

PURE NASH EQUILIBRIA IN CONCURRENT DETERMINISTIC GAMES

PATRICIA BOUYER¹, ROMAIN BRENGUIER², NICOLAS MARKEY¹, AND MICHAEL UMMELS³

¹ LSV – CNRS & ENS Cachan – France
e-mail address: {bouyer,markey}@lsv.ens-cachan.fr

² Université Libre de Bruxelles – Belgium
e-mail address: romain.brenguier@ulb.ac.be

³ Institute of Transportation Systems, German Aerospace Center – Germany
e-mail address: michael.ummels@dlr.de

ABSTRACT. We study pure-strategy Nash equilibria in multi-player concurrent deterministic games, for a variety of preference relations. We provide a novel construction, called the suspect game, which transforms a multi-player concurrent game into a two-player turn-based game which turns Nash equilibria into winning strategies (for some objective that depends on the preference relations of the players in the original game). We use that transformation to design algorithms for computing Nash equilibria in finite games, which in most cases have optimal worst-case complexity, for large classes of preference relations. This includes the purely qualitative framework, where each player has a single ω -regular objective that she wants to satisfy, but also the larger class of semi-quantitative objectives, where each player has several ω -regular objectives equipped with a preorder (for instance, a player may want to satisfy all her objectives, or to maximise the number of objectives that she achieves.)

1. INTRODUCTION

Games (and especially games played on graphs) have been intensively used in computer science as a powerful way of modelling interactions between several computerised systems [39, 24]. Until recently, more focus had been put on the study of purely antagonistic games (a.k.a. zero-sum games), which conveniently represent systems evolving in a (hostile) environment. In this zero-sum games setting, the objectives of both players are opposite: the aim of one player is to prevent the other player from achieving her own objective.

1998 ACM Subject Classification: I.2.2, F.1.1, F.3.1, F.4.1.

Key words and phrases: Nash equilibria, concurrent games, qualitative objectives, ordered objectives.

This article is an extended version of several works that appeared as [4, 5, 6]. Most of the work reported here was done while the second and fourth authors were students at LSV. This work has been partly supported by ESF LogiCCC project GASICS, ERC Starting Grant EQualIS (308087), ERC Starting Grant inVEST (279499), and EU FP7 project Cassting (601148).

Over the last ten years, games with non-zero-sum objectives have come into the picture: they allow for conveniently modelling complex infrastructures where each individual system tries to fulfil its own objectives, while still being subject to uncontrollable actions of the surrounding systems. As an example, consider a wireless network in which several devices try to send data: each device can modulate its transmit power, in order to maximise its bandwidth and reduce energy consumption as much as possible. In that setting, focusing only on optimal strategies for one single agent may be too narrow. Game-theoreticians have defined and studied many other solution concepts for such settings, of which Nash equilibrium [35] is the most prominent. A Nash equilibrium is a strategy profile where no player can improve the outcome of the game by unilaterally changing her strategy. In other terms, in a Nash equilibrium, each individual player has a satisfactory strategy. Notice that Nash equilibria need not exist or be unique, and are not necessarily optimal: Nash equilibria where all players lose may coexist with more interesting Nash equilibria. Finding *constrained* Nash equilibria (e.g., equilibria in which some players are required to win) is thus an interesting problem for our setting.

In this paper, we report on our recent contributions on the computation of Nash equilibria in concurrent games (preliminary works appeared as [4, 5, 6]). Concurrent games played on graphs are a general model for interactive systems, where the agents take their decision simultaneously. Therefore concurrent games subsume turn-based games, where in each state, only one player has the decision for the next move. One motivation for concurrent games is the study of *timed games* (which are games played on timed automata [3, 1]): the semantics of a timed game is naturally given as a concurrent game (the players all choose simultaneously a delay and an action to play, and the player with the shortest delay decides for the next move—this mechanism cannot be made turn-based since we cannot fix *a priori* the player who will choose the smallest delay); the region-based game abstraction which preserves Nash equilibria also requires the formalism of concurrent games [5, 7]. Multi-agent infrastructures can be viewed as distributed systems, which can naturally be modelled as concurrent games.

Our contributions. The paper focuses on concurrent deterministic games and on *pure* Nash equilibria, that is, strategy profiles which are deterministic (as opposed to randomised). In this work we assume strategies only depend on the set of states which is visited, and not on the actions that have been played. This is a partial-information hypothesis which we believe is relevant in the context of distributed systems, where only the effect of the actions can be seen by the players. We will discuss in more detail all these choices in the conclusion.

In the context exposed above, we develop a complete methodology for computing pure Nash equilibria in (finite) games. First, in Section 4, we propose a novel transformation of the multi-player concurrent game (with a preference relation for each player) into a two-player zero-sum turn-based game, which we call the *suspect game*. Intuitively, in the suspect game, one of the players suggests a global move (one action per player of the original game), with the aim to progressively build a Nash equilibrium; while the second player aims at proving that what the first player proposes is *not* a Nash equilibrium. This transformation can be applied to arbitrary concurrent games (even those with infinitely many states) and preference relations for the players, and it has the property that there is a correspondence between Nash equilibria in the original game and winning strategies in the transformed two-player turn-based game. The winning condition in the suspect game of course depends on the preference relations of the various players in the original game.

Objective	Value	(Constrained) Existence of Nash Eq.
Reachability	P-c. [32]	NP-c. (Sect. 5.1)
Safety	P-c. [32]	NP-c. (Sect. 5.2)
Büchi	P-c. [32]	P-c. (Sect. 5.3)
co-Büchi	P-c. [32]	NP-c. (Sect. 5.4)
Parity	$UP \cap co-UP$ [28]	P_{\parallel}^{NP} -c. ¹ (Sect. 5.6)
Streett	co-NP-c. [18]	P_{\parallel}^{NP} -h. and in PSPACE
Rabin	NP-c. [18]	P_{\parallel}^{NP} -c. (Sect. 5.6)
Muller	PSPACE-c. [27]	PSPACE-c.
Circuit	PSPACE-c. [27]	PSPACE-c. (Sect. 5.5)
Det. Büchi Automata	P-c.	PSPACE-h. (Sect. 5.7) and in EXPTIME
Det. Rabin Automata	NP-c.	PSPACE-h. and in EXPTIME (Sect. 5.7)

TABLE 1. Summary of the complexities for single objectives

Preorder	Value	Existence of NE	Constr. Exist. of NE
Maximise, Disj.	P-c. (Sect.6.2)	P-c. (Sect.6.2)	P-c. (Sect.6.2)
Subset	P-c. (Sect. 6.3)	P-c. (Sect.6.2)	P-c. (Sect.6.2)
Conj., Lexicogr.	P-c. (Sect. 6.3)	P-h., in NP (Sect. 6.4)	NP-c. (Sect. 6.4)
Counting	coNP-c. (Sect. 6.4)	NP-c. (Sect. 6.4)	NP-c. (Sect. 6.4)
Mon. Bool. Circuit	coNP-c. (Sect. 6.4)	NP-c. (Sect. 6.4)	NP-c. (Sect. 6.4)
Boolean Circuit	PSPACE-c. (Sect. 6.1)	PSPACE-c. (Sect. 6.1)	PSPACE-c. (Sect. 6.1)

TABLE 2. Summary of the results for ordered Büchi objectives

Preorder	Value	(Constrained) Exist. of NE
Disjunction, Maximise	P-c. (Sect. 7.2)	NP-c. (Sect. 7.2)
Subset	PSPACE-c. (Sect. 7.1)	NP-c. (Sect. 7.2)
Conjunction, Counting, Lexicogr.	PSPACE-c. (Sect. 7.1)	PSPACE-c. (Sect. 7.1)
(Monotonic) Boolean Circuit	PSPACE-c. (Sect. 7.1)	PSPACE-c. (Sect. 7.1)

TABLE 3. Summary of the results for ordered reachability objectives

Then, using that construction we develop (worst-case) optimal-complexity algorithms for deciding the existence of (constrained) Nash equilibria in *finite* games for various classes of preference relations. In Section 5, we focus on qualitative ω -regular objectives, i.e., preference relations are given by single objectives (which can be reachability, Büchi, parity, etc), and it is better for a player to satisfy her objective than to not satisfy her objective. We prove the whole set of results which are summarised in the second column of Table 1 (the first column summarises the complexity in the zero-sum two-player setting – called the value problem). Among the results obtained this way, the constrained Nash equilibrium existence problem is NP-complete in finite games with single reachability or safety objectives, while it is PTIME-complete for single Büchi objectives.

¹The complexity class P_{\parallel}^{NP} is defined in terms of Turing machine having access to an oracle; oracle are artificial devices that can solve a problem in constant time, thus hiding part of the complexity of the overall problem. The class P^{NP} is the class of problems that can be solved in polynomial time by a deterministic




FIGURE 1. A simple game-model for the wireless network

In Sections 6 and 7, we extend the previous qualitative setting to the *semi-quantitative* setting of ordered objectives. An ordered objective is a set of Büchi (or reachability) objectives and a preorder on this set. The preference relation given by such an ordered objective is then given by the value of the plays (w.r.t. the objectives) in that preorder. Preorders of interest are for instance conjunction, disjunction, lexicographic order, counting preorder, maximise preorder, subset preorder, or more generally preorders given as Boolean circuits. We provide algorithms for deciding the existence of Nash equilibria for ordered objectives, with (in most cases) optimal worst-case complexity. These algorithms make use of the suspect-game construction. The results are listed in Table 2 for Büchi objectives and in Table 3 for reachability objectives.

