the size of his family of subsets of colors) in the case of Muller objectives.
2. When the number $|\Pi|$ of players is fixed, for all these kinds of objectives, the existence of a play with payoff p from v can be solved in polynomial time.

The general approach to prove this lemma is the following one. A play with payoff p from v in a Boolean game \mathcal{G} is a play satisfying an objective Ω equal to the conjunction of objectives Obj_i (when $p_i = 1$) and of objectives $V^\omega \setminus \text{Obj}_i$ (when $p_i = 0$). It is nothing else than an infinite path in the underlying graph $G = (V, E)$ satisfying some particular ω -regular objective Ω . The existence of such paths is a well studied problem; we gather in the next proposition the known results that we need for proving Lemma 31. Recall that a *Generalized Reachability* (resp. *Generalized Büchi*) objective Ω is a conjunction of Reachability (resp. Büchi) objectives. Moreover, an objective Ω equal to a Boolean combination of Büchi objectives, called a *BC Büchi objective*, is defined as follows. Let F_1, \dots, F_ℓ be ℓ subsets of V , and ϕ be a Boolean formula over variables f_1, \dots, f_ℓ . We say that an infinite path ρ in G satisfies $(\phi, F_1, \dots, F_\ell)$ if the truth assignment

$$f_i = 1 \text{ if and only if } \text{Inf}(\rho) \cap F_i \neq \emptyset, \text{ and } f_i = 0 \text{ otherwise}$$

satisfies ϕ . All operators \vee, \wedge, \neg are allowed in a BC Büchi objective. However we denote by $|\phi|$ the size of ϕ equal to the number of disjunctions and conjunctions inside ϕ , and we say that the BC Büchi objective $(\phi, F_1, \dots, F_\ell)$ is of size $|\phi|$ and with ℓ variables.

Proposition 32. *Let $G = (V, E)$ be a graph, $v \in V$ be one of its vertices, and $\Omega \subseteq V^\omega$ be an objective. Then deciding the existence of an infinite path from v in G that satisfies Ω is*

- *in polynomial time for Ω equal to either a Streett objective, or an Explicit Muller objective, or the opposite of an Explicit Muller objective, or a conjunction of a Generalized Büchi objective and a co-Büchi objective,*
- *in $\mathcal{O}(2^\ell(|V| + |E|))$ time for Ω equal to a conjunction of a Generalized Reachability objective and a Safety objective, where ℓ is the number of reachability objectives,*
- *in $\mathcal{O}(2^L \cdot |\phi| + (L^L \cdot |V|)^5)$ time for a BC Büchi objective $\Omega = (\phi, F_1, \dots, F_\ell)$, where $L = 2^\ell$.*

Proof. Let Ω be an objective. If it is a Streett objective, then the result is proved in [26].

For the other objectives, we use known results about *two-player zero-sum games*, where player A aims at satisfying a certain objective Ω whereas player B tries to prevent him to satisfy it. A classical problem is to decide whether player A has a winning strategy that allows him to satisfy Ω against any strategy of player B , see for instance [12, 24]. When player A is the only one to play, the existence of a winning strategy for him is equivalent to the existence of a path satisfying Ω (see [12, Section 3.1]). This is exactly the problem that we want to solve. In the sequel of the proof, we mean by (G, Ω) a two-player zero-sum game, where player A (resp. player B) aims at satisfying Ω (resp. $V^\omega \setminus \Omega$).

If Ω is an Explicit Muller objective, then deciding the existence of a winning strategy for player A (resp. player B) (G, Ω) can be done in polynomial time by [27]. Thus the case where Ω is the opposite of an Explicit Muller objective is also proved (by exchanging players A and B).

Suppose that Ω is the conjunction of a Generalized Büchi objective and a co-Büchi objective. By a classic reduction (see [28, Theorem 12]), the game (G, Ω) can be polynomially transformed into a game (G', Ω') with an objective Ω' equal to the conjunction of a Büchi objective and a co-Büchi objective. The existence of a winning strategy for player A in the latter game can be tested in polynomial time [29].

Suppose that Ω is the conjunction of a Generalized Reachability objective and a Safety objective, such that ℓ is the number of Reachability objectives

and F is the set of vertices to be avoided in the Safety objective. We first treat separately the Safety objective by removing from G all the vertices of F . This can be done in $\mathcal{O}(|V| + |E|)$ time. In the resulting graph G' , we then test the existence of a winning strategy for player A in the game (G', Ω') with Ω' being the Generalized Reachability objective. This can be done in $\mathcal{O}(2^\ell(|V| + |E|))$ time [30].