Examples. Back to the earlier wireless network example, we can model a simple discretised version of it as follows. From a state, each device can increase (action 1) or keep unchanged (action 0) its power: the arena of the game is represented for two devices and two levels of power on Figure 1 (labels of states are power levels). This yields a new bandwidth allocation (which depends on the degradation due to the other devices) and a new energy consumption. The satisfaction of each device is measured as a compromise between energy consumption and bandwidth allocated, and it is given by a quantitative payoff function.² This can be transformed into Büchi conditions and a preorder on them. There are basically two families of pure Nash equilibria in this system: the one where the two players choose to go and stay forever in state (1, 1); and the one where the two players go to state (2, 2) and stay there forever.

We describe another example, the medium access control, that involves qualitative objectives. It was first given a game-theoretic model in [31]. Several users share the access to a wireless channel. During each slot, they can choose to either transmit or wait for the next slot. If too many users are emitting in the same slot, then they fail to send data. Each attempt to transmit costs energy to the players. They have to maximise the number of

Turing machine which has access to an oracle for solving NP problems. The class $P_{\parallel}^{\text{NP}}$ is the subclass where, instead of asking a sequence of (dependent) queries to the oracle, the Turing machine is only allowed to ask one set of queries. We refer to [36, 44] for more details.

²The (quantitative) payoff for player i can be expressed by $\text{payoff}_i = \frac{R}{\text{power}_i} \left(1 - e^{-0.5\gamma_i}\right)^L$ where γ_i is the signal-to-interference-and-noise ratio for player i , R is the rate at which the wireless system transmits the information in bits per seconds and L is the size of the packets in bits ([37]).




FIGURE 2. A simple game-model for the medium access control

successful attempts using the energy available to them. We give in Figure 2 a possible model for that protocol for two players and at most one attempt per player and a congestion of 2 (that is, the two players should not transmit at the same time): each state is labelled with the energy level of the two players, and the number of successful attempts of each of the player. There is several Nash equilibria, and they give payoff 1 to every player: it consists in going to state $(0, 1, 0, 1)$ by not simultaneously transmitting.

Related work. Game theory has been a very active area since the 1940's, with the pioneering works of Von Neumann, Morgenstern [43], Nash [35] and Shapley [38]. It has had numerous uses in various domains, ranging from economics to human sciences and logic. Equilibria are a central concept in (non-zero-sum) games, as they are meant to represent rational behaviours of the players. Many important results about existence of various kinds of equilibria in different kinds of games have been established [43, 35, 20].

For applications in logic and computer science, games played on graphs have received more focus; also, computer scientists have been mostly looking for algorithmic solutions for deciding the existence and effectively computing equilibria and ϵ -equilibria [15, 10, 40].

For two-player concurrent games with Büchi objectives, the existence of ϵ -equilibria (in randomised strategies) was proved by Chatterjee [10]. However, exact Nash equilibria need not exist; turn-based games with Büchi objectives are an important subclass where Nash equilibria (even in pure strategies) always exist [15]. When they exist, Nash equilibria need not be unique; equilibria where all the players lose can coexist with equilibria where some (or all) of them win. Ummels introduced *constrained* Nash equilibria, i.e., Nash equilibria where some players are required to win. In particular, he showed that the existence of constrained Nash equilibria can be decided in polynomial time for turn-based games with Büchi objectives [40]. In this paper, we extend this result to concurrent games, and to various classes of ω -regular winning objectives. For concurrent games with ω -regular objectives, the decidability of the constrained Nash equilibrium existence problem w.r.t. pure strategies was established by Fisman et al. [21], but their algorithm runs in doubly exponential time, whereas our algorithm runs in exponential time for objectives given as Büchi automata. Finally, Ummels and Wojtczak [42] proved that the existence of a Nash equilibrium in pure or randomised strategies is undecidable for *stochastic* games with reachability or Büchi objectives, which justifies our restriction to concurrent games without probabilistic transitions. They also proved a similar undecidability result for randomised Nash equilibria in non-stochastic games [41], hence we consider only pure-strategy Nash equilibria.

Several solution concepts have been defined and studied for games on graphs. In particular, *secure equilibria* [11, 17] are Nash equilibria where besides satisfying their primary objectives, the players try to prevent the other players from achieving their own (primary) objectives. Notice that our results in Sect. 6.4 and Sect. 7.1 do apply to such kinds of lexicographic combination of several objectives.

Temporal logics can also be used to express properties of games. While **ATL** [2] can mainly express only zero-sum properties, other logics such as **ATL** with strategy contexts (**ATL_{sc}**) [16] and **Strategy Logic (SL)** [14, 33] can be used to express rich properties in a non-zero-sum setting. In terms of complexity however, model checking for such logics has high complexity: Nash equilibria can be expressed using one quantifier alternation (an existential quantification over strategy profiles followed with a universal quantification over deviations); model checking this fragment of **ATL_{sc}** or **SL** is 2-EXPTIME-complete.

2. DEFINITIONS

2.1. General definitions. In this section, we fix some definitions and notations.

Preorders. We fix a non-empty set P . A *preorder* over P is a binary relation $\lesssim \subseteq P \times P$ that is reflexive and transitive. With a preorder \lesssim , we associate an *equivalence relation* \sim defined so that $a \sim b$ if, and only if, $a \lesssim b$ and $b \lesssim a$. The *equivalence class* of a , written $[a]_{\lesssim}$, is the set $\{b \in P \mid a \sim b\}$. We also associate with \lesssim a *strict partial order* \prec defined so that $a \prec b$ if, and only if, $a \lesssim b$ and $b \not\lesssim a$. A preorder \lesssim is said *total* if, for all elements $a, b \in P$, either $a \lesssim b$, or $b \lesssim a$. An element a in a subset $P' \subseteq P$ is said *maximal in P'* if there is no $b \in P'$ such that $a \prec b$; it is said *minimal in P'* if there is no $b \in P'$ such that $b \prec a$. A preorder is said *Noetherian* (or *upwards well-founded*) if any subset $P' \subseteq P$ has at least one maximal element. It is said *almost-well-founded* if any lower-bounded subset $P' \subseteq P$ has a minimal element.

Transition systems. A *transition system* is a pair $\mathcal{S} = \langle \text{States}, \text{Edg} \rangle$ where States is a set of states and $\text{Edg} \subseteq \text{States} \times \text{States}$ is the set of transitions. A *path* π in \mathcal{S} is a sequence $(s_i)_{0 \leq i < n}$ (where $n \in \mathbb{N}^+ \cup \{\infty\}$) of states such that $(s_i, s_{i+1}) \in \text{Edg}$ for all $i \leq n$. The *length* of π , denoted by $|\pi|$, is $n - 1$. The set of finite paths (also called *histories*) of \mathcal{S} is denoted by $\text{Hist}_{\mathcal{S}}$, the set of infinite paths (also called *plays*) of \mathcal{S} is denoted by $\text{Play}_{\mathcal{S}}$, and $\text{Path}_{\mathcal{S}} = \text{Hist}_{\mathcal{S}} \cup \text{Play}_{\mathcal{S}}$ is the set of all paths of \mathcal{S} . Given a path $\pi = (s_i)_{0 \leq i < n}$ and an integer $j < n$, the *j -th prefix* (resp. *j -th suffix*, *j -th state*) of π , denoted by $\pi_{\leq j}$ (resp. $\pi_{\geq j}$, $\pi_{=j}$), is the finite path $(s_i)_{0 \leq i < j+1}$ (resp. the path $(s_{j+i})_{0 \leq i < n-j}$, the state s_j). If $\pi = (s_i)_{0 \leq i < n}$ is a history, we write $\text{last}(\pi) = s_{|\pi|}$ for the last state of π . If π' is a path such that $(\text{last}(\pi), \pi'_{=0}) \in \text{Edg}$, then the *concatenation* $\pi \cdot \pi'$ is the path ρ s.t. $\rho_{=i} = \pi_{=i}$ for $i \leq |\pi|$ and $\rho_{=i} = \pi'_{=(i-1-|\pi|)}$ for $i > |\pi|$. In the sequel, we write $\text{Hist}_{\mathcal{S}}(s)$, $\text{Play}_{\mathcal{S}}(s)$ and $\text{Path}_{\mathcal{S}}(s)$ for the respective subsets of paths starting in state s . If π is a play, $\text{Occ}(\pi) = \{s \mid \exists j. \pi_{=j} = s\}$ is the sets of states that appears at least once along π and $\text{Inf}(\pi) = \{s \mid \forall i. \exists j \geq i. \pi_{=j} = s\}$ is the set of states that appears infinitely often along π .

2.2. Concurrent games. Our definition of concurrent games extends the definition in [2] by allowing for more than two players, each of them having a preorder over plays.