If Ω is a BC Büchi objective $(\phi, F_1, \dots, F_\ell)$, then deciding the existence of a winning strategy for player A in the game (G, Ω) can be done in $\mathcal{O}(2^L \cdot |\phi| + (L^L \cdot |V|)^5)$ time with $L = 2^\ell$ by [31]. \square

Proof of Lemma 31. Let us prove Part 1 of the lemma. A play with payoff p from v in \mathcal{G} is a play satisfying an objective Ω equal to the conjunction of objectives Obj_i (when $p_i = 1$) and of objectives $V^\omega \setminus \text{Obj}_i$ (when $p_i = 0$). For each type of Boolean objectives Obj_i , we first explain what kind of objective Ω we obtain and we then apply Proposition 32.

- Consider a Boolean game \mathcal{G} with Parity objectives. In this case, as Obj_i is a Parity objective for all $i \in \Pi$, each $V^\omega \setminus \text{Obj}_i$ is again a Parity objective, and Ω is thus a conjunction of Parity objectives which is a Streett objective [32]. Therefore the existence of a play with payoff p in \mathcal{G} can be tested in polynomial time by Proposition 32.
- Consider the case of Büchi objectives. Then, the intersection of Büchi objectives Obj_i (when $p_i = 1$) is a Generalized Büchi objective and the intersection of co-Büchi objectives $V^\omega \setminus \text{Obj}_i$ (when $p_i = 0$) is again a co-Büchi objective. Hence Ω is the conjunction of a Generalized Büchi objective and a co-Büchi objective. The existence of a play with payoff p in \mathcal{G} can be tested in polynomial time by Proposition 32.

Notice that the case of Boolean games with co-Büchi objectives is solved exactly in the same way. Indeed we have the same kind of objective Ω since Obj_i is a co-Büchi objective if and only if $V^\omega \setminus \text{Obj}_i$ is a Büchi objective.

- Consider a Boolean game with Reachability objectives. The intersection of Reachability objectives Obj_i (when $p_i = 1$) is a Generalized Reachability objective and the intersection of Safety objectives $V^\omega \setminus \text{Obj}_i$ (when $p_i = 0$) is again a Safety objective. The existence of a play with payoff p in \mathcal{G} can be tested in $\mathcal{O}(2^{|\Pi|}(|V| + |E|))$ time by Proposition 32 as there are at most $|\Pi|$ Reachability objectives.

The case of Boolean games with Safety objectives is solved in the same way.

- Consider a Boolean game with Rabin objectives (with the related families $(G_j^i, R_j^i)_{1 \leq j \leq k_i, i \in \Pi}$). In this case, the objective Ω is the conjunction of Rabin objectives (when $p_i = 1$) and of Streett objectives (when $p_i = 0$), that is, Ω is a BC Büchi objective $(\phi, (G_j^i, R_j^i)_{1 \leq j \leq k_i, i \in \Pi})$ such that

$$\phi = \bigwedge_{i|p_i=1} \bigvee_{j=1}^{k_i} (g_j^i \wedge \neg r_j^i) \quad \wedge \quad \bigwedge_{i|p_i=0} \bigwedge_{j=1}^{k_i} (\neg g_j^i \vee r_j^i) \quad (12)$$

In this formula, the variable g_j^i (resp. r_j^i) is associated with the set G_j^i (resp. R_j^i), and ϕ has size $\mathcal{O}(\sum_{i=1}^{|\Pi|} 2 \cdot k_i)$ and has $\sum_{i=1}^{|\Pi|} 2 \cdot k_i$ variables. The announced complexity for deciding the existence of a play with payoff p follows from Proposition 32.

The case of Boolean games with Streett objectives is solved in the same way.