Definition 2.1. A *concurrent game* is a tuple $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$, where States is a finite non-empty set of states, Agt is a finite set of players, Act is a finite set of actions, and

- $\text{Mov}: \text{States} \times \text{Agt} \rightarrow 2^{\text{Act}} \setminus \{\emptyset\}$ is a mapping indicating the actions available to a given player in a given state;
- $\text{Tab}: \text{States} \times \text{Act}^{\text{Agt}} \rightarrow \text{States}$ associates, with a given state and a given move of the players (i.e., an element of Act^{Agt}), the state resulting from that move;
- for each $A \in \text{Agt}$, \preceq_A is a preorder over States^ω , called the preference relation of player A .

Figure 3 displays an example of a finite concurrent game. Transitions are labelled with the moves that trigger them. We say that a *move* $m_{\text{Agt}} = \langle m_A \rangle_{A \in \text{Agt}} \in \text{Act}^{\text{Agt}}$ is *legal* at s if $m_A \in \text{Mov}(s, A)$ for all $A \in \text{Agt}$. A game is *turn-based* if for each state the set of allowed moves is a singleton for all but at most one player.

In a concurrent game \mathcal{G} , whenever we arrive at a state s , the players simultaneously select an available action, which results in a legal move m_{Agt} ; the next state of the game is then $\text{Tab}(s, m_{\text{Agt}})$. The same process repeats *ad infinitum* to form an infinite sequence of states.

In the sequel, as no ambiguity will arise, we may abusively write \mathcal{G} for its underlying transition system $(\text{States}, \text{Edg})$ where $\text{Edg} = \{(s, s') \in \text{States} \times \text{States} \mid \exists m_{\text{Agt}} \in \prod_{A \in \text{Agt}} \text{Mov}(s, A) \text{ s.t. } \text{Tab}(s, m_{\text{Agt}}) = s'\}$. The notions of paths and related concepts in concurrent games follow from this identification.

Remark 2.2 (Representation of finite games). In this paper, for finite games, we will assume an explicit encoding of the transition function Tab . Hence, its size, denoted $|\text{Tab}|$, is equal to $\sum_{s \in \text{States}} \prod_{A \in \text{Agt}} |\text{Mov}(s, A)| \cdot \lceil \log(|\text{States}|) \rceil$. Note that it can be exponential with respect to the number of players. A symbolic encoding of the transition table has been proposed in [30], in the setting of ATL model checking. This makes the problem harder, as the input is more succinct (see Remark 5.1 and Proposition 5.2 for a formal statement). We would also have a blowup in our setting, and prefer to keep the explicit representation in order to be able to compare with existing results. Notice that, as a matter of fact, there is no way to systematically avoid an explosion: as there are $|\text{States}|^{|\text{States}| \cdot |\text{Act}|^{|\text{Agt}|}}$ possible transition functions, for any encoding there is one function whose encoding will have size at least $\lceil \log(|\text{States}|) \rceil \cdot |\text{States}| \cdot |\text{Act}|^{|\text{Agt}|}$. The total size of the game, is then

$$|\mathcal{G}| = |\text{States}| + |\text{States}| \cdot |\text{Agt}| \cdot |\text{Act}| + \sum_{s \in \text{States}} \prod_{A \in \text{Agt}} |\text{Mov}(s, A)| \cdot \lceil \log(|\text{States}|) \rceil + \sum_{A \in \text{Agt}} |\preceq_A|.$$

The size of a preference relation \preceq_A will depend on how it is encoded, and we will make it precise when it is relevant. This is given in Section 2.5.

Definition 2.3. Let \mathcal{G} be a concurrent game, and $A \in \text{Agt}$. A *strategy* for A is a mapping $\sigma_A: \text{Hist}_{\mathcal{G}} \rightarrow \text{Act}$ such that $\sigma_A(\pi) \in \text{Mov}(\text{last}(\pi), A)$ for all $\pi \in \text{Hist}_{\mathcal{G}}$. A strategy σ_P for a coalition $P \subseteq \text{Agt}$ is a tuple of strategies, one for each player in P . We write $\sigma_P = (\sigma_A)_{A \in P}$




FIGURE 3. Representation of a two-player concurrent game

for such a strategy. A *strategy profile* is a strategy for Agt. We write $\text{Strat}_{\mathcal{G}}^P$ for the set of strategies of coalition P , and $\text{Prof}_{\mathcal{G}} = \text{Strat}_{\mathcal{G}}^{\text{Agt}}$.

Note that, in this paper, we only consider *pure* (i.e., non-randomised) strategies. This is actually crucial in all the constructions we give (lasso representation in Subsection 3.1 and suspect-game construction in Section 4). Notice also that our strategies are based on the sequences of visited states (they map sequences of states to actions), which is realistic when considering multi-agent systems. In some settings, it is more usual to base strategies on the sequences of actions played by all the players. When dealing with Nash equilibria, this makes a big difference: strategies based on actions can immediately detect which player(s) deviated from their strategy; strategies based on states will only detect deviations because an unexpected state is visited, without knowing which player(s) is responsible for the deviation. Our construction precisely amounts to keeping track of a list of *suspects* for some deviation.

Let \mathcal{G} be a game, P a coalition, and σ_P a strategy for P . A path π is *compatible* with the strategy σ_P if, for all $k < |\pi|$, there exists a move m_{Agt} such that

- (1) m_{Agt} is legal at $\pi_{=k}$,
- (2) $m_A = \sigma_A(\pi_{\leq k})$ for all $A \in P$, and
- (3) $\text{Tab}(\pi_{=k}, m_{\text{Agt}}) = \pi_{=k+1}$.

We write $\text{Out}_{\mathcal{G}}(\sigma_P)$ for the set of paths (called the *outcomes*) in \mathcal{G} that are compatible with strategy σ_P of P . We write $\text{Out}_{\mathcal{G}}^f$ (resp. $\text{Out}_{\mathcal{G}}^\infty$) for the finite (resp. infinite) outcomes, and $\text{Out}_{\mathcal{G}}(s, \sigma_P)$, $\text{Out}_{\mathcal{G}}^f(s, \sigma_P)$ and $\text{Out}_{\mathcal{G}}^\infty(s, \sigma_P)$ for the respective sets of outcomes of σ_P with initial state s . Notice that any strategy profile has a single infinite outcome from a given state. In the sequel, when given a strategy profile σ_{Agt} , we identify $\text{Out}(s, \sigma_{\text{Agt}})$ with the unique path it contains.

A concurrent game involving only two players (A and B , say) is *zero-sum* if, for any two plays π and π' , it holds $\pi \lesssim_A \pi'$ if, and only if, $\pi' \lesssim_B \pi$. Such a setting is purely antagonistic, as both players have opposite objectives. The most relevant concept in such a setting is that of *winning strategies*, where the aim is for one player to achieve her objectives *whatever the other players do*. In *non-zero-sum* games, winning strategies are usually too restricted, and the most relevant concepts are *equilibria*, which correspond to strategies that *satisfy* (which can be given several meanings) all the players. One of the most studied notion of equilibria is *Nash equilibria* [35], which we now introduce.

2.3. Nash equilibria. We begin with introducing some vocabulary. When $\pi \lesssim_A \pi'$, we say that π' is *at least as good* as π for A . We say that a strategy σ_A for A *ensures* π if every outcome of σ_A is at least as good as π for A , and that A *can ensure* π when such a strategy exists.

Given a move m_{Agt} and an action m' for some player A , we write $m_{\text{Agt}}[A \mapsto m']$ for the move n_{Agt} with $n_B = m_B$ when $B \neq A$ and $n_A = m'$. This is extended to strategies in the natural way.

Definition 2.4. Let \mathcal{G} be a concurrent game and let s be a state of \mathcal{G} . A *Nash equilibrium* of \mathcal{G} from s is a strategy profile $\sigma_{\text{Agt}} \in \text{Prof}_{\mathcal{G}}$ such that $\text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma']) \lesssim_A \text{Out}(s, \sigma_{\text{Agt}})$ for all players $A \in \text{Agt}$ and all strategies $\sigma' \in \text{Strat}^A$.

So, Nash equilibria are strategy profiles where no single player has an incentive to unilaterally deviate from her strategy.

Remark 2.5. Our definition of a Nash equilibrium requires any deviation to be worse or equivalent to the equilibrium. Another possible definition would have been to ask any deviation to be no better than the equilibrium. Those two definitions yield different notions of Nash equilibria (unless the preorders are total), as illustrated in Figure 4: the black node n represents $\text{Out}(s, \sigma_{\text{Agt}})$, the light-gray area contains the nodes n' such that $n' \lesssim n$, while the dark-gray area contains the nodes n' for which $n \not\lesssim n'$.

This alternative definition would also be meaningful, and the techniques we develop in this paper could be adapted to handle such a variant.

In this paper we will give a general construction that relates Nash equilibria in a game (which can be infinite) and winning strategies in a two-player turn-based game (called the *suspect game*), it is presented in Section 4. We will then be mostly interested in solving the decision problems that we define next, when games are finite.

2.4. Decision problems we will consider. Given a concurrent game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$ and a state $s \in \text{States}$, we consider the following problems:

- *Value problem:* Given a player A and a play π , is there a strategy σ_A for player A such that for any outcome ρ in \mathcal{G} from s of σ_A , it holds $\pi \lesssim_A \rho$?
- *NE Existence problem:* Does there exist a Nash equilibrium in \mathcal{G} from s ?
- *Constrained NE existence problem:* Given two plays π_A^- and π_A^+ for each player A , does there exist a Nash equilibrium in \mathcal{G} from s whose outcome π satisfies $\pi_A^- \lesssim_A \pi \lesssim_A \pi_A^+$ for all $A \in \text{Agt}$?

We will focus on decidability and complexity results of these three problems when games are finite, for various classes of preference relations. Complexity results will heavily rely on what preorders we allow for the preference relation and how they are represented. We have already discussed the representation of the game structure in Remark 2.2. We define and discuss now the various preference relations we will study, and explain how we encode the various inputs to the problems.

2.5. Focus on the preference relations we will consider. We define the various classes of preference relations we will focus on in the rest of the paper. We begin with single-objective preference relations, and we then define a more general class of ordered objectives. We fix a game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$.




FIGURE 4. Two different notions of *improvements* for a non-total order.

2.5.1. Single-objective preference relations.

Definition 2.6. An *objective* (or *winning condition*) is an arbitrary set of plays. A preference relation \preceq_A is *single-objective* whenever there exists an objective Ω_A such that: $\rho \preceq_A \rho'$ if, and only if, $\rho' \in \Omega_A$ (we then say that ρ' is winning for A) or $\rho \notin \Omega_A$ (we then say that ρ is losing for A).

The setting of single-objective preference relations is purely qualitative, since a player can only win (in case the outcome is in her objective), or lose (otherwise).

An objective Ω can be specified in various ways. Next we will consider the following families of ω -regular objectives:

- A **reachability objective** is given by a target set $T \subseteq \text{States}$ and the corresponding set of winning plays is defined by

$$\Omega_T^{\text{Reach}} = \{\rho \in \text{Play} \mid \text{Occ}(\rho) \cap T \neq \emptyset\}.$$

- A **safety objective** is given by a target set $T \subseteq \text{States}$ and the corresponding set of winning plays is defined by

$$\Omega_T^{\text{Safety}} = \{\rho \in \text{Play} \mid \text{Occ}(\rho) \cap T = \emptyset\}.$$

- A **Büchi objective** is given by a target set $T \subseteq \text{States}$ and the corresponding set of winning plays is defined by

$$\Omega_T^{\text{Büchi}} = \{\rho \in \text{Play} \mid \text{Inf}(\rho) \cap T \neq \emptyset\}.$$

- A **co-Büchi objective** is given by a target set $T \subseteq \text{States}$ and the corresponding set of winning plays is defined by

$$\Omega_T^{\text{co-Büchi}} = \{\rho \in \text{Play} \mid \text{Inf}(\rho) \cap T = \emptyset\}.$$

- A **parity objective** is given by a priority function $p: \text{States} \mapsto \llbracket 0, d \rrbracket$ (where $\llbracket 0, d \rrbracket = [0, d] \cap \mathbb{Z}$) with $d \in \mathbb{N}$, and the corresponding set of winning plays is defined by

$$\Omega_p^{\text{Parity}} = \{\rho \in \text{Play} \mid \min(\text{Inf}(p(\rho))) \text{ is even}\}.$$

- A **Streett objective** is given by a tuple $(Q_i, R_i)_{i \in \llbracket 1, k \rrbracket}$ and the corresponding set of winning plays is defined by

$$\Omega_{(Q_i, R_i)_{i \in \llbracket 1, k \rrbracket}}^{\text{Streett}} = \{\rho \in \text{Play} \mid \forall i. \text{Inf}(\rho) \cap Q_i \neq \emptyset \Rightarrow \text{Inf}(\rho) \cap R_i \neq \emptyset\}.$$

- A **Rabin objective** is given by a tuple $(Q_i, R_i)_{i \in \llbracket 1, k \rrbracket}$ and the corresponding set of winning plays is defined by

$$\Omega_{(Q_i, R_i)_{i \in \llbracket 1, k \rrbracket}}^{\text{Rabin}} = \{\rho \in \text{Play} \mid \exists i. \text{Inf}(\rho) \cap Q_i \neq \emptyset \wedge \text{Inf}(\rho) \cap R_i = \emptyset\}.$$

- A **Muller objective** is given by a finite set C , a coloring function $c: \text{States} \mapsto C$, and a set $\mathcal{F} \subseteq 2^C$. The corresponding set of winning plays is then defined by

$$\Omega_{c, \mathcal{F}}^{\text{Muller}} = \{\rho \in \text{Play} \mid \text{Inf}(c(\rho)) \in \mathcal{F}\}.$$

We will also consider the following other types of objectives:




FIGURE 5. Boolean circuit defining the condition that either ℓ_3 appears infinitely often, or if ℓ_1 appears infinitely often then so does ℓ_2 .

- A **circuit objective** is given by a boolean circuit C with the set States as input nodes and one output node. A play ρ is winning if and only if C evaluates to true when the input nodes corresponding to states in $\text{Inf}(\rho)$ are set to **true**, and all other input nodes are set to **false**. We write $\Omega_C^{\text{Circuit}}$ for the set of winning plays.

Figure 5 displays an example of a circuit for the game of Figure 3: this Boolean circuit defines the condition that either ℓ_3 appears infinitely often, or if ℓ_1 appears infinitely often then so does ℓ_2 .

- A **deterministic Büchi automaton objective** is given by a deterministic Büchi automaton $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, R \rangle$, with $\Sigma = \text{States}$. Then the corresponding set of winning plays is defined by

$$\Omega_{\mathcal{A}}^{\text{det-Büchi-aut}} = \mathcal{L}(\mathcal{A}).$$

- A **deterministic Rabin automaton objective** is given by a deterministic Rabin automaton $\mathcal{A} = \langle Q, \Sigma, \delta, q_0, (E_i, F_i)_{i \in [1, k]} \rangle$, with $\Sigma = \text{States}$. Then the corresponding set of winning plays is defined by

$$\Omega_{\mathcal{A}}^{\text{det-Rabin-aut}} = \mathcal{L}(\mathcal{A}).$$

- A **Presburger-definable objective** is given by a Presburger formula ϕ with free variables $(X_s)_{s \in \text{States}}$. The corresponding set of winning plays is defined by

$$\Omega_{\phi}^{\text{Presb}} = \{\rho \in \text{Play} \mid \phi(\#s(\rho))_{s \in \text{States}} = 0\}$$

where $\#s(\rho)$ is the number of occurrences³ of state s along ρ .

Encodings. For complexity issues we now make explicit how the various objectives are encoded:

- Reachability, safety, Büchi and co-Büchi objectives are given by a set $T \subseteq \text{States}$, they can therefore be encoded using $|\text{States}|$ bits.
- For parity objectives, we assume without loss of generality that $d \leq 2 \cdot |\text{States}|$. The priority function has then size at most $|\text{States}| \cdot \lceil \log(2 \cdot |\text{States}| + 1) \rceil$.
- Street and Rabin objectives are given by tuples $(Q_i, R_i)_{i \in [1, k]}$. Their sizes are given by: $\sum_{i \in [1, k]} |Q_i| \lceil \log(|\text{States}|) \rceil$.
- Muller objectives are given by a coloring function and a set \mathcal{F} . Its size is $|\text{States}| \cdot \lceil \log(|C|) \rceil + |\mathcal{F}| \cdot \lceil \log(|C|) \rceil$. Note that thanks to the coloring function, this encoding can be exponentially more succinct than an explicit representation such as the one considered in [26].

³By convention, if $s \in \text{Inf}(\rho)$, and variable X_s appears in ϕ , then $\rho \notin \Omega_{\phi}^{\text{Presb}}$.

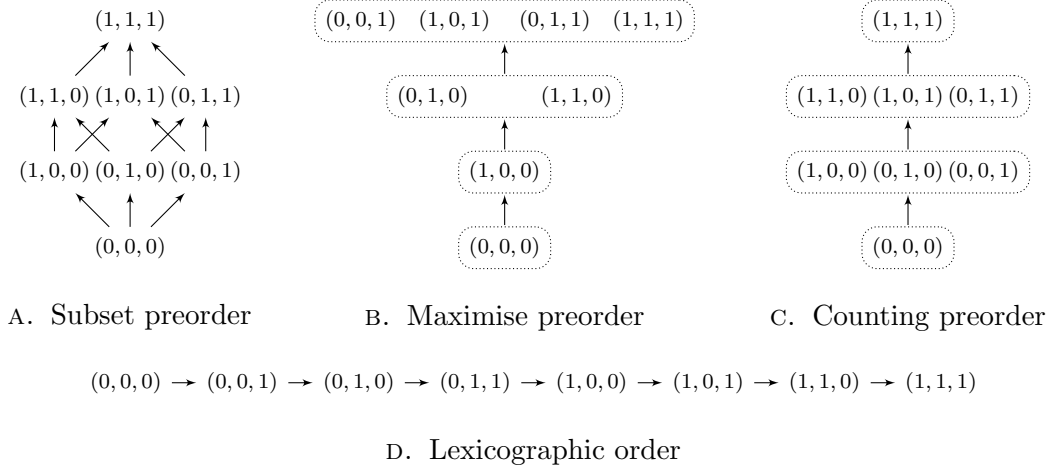


FIGURE 6. Examples of preorders (for $n = 3$): dotted boxes represent equivalence classes for the relation \sim , defined as $a \sim b \Leftrightarrow a \lesssim b \wedge b \lesssim a$; arrows represent the preorder relation \lesssim quotiented by \sim .