- The case of Boolean games with Muller objectives (with the related coloring functions $\Omega_i : V \rightarrow \{1, \dots, d_i\}$ and families $\mathcal{F}_i \subseteq 2^{\Omega_i(V)}$, $i \in \Pi$) is treated as in the previous item. Indeed a play satisfies the Muller objective Obj_i if there exists an element F of \mathcal{F}_i such that all colors of F are seen infinitely often along the play while no other color is seen infinitely often. Therefore, as the objective Ω is a conjunction of Muller objectives and of the opposite of Muller objectives, Ω is a BC Büchi objective $(\phi, (F_c^i)_{c \in \{1, \dots, d_i\}, i \in \Pi})$ described by the following formula ϕ

$$\phi = \bigwedge_{i|p_i=1} \bigvee_{F \in \mathcal{F}_i} \left(\bigwedge_{c \in F} f_c^i \wedge \bigwedge_{c \notin F} \neg f_c^i \right) \quad \wedge \quad \bigwedge_{i|p_i=0} \bigwedge_{F \in \mathcal{F}_i} \left(\bigvee_{c \in F} \neg f_c^i \vee \bigvee_{c \notin F} f_c^i \right) \quad (13)$$

In this formula, the variable f_c^i is associated with the subset $F_c^i = \{v \in V \mid \Omega_i(v) = c\}$ of vertices colored by color $c \in \{1, \dots, d_i\}$, $i \in \Pi$. This formula has size $\mathcal{O}(\sum_{i=1}^{|\Pi|} |\mathcal{F}_i| \cdot d_i)$ and has $\sum_{i=1}^{|\Pi|} d_i$ variables.

- It remains to treat the case of Boolean games with Explicit Muller objectives (with the related families $\mathcal{F}_i \subseteq 2^V$, $i \in \Pi$). The approach is a little different in a way to get a polynomial algorithm. By definition,

there exists a play ρ with payoff p if and only if for all i , $F = \text{Inf}(\rho) \in \mathcal{F}_i$ exactly when $p_i = 1$.

If $p \neq (0, \dots, 0)$, such potential sets F can be computed as follows. Initially let Ω be an empty set. Then for each $F \in \cup_{i \in \Pi} \mathcal{F}_i$, we compute $q \in \{0, 1\}^{|\Pi|}$ such that $q_i = 1$ if and only if $F \in \mathcal{F}_i$, and if $p = q$ we add F to Ω . Notice that Ω can be computed in polynomial time. Hence to test the existence of a play with payoff p in \mathcal{G} , we test the existence of a path in G satisfying the Explicit Muller objective Ω . This can be done in polynomial time by Proposition 32.

If $p = (0, \dots, 0)$, there exists a play ρ with payoff p if and only if no $F \in \Omega' = \cup_{i \in \Pi} \mathcal{F}_i$ is equal to $\text{Inf}(\rho)$, *i.e.*, if and only if there exists a path in G satisfying the opposite of the Explicit Muller objective Ω' . This can be tested in polynomial time by Proposition 32.

We now turn to Part 2 of the lemma. Suppose that the number $|\Pi|$ of players is fixed. In case of Boolean games with Reachability, Safety, Büchi, co-Büchi, Explicit Muller, and Parity objectives, by Part 1 of the lemma, we get a polynomial time algorithm for deciding whether there exists a play with payoff p . In case of Boolean games with Rabin, Streett, and Muller objectives, we need another argument in view of the complexities of Part 1 of the lemma.

Let us begin with Rabin objectives (the argument is similar for Streett objectives). Recall that we are faced with a BC Büchi objective described by formula ϕ given in (12). This formula is a conjunction of disjunctions of subformulas of the form either $g_j^i \wedge \neg r_j^i$, or $\neg g_j^i$, or r_j^i . It can be rewritten as a disjunction of conjunctions of those subformulas: $\bigvee_r \bigwedge_s \phi_{rs}$. In the latter formula,

- the number of subformulas $\bigwedge_s \phi_{rs}$ is polynomial since it is equal to $\prod_{i|p_i=1} k_i \cdot \prod_{i|p_i=0} 2k_i$, and
- each subformula $\bigwedge_s \phi_{rs}$ describes a polynomial conjunction of Büchi and co-Büchi objectives, that is, an objective Ω_r equal to the conjunction of a Generalized Büchi objective and a co-Büchi objective.

By Proposition 32, testing whether there exists a path satisfying Ω_r can be done in polynomial time. Therefore testing the existence of a path satisfying the objective described by ϕ reduces to a polynomial number of tests (disjunction \bigvee_r) that can be done in polynomial time (objective Ω_r).