- The size of objectives given by circuits, deterministic automata or Presburger formulas is that of the corresponding circuits, deterministic automata or Presburger formulas.

Encodings of thresholds in inputs of the value and the constrained NE existence problems. For all the objectives except for those given by automata, whether a play ρ satisfies the objective or not only depends on the sets $\text{Occ}(\rho)$ and $\text{Inf}(\rho)$. The various thresholds will therefore be encoded as such pairs (Occ, Inf) .

For deterministic-automata objectives, the thresholds will be also encoded as pairs of sets of states of the objectives, representing respectively the set of states which are visited and the set of states which are visited infinitely often.

For the Boolean circuit objectives, whether a play ρ satisfies the objective or not only depends on the set $\text{Inf}(\rho)$. Therefore we will use as encoding for the threshold a single set Inf .

For the Presburger formulas objectives, we will use as encoding for the thresholds the Parikh image of the play (i.e., the number of visits to each of the states).

2.5.2. Ordered objectives. We now turn to a more general class of preference relations, allowing for a *semi-quantitative setting*.

Definition 2.7. An *ordered objective* is a pair $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$, where, for every $1 \leq i \leq n$, Ω_i is an objective, and \lesssim is a preorder on $\{0, 1\}^n$. A play ρ is assigned a *payoff vector* w.r.t. that ordered objective, which is defined as $\text{payoff}_\omega(\rho) = \mathbf{1}_{\{i | \rho \in \Omega_i\}} \in \{0, 1\}^n$ (where $\mathbf{1}_S$ is the vector v such that $v_i = 1 \Leftrightarrow i \in S$). The corresponding preference relation \lesssim_ω is then defined by $\rho \lesssim_\omega \rho'$ if, and only if, $\text{payoff}_\omega(\rho) \lesssim \text{payoff}_\omega(\rho')$.

There are many ways of specifying a preorder. We define below the preorders on $\{0, 1\}^n$ that we consider in the sequel. Figure 6 displays four such preorders for $n = 3$. For the purpose of these definitions, we assume that $\max \emptyset = -\infty$.

- (1) *Conjunction*: $v \lesssim w$ if, and only if, either $v_i = 0$ for some $1 \leq i \leq n$, or $w_i = 1$ for all $1 \leq i \leq n$. This corresponds to the case where a player wants to achieve all her objectives.
- (2) *Disjunction*: $v \lesssim w$ if, and only if, either $v_i = 0$ for all $1 \leq i \leq n$, or $w_i = 1$ for some $1 \leq i \leq n$. The aim here is to satisfy at least one objective.
- (3) *Counting*: $v \lesssim w$ if, and only if, $|\{i \mid v_i = 1\}| \leq |\{i \mid w_i = 1\}|$. The aim is to maximise the number of conditions that are satisfied;
- (4) *Subset*: $v \lesssim w$ if, and only if, $\{i \mid v_i = 1\} \subseteq \{i \mid w_i = 1\}$: in this setting, a player will always struggle to satisfy a larger (for inclusion) set of objectives.
- (5) *Maximise*: $v \lesssim w$ if, and only if, $\max\{i \mid v_i = 1\} \leq \max\{i \mid w_i = 1\}$. The aim is to maximise the highest index of the objectives that are satisfied.
- (6) *Lexicographic*: $v \lesssim w$ if, and only if, either $v = w$, or there is $1 \leq i \leq n$ such that $v_i = 0$, $w_i = 1$ and $v_j = w_j$ for all $1 \leq j < i$.
- (7) *Boolean Circuit*: given a Boolean circuit, with input from $\{0, 1\}^{2n}$, $v \lesssim w$ if, and only if, the circuit evaluates 1 on input $v_1 \dots v_n w_1 \dots w_n$.
- (8) *Monotonic Boolean Circuit*: same as above, with the restriction that the input gates corresponding to v are negated, and no other negation appear in the circuit.

In terms of expressiveness, any preorder over $\{0, 1\}^n$ can be given as a Boolean circuit: for each pair (v, w) with $v \lesssim w$, it is possible to construct a circuit whose output is 1 if, and only if, the input is $v_1 \dots v_n w_1 \dots w_n$; taking the disjunction of all these circuits we obtain a Boolean circuit defining the preorder. Its size can be bounded by $2^{2n+3}n$, which is exponential in general. But all the above examples ((1)-(6)) can be specified with a circuit of polynomial size. In Figure 7 we give a polynomial-size Boolean circuit for the subset preorder. In the following, for complexity issues, we will assume that the encoding of all preorders (1)-(6) takes constant size, and that the size of the preorder when it is given as a Boolean circuit is precisely the size of the circuit for input size n , where n is the number of objectives.

A preorder \lesssim is *monotonic* if it is compatible with the subset ordering, i.e. if $\{i \mid v_i = 1\} \subseteq \{i \mid w_i = 1\}$ implies $v \lesssim w$. Hence, a preorder is monotonic if fulfilling more objectives never results in a lower payoff. All our examples of preorders except for the Boolean circuit preorder are monotonic. Moreover, any monotonic preorder can be expressed as a monotonic Boolean circuit: for a pair (v, w) with $v \lesssim w$, we can build a circuit whose output is 1 if, and only if, the input is $v_1 \dots v_n w_1 \dots w_n$. We can require this circuit to have negation at the leaves. Indeed, if the input w_j appears negated, and if $w_j = 0$, then by monotonicity, also the input (v, \tilde{w}) is accepted, with $\tilde{w}_i = w_i$ when $i \neq j$ and $\tilde{w}_j = 1$. Hence the negated input gate can be replaced with **true**. Similarly for positive occurrences of any v_j . Hence any monotonic preorder can be written as a monotonic Boolean circuit. Notice that with Definition 2.4, any Nash equilibrium σ_{Agt} for the subset preorder is also a Nash equilibrium for any monotonic preorder.

Next we will be interested in two kinds of ordered objectives, *ordered reachability objectives*, where all objectives are supposed to be reachability objectives, and *ordered Büchi objectives*, where all objectives are supposed to be Büchi objectives. Note that other classical objectives (parity, Streett, Rabin, Muller, etc.) can be equivalently described with a preorder given by a polynomial-size Boolean circuit over Büchi objectives. For instance, each set of a Muller condition can be encoded as a conjunction of Büchi and co-Büchi conditions.

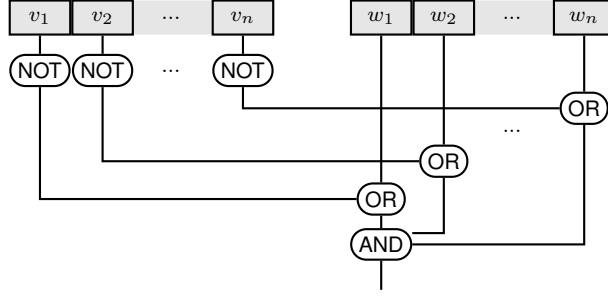


FIGURE 7. Boolean circuit defining the subset preorder

For ordered reachability (resp. Büchi) objectives, thresholds used as inputs to the various decision problems will be given by the set of states that are visited (resp. visited infinitely often).

In Sections 6 and 7, we will be interested in games where, for every player A , the preference relation \lesssim_A is given by an ordered objective $\omega_A = \langle (\Omega_i^A)_{1 \leq i \leq n_A}, \lesssim_A \rangle$. We will then write payoff_A instead of payoff_{ω_A} for the payoffs, and if ρ is a play, $\text{payoff}(\rho) = (\text{payoff}_A(\rho))_{A \in \text{Agt}}$.

2.6. Undecidability of all three problems for single Presburger-definable objectives. We end this section with an undecidability result in the quite general setting of Presburger-definable preference relations.

Theorem 2.8. *The value, NE existence and constrained NE existence problems are undecidable for finite games with preference relations given by Presburger-definable qualitative objectives.*

Proof. We first prove the result for the constrained NE existence problem, by encoding a two-counter machine. We fix a two-counter machine, and assume without loss of generality that the halting state is preceded by a non-zero test for the two counters (hence if the machine halts, the two counters have a positive value in the halting state).

We begin with defining a family of preorders. Fix two sets of states S and T ; a play is said $(S = T)$ -winning if the number of visits to S equals the number of visits to T , and both are finite. Formally, $\pi \lesssim_{S=T} \pi'$ whenever π is not $(S = T)$ -winning, or π' is.