The argument for Muller objectives is similar. The formula ϕ given in (13) is a conjunction of disjunctions of subformulas of the form either $\bigwedge_{c \in F} f_c^i \wedge \bigwedge_{c \notin F} \neg f_c^i$, or $\neg f_c^i$, or f_c^i . Rewriting ϕ as a disjunction of conjunctions of those subformulas, that is, as $\bigvee_r \bigwedge_s \phi_{rs}$, we get again a polynomial number ($= \prod_{i|p_i=1} |\mathcal{F}_i| \cdot \prod_{i|p_i=0} (|\mathcal{F}_i| \cdot d_i)$) of subformulas $\bigwedge_s \phi_{rs}$, each of them describing a polynomial conjunction of Büchi and co-Büchi objectives. \square

5.3. P-completeness for Explicit Muller objectives and FPT for the other objectives

With the complexity of the decision algorithm given in Lemma 30 and the study of *path* made in Lemma 31, we are now ready to show that the constraint problem is P-complete for Explicit Muller objectives, and that for the other objectives, it is fixed parameter tractable and becomes polynomial when the number $|\Pi|$ of players is fixed.

Theorem 33. *The constraint problem for multiplayer Boolean games with Explicit Muller objectives is P-complete.*

Proof. We denote by $\mathcal{F}_i \subseteq 2^V$, the family of each player $i \in \Pi$ for his Explicit Muller objective. Let us prove the P-easiness. By Lemma 30, the decision algorithm is in $\mathcal{O}(m^3 \cdot |V| \cdot |\Pi| \cdot \text{init} \cdot \text{path} \cdot (|V| + |E|))$ time, where $m = \max_{v \in V} |\mathbf{P}_0(v)|$, *init* is the time complexity of computing $\mathbf{P}_0(v)$ for some given vertex v , and *path* is the time complexity of testing whether there exists a play with a given payoff p from a given vertex v . Lemma 31 states that *path* is polynomial for Boolean games with Explicit Muller objectives. To establish the P-easiness, it remains to prove that m and *init* are also polynomial. First, if there exists a play ρ with payoff p from v , then either $\text{Inf}(\rho) \in \cup_{i \in \Pi} \mathcal{F}_i$ or $p = (0, \dots, 0)$. Thus

$$\mathbf{P}_0(v) \subseteq P = \{q \in \{0, 1\}^{|\Pi|} \mid \exists F \in \cup_{i \in \Pi} \mathcal{F}_i, q_i = 1 \Leftrightarrow F \in \mathcal{F}_i\} \cup \{(0, \dots, 0)\}.$$

(This kind of argument was already used in the proof of Lemma 31 for Explicit Muller objectives). Therefore, $|\mathbf{P}_0(v)| \leq |P| \leq |\cup_{i \in \Pi} \mathcal{F}_i| + 1$, showing that m is polynomial. Second, to compute $\mathbf{P}_0(v)$, we check for each $p \in P$ whether there exists a play with payoff p from v . It follows that *init* is polynomial by Lemma 31.

The P-hardness is obtained thanks to a reduction from the AND-OR graph reachability problem that is P-complete [33]. Indeed, the P-hardness of the constraint problem for SPEs (instead of weak SPEs) in Boolean games

with Reachability objectives is proved in [25, Corollary 6.22] thanks to such a reduction, and it is not difficult to see that the same reduction also holds for weak SPEs and Explicit Muller objectives. \square

We recall that a *parameterized language* L is a subset of $\Sigma^* \times \mathbb{N}$, where Σ is a finite alphabet, the second component being the parameter of the language. It is called *fixed parameter tractable* if there is an algorithm that determines whether $(x, t) \in L$ in $f(t) \cdot |x|^c$ time, where c is a constant independent of the parameter t and f is a computable function depending on t only. We also say that L belongs to the class FPT. Intuitively, a language is in FPT if there is an algorithm running in polynomial time with respect to the input size times some computable function on the parameter. We refer the interested reader to [34] for more details on parameterized complexity.

Theorem 34. *Let \mathcal{G} be a Boolean game.*

1. *The constraint problem is in FPT for Reachability, Safety, Büchi, co-Büchi, Parity, Muller, Rabin, and Streett objectives. The parameters are*
 - *the number $|\Pi|$ of players for Reachability, Safety, Büchi, co-Büchi, and Parity objectives,*
 - *the number $|\Pi|$ of players and the numbers k_i , $i \in \Pi$, of pairs $(G_j^i, R_j^i)_{1 \leq j \leq k_i}$, for Rabin and Streett objectives, and*
 - *the number $|\Pi|$ of players, the numbers d_i , $i \in \Pi$, of colors and the sizes $|\mathcal{F}_i|$, $i \in \Pi$, of the families of subsets of colors for Muller objectives.*
2. *When the number $|\Pi|$ of players is fixed, for all these kinds of objectives, the constraint problem can be solved in polynomial time.*

Notice that in this theorem, to obtain fixed parameter tractability for Rabin, Streett, and Muller objectives, in addition to the number of players, we also have to consider the parameter equal to the size of the objective description. Nevertheless, when the number of players is fixed, we get polynomial time complexity for all types of objectives.