We use such preorders to encode the acceptance problem for two-counter machines: the value of counter c_1 is encoded as the difference between the number of visits to S_1 and T_1 , and similarly for counter c_2 . Incrementing counter c_i thus consists in visiting a state in S_i , and decrementing consists in visiting T_i ; in other terms, if instruction q_k of the two-counter machine consists in incrementing c_1 and jumping to $q_{k'}$, then the game will have a transition from some state q_k to a state in S_1 , and a transition from there to $q_{k'}$. The game involves three players: A_1 , A_2 and B . The aim of player A_1 (resp. A_2) is to visit S_1 and T_1 (resp. S_2 and T_2) the same number of times: player A_i 's preference is $\lesssim_{S_i=T_i}$. The aim of player B is to reach the state corresponding to the halting state of the two-counter machine. Due to the assumption on the two-counter machine, if B wins, then both A_1 and A_2 lose.

It remains to encode the zero-test: this is achieved by the module of Figure 8. In this module, player B tries to avoid the three sink states (marked in grey), since this would prevent her from reaching her goal. When entering the module, player B has to choose one of the available branches: if she decides to go to $u_i^{\neq 0}$, then A_i could take the play into the self-loop, which is winning for her if S_i and T_i have been visited the same number of times in the history of this path, which corresponds to having $c_i = 0$; hence player B should play to $u_i^{\neq 0}$ only if $c_i \neq 0$, so that A_i has no interest in going to this self-loop.

Similarly, if player B decides to go to $u_i^{=0}$, player A_i has the opportunity to “leave” the main stream of the game, and go to s_i or t_i (obviously $s_i \in S_i$ and $t_i \in T_i$). If the numbers of visits to S_i and T_i up to that point are different, then player A_i has the opportunity to make both numbers equal, and to win. Conversely, if both numbers are equal (i.e., $c_i = 0$), then going to s_i or t_i will be losing for A_i , whatever happens from there. Hence, if $c_i = 0$ when entering the module, then player B should go to $u_i^{=0}$.




FIGURE 8. Testing whether $c_i = 0$.

One can then easily show that the two-counter machine stops if, and only if, there is a Nash equilibrium in the resulting game \mathcal{G} , in which player B wins and players A_1 and A_2 lose. Indeed, assume that the machine stops, and consider the strategies where player B plays (in the first state of the test modules) according to the value of the corresponding counter, and where players A_1 and A_2 always keep the play in the main stream of the game. Since the machine stops, player B wins, while players A_1 and A_2 lose. Moreover, none of them has a way to improve their payoff: since player B plays according to the values of the counters, players A_1 and A_2 would not benefit from deviating from their above strategies. Conversely, if there is such a Nash equilibrium, then in any visited test module, player B always plays according to the values of the counters: otherwise, player A_1 (or A_2) would have the opportunity to win the game. By construction, this means that the run of the Nash equilibrium corresponds to the execution of the two-counter machine. As player B wins, this execution reaches the halting state.

Finally, it is not difficult to adapt this reduction to involve only two players: players A_1 and A_2 would be replaced by one single player A , in charge of ensuring that both conditions (for c_1 and c_2) are fulfilled. This requires minor changes to the module for testing $c_i = 0$: when leaving the main stream of the game in a module for testing counter c_i , player A should be given the opportunity (after the grey state) to visit states S_{3-i} or T_{3-i} in order to adjust that part of her objective.

By changing the winning condition for Player B , the game \mathcal{G} can also be made zero-sum: for this, B must lose if the play remains in the main stream forever without visiting the final state; otherwise, B loses if the number of visits to s_i and t_i are finite and equal for both $i = 1$ and $i = 2$; B wins in any other case. The objective of player A is opposite. It is not difficult to modify the proof above for showing that the two-counter machine halts if, and only if, player B has a winning strategy in this game.

Finally, by adding a small initial module depicted on Figure 9 to this zero-sum version of the game \mathcal{G} , one can encode the halting problem for two-counter machines to the NE existence problem. Indeed, in the zero-sum game, there is exactly one Nash equilibrium, with only two possible payoffs (either A wins, or B wins). Now, assuming that A loses and B wins in state s_1 , then there is a (pure) Nash equilibrium in the game extended with the initial module if, and only if, player B wins in the zero-sum game above. \square




FIGURE 9. Extending the game with an initial concurrent module

3. PRELIMINARY RESULTS

This section contains general results that will be applied later in various settings. In each of the statements, we give the restrictions on the games and on the preference relations that should be satisfied.

3.1. Nash equilibria as lasso runs. We first characterise outcomes of Nash equilibria as ultimately periodic runs, in the case where preference relations only depend on the set of states that are visited, and on the set of states that are visited infinitely often. Note that ω -regular conditions satisfy this hypothesis, but Presburger relations such as the ones used for proving Theorem 2.8 do not.

Proposition 3.1. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ be a **finite concurrent game** such that, for every player A , it holds⁴ $\rho \sim_A \rho'$ as soon as $\text{Inf}(\rho) = \text{Inf}(\rho')$ and $\text{Occ}(\rho) = \text{Occ}(\rho')$. Let $\rho \in \text{Play}$. If there is a Nash equilibrium with outcome ρ , then there is a Nash equilibrium with outcome ρ' of the form $\pi \cdot \tau^\omega$ such that $\rho \sim_A \rho'$, and where $|\pi|$ and $|\tau|$ are bounded by $|\text{States}|^2$.*

Proof. Let σ_{Agt} be a Nash equilibrium from some state s , and ρ be its outcome. We define a new strategy profile σ'_{Agt} , whose outcome from s is ultimately periodic, and then show that σ'_{Agt} is a Nash equilibrium from s .

To begin with, we inductively construct a history $\pi = \pi_0 \pi_1 \dots \pi_n$ that is not too long and visits precisely those states that are visited by ρ (that is, $\text{Occ}(\pi) = \text{Occ}(\rho)$).

The initial state is $\pi_0 = \rho_0 = s$. Then we assume we have constructed $\pi_{\leq k} = \pi_0 \dots \pi_k$ which visits exactly the same states as $\rho_{\leq k'}$ for some k' . If all the states of ρ have been visited in $\pi_{\leq k}$ then the construction is over. Otherwise there is an index i such that ρ_i does not appear in $\pi_{\leq k}$. We therefore define our next target as the smallest such i : we let $t(\pi_{\leq k}) = \min\{i \mid \forall j \leq k. \pi_j \neq \rho_i\}$. We then look at the occurrence of the current state π_k that is the closest to the target in ρ : we let $c(\pi_{\leq k}) = \max\{j < t(\pi_{\leq k}) \mid \pi_k = \rho_j\}$. Then we emulate what happens at that position by choosing $\pi_{j+1} = \rho_{c(\pi_{\leq j})+1}$. Then π_{k+1} is either the target, or a state that has already been seen before in $\pi_{\leq k}$, in which case the resulting $\pi_{\leq k+1}$ visits exactly the same states as $\rho_{\leq c(\pi_{\leq k})+1}$.

At each step, either the number of remaining targets strictly decreases, or the number of remaining targets is constant but the distance to the next target strictly decreases. Therefore

⁴We recall that $\rho \sim_A \rho'$ if, and only if, $\rho \preceq_A \rho'$ and $\rho' \preceq_A \rho$.

the construction terminates. Moreover, notice that between two targets we do not visit the same state twice, and we visit only states that have already been visited, plus the target. As the number of targets is bounded by $|\text{States}|$, we get that the length of the path π constructed thus far is bounded by $1 + |\text{States}| \cdot (|\text{States}| - 1)/2$.

Using similar ideas, we now inductively construct $\tau = \tau_0\tau_1 \dots \tau_m$, which visits precisely those states which are seen infinitely often along ρ , and which is not too long. Let l be the least index after which the states visited by ρ are visited infinitely often, i.e. $l = \min\{i \in \mathbb{N} \mid \forall j \geq i. \rho_j \in \text{Inf}(\rho)\}$. The run $\rho_{\geq l}$ is such that its set of visited states and its set of states visited infinitely often coincide. We therefore define τ in the same way we have defined π above, but for play $\rho_{\geq l}$. As a by-product, we also get $c(\tau_{\leq k})$, for $k < m$.

We now need to glue π and τ together, and to ensure that τ can be glued to itself, so that $\pi \cdot \tau^\omega$ is a real run. We therefore need to link the last state of π with the first state of τ (and similarly the last state of τ with its first state). This possibly requires appending some more states to π and τ : we fix the target of π and τ to be τ_0 , and apply the same construction as previously until the target is reached. The total length of the resulting paths π' and τ' is bounded by $1 + (|\text{States}| - 1) \cdot (|\text{States}| + 2)/2$ which is less than $|\text{States}|^2$.

We let $\rho' = \pi' \cdot \tau'^\omega$, and abusively write $c(\rho'_{\leq k})$ for $c(\pi'_{\leq k})$ if $k \leq |\pi'|$ and $c(\tau'_{\leq k'})$ with $k' = (k - 1 - |\pi'|) \bmod |\tau'|$ otherwise. We now define our new strategy profile, having ρ' as outcome from s . Given a history h :

- if h followed the expected path, i.e., $h = \rho'_{\leq k}$ for some k , we mimic the strategy at $c(h)$: $\sigma'_{\text{Agt}}(h) = \sigma_{\text{Agt}}(\rho'_{c(h)})$. This way, ρ' is the outcome of σ'_{Agt} from s .
- otherwise we take the longest prefix $h_{\leq k}$ that is a prefix of ρ' , and define $\sigma'_{\text{Agt}}(h) = \sigma_{\text{Agt}}(\rho'_{c(h_{\leq k})} \cdot h_{\geq k+1})$.

We now show that σ'_{Agt} is a Nash equilibrium. Assume that one of the players changes her strategy while playing according to σ'_{Agt} : either the resulting outcome does not deviate from $\pi \cdot \tau^\omega$, in which case the payoff of that player is not improved; or it deviates at some point, and from that point on, σ'_{Agt} follows the same strategies as in σ_{Agt} . Assume that the resulting outcome is an improvement over ρ' for the player who deviated. The suffix of the play after the deviation is the suffix of a play of σ_{Agt} after a deviation by the same player. By construction, both plays have the same sets of visited and infinitely-visited states. Hence we have found an advantageous deviation from σ_{Agt} for one player, contradicting the fact that σ_{Agt} is a Nash equilibrium. \square

3.2. Encoding the value problem as a constrained NE existence problem. We now give a reduction that will be used to infer hardness results for the constrained NE existence problem from the hardness of the value problem (as defined in Section 2.4): this will be the case when the hardness proof for the value problem involves the construction of a game satisfying the hypotheses of the proposition.

Proposition 3.2. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ be a two-player zero-sum game played between players A and B , such that:*

- *the preference relation \preceq_A for player A is total, Noetherian and almost-well-founded (see Section 2.1);*
- *\mathcal{G} is determined, i.e., for all play π :*

$$[\exists \sigma_A. \forall \sigma_B. \pi \preceq_A \text{Out}(\sigma_A, \sigma_B)] \Leftrightarrow [\forall \sigma_B. \exists \sigma_A. \pi \preceq_A \text{Out}(\sigma_A, \sigma_B)].$$

Let \mathcal{G}' be the (non-zero-sum) game obtained from \mathcal{G} by replacing the preference relation of player B by the one where all plays are equivalent. Then, for every state s , for every play π from s , the two following properties are equivalent:

- (i) there is a Nash equilibrium in \mathcal{G}' from s with outcome ρ such that $\pi \not\lesssim_A \rho$;
- (ii) player A cannot ensure π from s in \mathcal{G} .

Proof. In this proof, σ_A and σ'_A (resp. σ_B and σ'_B) refer to player- A (resp. player- B) strategies. Furthermore we will write $\text{Out}(\sigma_A, \sigma_B)$ instead of $\text{Out}_{\mathcal{G}}(s, (\sigma_A, \sigma_B))$.

We first assume there is a Nash equilibrium (σ_A, σ_B) in \mathcal{G}' from s such that $\pi \not\lesssim_A \text{Out}(\sigma_A, \sigma_B)$. Since \lesssim_A is total, $\text{Out}(\sigma_A, \sigma_B) \prec_A \pi$. Consider a strategy σ'_A of player A in \mathcal{G} . As (σ_A, σ_B) is a Nash equilibrium, it holds that $\text{Out}(\sigma'_A, \sigma_B) \lesssim_A \text{Out}(\sigma_A, \sigma_B)$, which implies $\text{Out}(\sigma'_A, \sigma_B) \prec_A \pi$. We conclude that condition (ii) holds.

Assume now property (ii). As the preference relation is Noetherian, we can select π^+ which is the largest element for \lesssim_A which can be ensured by player A . Let σ_A be a corresponding strategy: for every strategy σ_B , $\pi^+ \lesssim_A \text{Out}(\sigma_A, \sigma_B)$. Towards a contradiction, assume now that for every strategy σ'_B , there exists a strategy σ'_A such that $\pi^+ \prec_A \text{Out}(\sigma'_A, \sigma'_B)$. Consider the set S of such outcomes, and define π' as its minimal element (this is possible since the order \lesssim_A is almost-well-founded). Notice then that $\pi^+ \prec_A \pi'$, and also that for every strategy σ'_B , there exists a strategy σ'_A such that $\pi' \lesssim_A \text{Out}(\sigma'_A, \sigma'_B)$. Then, as the game is determined, we get that there exists some strategy σ'_A such that for all strategy σ'_B , it holds that $\pi' \lesssim_A \text{Out}(\sigma'_A, \sigma'_B)$. In particular, strategy σ'_A ensures π' , which contradicts the maximality of π^+ . Therefore, there is some strategy σ'_B for which for every strategy σ'_A , $\pi^+ \not\lesssim_A \text{Out}(\sigma'_A, \sigma'_B)$, which means $\text{Out}(\sigma'_A, \sigma'_B) \lesssim_A \pi^+$. We show now that (σ_A, σ'_B) is a witness for property (i). We have seen on the one hand that $\pi^+ \lesssim_A \text{Out}(\sigma_A, \sigma'_B)$, and on the other hand that $\text{Out}(\sigma_A, \sigma'_B) \lesssim_A \pi^+$. By hypothesis, $\pi^+ \prec_A \pi$, which yields $\text{Out}(\sigma_A, \sigma'_B) \prec_A \pi$. Pick another strategy σ'_A for player A . We have seen that $\text{Out}(\sigma'_A, \sigma'_B) \lesssim_A \pi^+$, which implies $\text{Out}(\sigma'_A, \sigma'_B) \lesssim_A \text{Out}(\sigma_A, \sigma'_B)$. This concludes the proof of (i). \square

Remark 3.3. Any finite total preorder is obviously Noetherian and almost-well-founded. Also, any total preorder isomorphic to the set of non-positive integers is Noetherian and almost-well-founded. On the other hand, a total preorder isomorphic to $\{1/n \mid n \in \mathbb{N}^+\}$ is Noetherian but not almost-well-founded.

3.3. Encoding the value problem as a NE existence problem. We prove a similar result for the NE existence problem. In this reduction however, we have to modify the game by introducing a truly concurrent move at the beginning of the game. This is necessary since for turn-based games with ω -regular winning conditions, there always exists a Nash equilibrium [15], hence the NE existence problem would be trivial.

Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$ be a two-player zero-sum game, with players A and B . Given a state s of \mathcal{G} and a play π from s , we define a game \mathcal{G}_π by adding two states s_0 and s_1 , in the very same way as in Figure 9, on page 15. From s_0 , A and B play a matching-penny game to either go to the sink state s_1 , or to the state s in the game \mathcal{G} . We assume the same hypotheses than in Proposition 3.2 for the preference relation \lesssim_A . Let π^+ be in the highest equivalence class for \lesssim_A smaller than π (it exists since \lesssim_A is Noetherian). In \mathcal{G}_π , player B prefers runs that end in s_1 : formally, the preference relation \lesssim_B^π of player B in \mathcal{G}_π is given by $\pi' \lesssim_B^\pi \pi'' \Leftrightarrow \pi'' = s_0 \cdot s_1^\omega \vee \pi' \neq s_0 \cdot s_1^\omega$. On the other hand,

player A prefers a path of \mathcal{G} over going to s_1 , if and only if, it is at least as good as π : formally, the preference relation \lesssim_A^π for player A in \mathcal{G}_π is given by $s_0 \cdot \pi' \lesssim_A^\pi s_0 \cdot \pi'' \Leftrightarrow \pi' \lesssim_A \pi''$, and $s_0 \cdot s_1^\omega \sim_A'' s_0 \cdot \pi^+$.

Proposition 3.4. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$ be a two-player zero-sum game, with players A and B , such that:*

- *the preference relation \lesssim_A for player A is total, Noetherian and almost-well-founded;*
- *\mathcal{G} is determined.*

Let s be a state and π be a play in \mathcal{G} from s . Consider the game \mathcal{G}_π defined above. Then the following two properties are equivalent:

- (i) *there is a Nash equilibrium in \mathcal{G}_π from s_0 ;*
- (ii) *player A cannot ensure π from s in \mathcal{G} .*

In particular, in a given class of games, if the hardness proof of the value problem involves a game which satisfies the hypotheses of the proposition, and if \mathcal{G}_π belongs to that class, then the NE existence problem is at least as hard as the complement of the value problem.

Proof. Assume that player A cannot ensure at least π from s in \mathcal{G} , then according to Proposition 3.2, there is a Nash equilibrium (σ_A, σ_B) in the game \mathcal{G}' of Proposition 3.2 with outcome ρ such that $\pi \not\lesssim_A \rho$. Consider the strategy profile $(\sigma_A^\pi, \sigma_B^\pi)$ in \mathcal{G}_π that consists in playing the same action for both players in s_0 , and then if the path goes to s , to play according to (σ_A, σ_B) . Player B gets her best possible payoff under that strategy profile. If A could change her strategy to get a payoff better than $s_0 \cdot \pi^+$, then it would induce a strategy in \mathcal{G}' giving her a payoff better than ρ (when played with strategy σ_B), which contradicts the fact that (σ_A, σ_B) is a Nash equilibrium in \mathcal{G}' . Therefore, $(\sigma_A^\pi, \sigma_B^\pi)$ is a Nash equilibrium in \mathcal{G}_π .