Proof of Theorem 34. We again use Lemmas 30 and 31 as in the previous proof. By Lemma 30, the complexity of the fixpoint algorithm is in $\mathcal{O}(m^3 \cdot |V| \cdot |\Pi| \cdot \text{init} \cdot \text{path} \cdot (|V| + |E|))$ time. The complexity *path* is studied in Part

1 (resp. in Part 2 when $|\Pi|$ is fixed) of Lemma 31, and we have $m \leq 2^{|\Pi|}$ and $init = \mathcal{O}(2^{|\Pi|} \cdot path)$. Hence we get both parts of Theorem 34 as a consequence of these results. \square

6. Conclusion and future work

In this paper, we have studied the computational complexity of the constraint problem for weak SPEs. We were able to obtain precise complexities for all the classical classes of ω -regular objectives (see Table 1), with one exception: we have proved NP-membership for Büchi objectives and failed to prove NP-hardness. We have also shown that the constraint problem can be solved in polynomial time when the number of players is fixed. Finally, we have provided some fixed parameter tractable algorithms when the number of players is considered as a parameter of the problem, for Reachability, Safety, Büchi, Co-Büchi, and Parity objectives. For the other Rabin, Streett, and Muller objectives, we also had to consider the size of the objective description as a parameter to obtain fixed parameter tractability. In a future work, we want to understand if the use of this second parameter is really necessary.

By characterizing the exact complexity of the constraint problem for Reachability and Safety objectives, we have obtained that this problem for SPEs (as for weak SPEs) is PSPACE-complete for those objectives. In the future, we intend to investigate the complexity of the other classes of ω -regular objectives for SPEs. It would be also interesting to extend the study to quantitative games. For instance the constraint problem for (weak) SPEs in reachability quantitative games is decidable [17] but its complexity is unknown.

References

- [1] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: POPL, ACM Press, 1989, pp. 179–190 (1989).
- [2] W. Thomas, On the synthesis of strategies in infinite games, in: STACS, Vol. 900 of LNCS, Springer, 1995, pp. 1–13 (1995).
- [3] D. Berwanger, Admissibility in infinite games, in: STACS, Vol. 4393 of LNCS, Springer, 2007, pp. 188–199 (2007).

- [4] R. Brenguier, L. Clemente, P. Hunter, G. A. Pérez, M. Randour, J. Raskin, O. Sankur, M. Sassolas, Non-zero sum games for reactive synthesis, in: LATA, Vol. 9618 of LNCS, Springer, 2016, pp. 3–23 (2016).
- [5] R. Brenguier, J.-F. Raskin, O. Sankur, Assume-admissible synthesis, in: CONCUR, LIPIcs 42, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015, pp. 100–113 (2015).
- [6] R. Brenguier, J. Raskin, M. Sassolas, The complexity of admissibility in omega-regular games, in: CSL-LICS, ACM, 2014, pp. 23:1–23:10 (2014).
- [7] V. Bruyère, N. Meunier, J. Raskin, Secure equilibria in weighted games, in: CSL-LICS, ACM, 2014, pp. 26:1–26:26 (2014).
- [8] K. Chatterjee, T. A. Henzinger, M. Jurdzinski, Games with secure equilibria, Theoretical Computer Science 365 (2006) 67–82 (2006).
- [9] D. Fisman, O. Kupferman, Y. Lustig, Rational synthesis, in: TACAS, Vol. 6015 of LNCS, Springer, 2010, pp. 190–204 (2010).
- [10] O. Kupferman, G. Perelli, M. Y. Vardi, Synthesis with rational environments, in: EUMAS, LNCS 8953, Springer, 2014, pp. 219–235 (2014).
- [11] M. Ummels, Rational behaviour and strategy construction in infinite multiplayer games, in: FSTTCS, Vol. 4337 of LNCS, Springer, 2006, pp. 212–223 (2006).
- [12] V. Bruyère, Computer aided synthesis: A game-theoretic approach, in: DLT, Vol. 10396 of LNCS, Springer, 2017, pp. 3–35 (2017).
- [13] E. Grädel, M. Ummels, Solution Concepts and Algorithms for Infinite Multiplayer Games, in: New Perspectives on Games and Interaction, Vol. 4, Amsterdam University Press, 2008, pp. 151–178 (2008).
- [14] J. F. Nash, Equilibrium points in n -person games, in: PNAS, Vol. 36, National Academy of Sciences, 1950, pp. 48–49 (1950).
- [15] M. Osborne, An introduction to game theory, Oxford Univ. Press, 2004 (2004).

- [16] T. Brihaye, V. Bruyère, J. De Pril, H. Gimbert, On subgame perfection in quantitative reachability games, *Logical Methods in Computer Science* 9 (2012).
- [17] T. Brihaye, V. Bruyère, N. Meunier, J. Raskin, Weak subgame perfect equilibria and their application to quantitative reachability, in: *CSL*, Vol. 41 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015, pp. 504–518 (2015).
- [18] V. Bruyère, S. L. Roux, A. Pauly, J. Raskin, On the existence of weak subgame perfect equilibria, in: *FOSSACS*, Vol. 10203 of *LNCS*, Springer, 2017, pp. 145–161 (2017).
- [19] T. Brihaye, V. Bruyère, A. Goeminne, J. Raskin, M. van den Bogaard, The complexity of subgame perfect equilibria in quantitative reachability games, in: *CONCUR*, Vol. 140 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 13:1–13:16 (2019).
- [20] J. Flesch, J. Kuipers, A. Mashiah-Yaakovi, G. Schoenmakers, E. Solan, K. Vrieze, Perfect-information games with lower-semicontinuous payoffs, *Mathematics of Operation Research* 35 (2010) 742–755 (2010).
- [21] D. Fudenberg, J. Tirole, *Game Theory*, MIT Press, 1991 (1991).
- [22] P. Bouyer, R. Brenguier, N. Markey, M. Ummels, Pure Nash equilibria in concurrent deterministic games, *Logical Methods in Computer Science* 11 (2) (2015).
- [23] M. Ummels, The complexity of Nash equilibria in infinite multiplayer games, in: *FOSSACS*, Vol. 4962 of *LNCS*, Springer, 2008, pp. 20–34 (2008).
- [24] E. Grädel, W. Thomas, T. Wilke (Eds.), *Automata, Logics, and Infinite Games: A Guide to Current Research*, Vol. 2500 of *LNCS*, Springer, 2002.
- [25] M. Ummels, *Rational Behaviour and Strategy Construction in Infinite Multiplayer Games*, Diploma thesis, RWTH Aachen (2005).
URL <http://www.logic.rwth-aachen.de/~ummels/diplom.pdf>

- [26] E. A. Emerson, C. Lei, Modalities for model checking: Branching time logic strikes back, *Science of Computer Programming* 8 (3) (1987) 275–306 (1987).
- [27] F. Horn, Explicit Muller games are PTIME, in: *FSTTCS*, Vol. 2 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008, pp. 235–243 (2008).
- [28] E. M. Clarke, T. A. Henzinger, H. Veith, R. Bloem (Eds.), *Handbook of Model Checking*, Springer, 2018 (2018).
- [29] L. de Alfaro, M. Faella, An accelerated algorithm for 3-color parity games with an application to timed games, in: *CAV*, Vol. 4590 of *LNCS*, Springer, 2007, pp. 108–120 (2007).
- [30] N. Fijalkow, F. Horn, Les jeux d’accessibilité généralisée, *Technique et Science Informatiques* 32 (9-10) (2013) 931–949 (2013).
- [31] V. Bruyère, Q. Hautem, J. Raskin, Parameterized complexity of games with monotonically ordered omega-regular objectives, in: *CONCUR 2018*, Vol. 118 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018, pp. 29:1–29:16 (2018).
- [32] K. Chatterjee, T. A. Henzinger, N. Piterman, Generalized parity games, in: *FOSSACS*, Vol. 4423 of *LNCS*, Springer, 2007, pp. 153–167 (2007).
- [33] N. Immerman, Number of quantifiers is better than number of tape cells, *Journal of Computer and System Sciences* 22 (1981) 384–406 (1981).
- [34] R. G. Downey, M. R. Fellows, *Parameterized Complexity*, Springer Publishing Company, Incorporated, 2012 (2012).