Conversely, assume that A can ensure π from s in \mathcal{G} , and assume towards a contradiction that there is a Nash equilibrium $(\sigma_A^\pi, \sigma_B^\pi)$ in \mathcal{G}_π from s_0 . Then $\text{Out}_{\mathcal{G}_\pi}(\sigma_A^\pi, \sigma_B^\pi)$ does not end in s_1 , otherwise player A could improve by switching to s and then playing according to a strategy which ensures π . Also, $\text{Out}_{\mathcal{G}_\pi}(\sigma_A^\pi, \sigma_B^\pi)$ cannot end in \mathcal{G} either, otherwise player B would improve by switching to s_1 . We get that there is no Nash equilibrium in \mathcal{G}_π from s_0 , which concludes the proof. \square

3.4. Encoding the constrained NE existence problem as an NE existence problem. The next proposition makes a link between the existence of a Nash equilibrium where a player gets a payoff larger than some bound and the (unconstrained) existence of a Nash equilibrium in a new game. This will allow, in some specific cases, to infer hardness results from the constrained NE existence problem to the NE existence problem.

The construction is inspired by the previous one, but it applies to a game with at least two players, and it applies to any two selected players as follows. Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$ be a concurrent game, s be a state of \mathcal{G} , ρ be a play from s , and A_i and A_j be two distinct players. We define the new game $E(\mathcal{G}, A_i, A_j, \rho)$ again in the same way as on Figure 9. Now, in s_0 , the two players A_i and A_j play a matching-penny game to either go to the sink state s_1 , or to state s in game \mathcal{G} .

For player A_j , the preference relation in $E(\mathcal{G}, A_i, A_j, \rho)$ is given by \lesssim'_{A_j} such that $s_0 \cdot s_1^\omega \prec'_{A_j} s_0 \cdot \pi$ and $s_0 \cdot \pi \lesssim'_{A_j} s_0 \cdot \pi' \Leftrightarrow \pi \lesssim_{A_j} \pi'$, for any path π and π' from s in \mathcal{G} . For player A_i the preference relation is $s_0 \cdot \pi \lesssim'_{A_i} s_0 \cdot \pi' \Leftrightarrow \pi \lesssim_{A_i} \pi'$, for any path π and

π' from s in \mathcal{G} , and $s_0 \cdot s_1^\omega \sim_{A_i} s_0 \cdot \rho$. For any other player A_k , the preference relation $E(\mathcal{G}, A_i, A_j, \rho)$ is given by $s_0 \cdot \pi \succ'_{A_k} s_0 \cdot \pi' \Leftrightarrow \pi \succ_{A_k} \pi'$ for any path π and π' from s in \mathcal{G} , and $s_0 \cdot s_1^\omega \sim_{A_k} s_0 \cdot \rho$.

Proposition 3.5. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\succ_A)_{A \in \text{Agt}} \rangle$ be a concurrent game, let s be a state of \mathcal{G} , and A_i and A_j be two distinct players participating to \mathcal{G} . Pick two plays π and ρ from s such that $\rho \succ_{A_i} \pi$. If there is a Nash equilibrium in \mathcal{G} whose outcome is π , then there is a Nash equilibrium in $E(\mathcal{G}, A_i, A_j, \rho)$ whose outcome is $s_0 \cdot \pi$. Reciprocally, if there is a Nash equilibrium in $E(\mathcal{G}, A_i, A_j, \rho)$ whose outcome is $s_0 \cdot \pi$, then there is a Nash equilibrium in \mathcal{G} whose outcome is π .*

Proof. Assume that there is a Nash equilibrium σ_{Agt} in \mathcal{G} with outcome π such that $\rho \succ_{A_i} \pi$. Then $s_0 \cdot s_1^\omega \succ'_{A_i} s_0 \cdot \pi$. Consider the strategy profile in $E(\mathcal{G}, A_i, A_j, \rho)$ that consists for A_i and A_j in playing different actions in s_0 and when the path goes to s , to play according to σ_{Agt} . Players A_i and A_j have no interest in changing their strategies in s_0 , since for A_j all plays of \mathcal{G} are better than $s_0 \cdot s_1^\omega$, and for A_i the play $s_0 \cdot \pi$ is better than $s_0 \cdot s_1^\omega$. Hence, this is a Nash equilibrium in game $E(\mathcal{G}, A_i, A_j, \rho)$.

Reciprocally, if there is a Nash equilibrium in $E(\mathcal{G}, A_i, A_j, \rho)$, its outcome cannot end in s_1 , since A_j would have an interest in changing her strategy in s_0 (all plays of \mathcal{G} are then better for her). The strategies followed from s thus defines a Nash equilibrium in \mathcal{G} . \square

If we consider a class of games such that $E(\mathcal{G}, A_i, A_j, \rho)$ belongs to that class when \mathcal{G} does, then the NE existence problem is then at least as hard as the constrained NE existence problem. Note however that the reduction assumes lower bounds on the payoffs, and we do not have a similar result for upper bounds on the payoffs. For instance, as we will see in Section 6, for a conjunction of Büchi objectives, we do not know whether the NE existence problem is in P (as the value problem) or NP-hard (as is the existence of an equilibrium where all the players are losing).

4. THE SUSPECT GAME

In this section, we construct an abstraction of a multi-player game \mathcal{G} as a two-player zero-sum game \mathcal{H} , such that there is a correspondence between Nash equilibria in \mathcal{G} and winning strategies in \mathcal{H} (formalised in forthcoming Theorem 4.5). This transformation does not require the game to be finite and is conceptually much deeper than the reductions given in the previous section; it will allow us to use algorithmic techniques from zero-sum games to compute Nash equilibria and hence solve the value and (constrained) NE existence problems in various settings.

4.1. Construction of the suspect game. We fix a concurrent game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\succ_A)_{A \in \text{Agt}} \rangle$ for the rest of the section, and begin with introducing a few extra definitions.

Definition 4.1. A strategy profile σ_{Agt} is a *trigger profile* for a play π from some state s if, for every player $A \in \text{Agt}$, for every strategy σ'_A of player A , the path π is at least as good as the outcome of $\sigma_{\text{Agt}}[A \mapsto \sigma'_A]$ from s (that is, $\text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A]) \succ_A \pi$).

The following result is folklore and a direct consequence of the definition:

Lemma 4.2. *A Nash equilibrium is a trigger profile for its outcome. Reciprocally, a strategy profile which is trigger profile for its outcome is a Nash equilibrium.*

Definition 4.3 ([4]). Given two states s and s' , and a move m_{Agt} , the set of *suspect players* for (s, s') and m_{Agt} is the set

$$\text{Susp}((s, s'), m_{\text{Agt}}) = \{A \in \text{Agt} \mid \exists m' \in \text{Mov}(s, A). \text{Tab}(s, m_{\text{Agt}}[A \mapsto m']) = s'\}.$$

Given a path ρ and a strategy profile σ_{Agt} , the set of suspect players for ρ and σ_{Agt} is the set of players that are suspect along each transition of ρ , i.e., it is the set

$$\text{Susp}(\rho, \sigma_{\text{Agt}}) = \left\{ A \in \text{Agt} \mid \forall i < |\rho|. A \in \text{Susp}((\rho_{=i}, \rho_{=i+1}), \sigma_{\text{Agt}}(\rho_{\leq i})) \right\}.$$

Intuitively, player $A \in \text{Agt}$ is a suspect for transition (s, s') and move m_{Agt} if she can unilaterally change her action to activate the transition (s, s') : if $s' \neq \text{Tab}(s, m_{\text{Agt}})$, then this may be due to a deviation from m_{Agt} of any of the players in the set $\text{Susp}((s, s'), m_{\text{Agt}})$, and no one else. If $s' = \text{Tab}(s, m_{\text{Agt}})$, it may simply be the case that no one has deviated, so everyone is a potential suspect for the next moves. Similarly, we easily infer that player A is in $\text{Susp}(\rho, \sigma_{\text{Agt}})$ if, and only if, there is a strategy σ'_A such that $\text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A]) = \rho$.

Note that the notion of suspect players requires moves and arenas to be deterministic, and therefore everything which follows assumes the restriction to pure strategy profiles and to deterministic game structures.

We fix a play π in \mathcal{G} . From game \mathcal{G} and play π , we build the *suspect game* $\mathcal{H}(\mathcal{G}, \pi)$, which is a two-player turn-based game defined as follows. The players in $\mathcal{H}(\mathcal{G}, \pi)$ are named **Eve** and **Adam**. Since $\mathcal{H}(\mathcal{G}, \pi)$ is turn-based, its state space can be written as the disjoint union of the set V_{\exists} controlled by **Eve**, which is (a subset of) $\text{States} \times 2^{\text{Agt}}$, and the set V_{\forall} controlled by **Adam**, which is (a subset of) $\text{States} \times 2^{\text{Agt}} \times \text{Act}^{\text{Agt}}$. The game is played in the following way: from a configuration (s, P) in V_{\exists} , **Eve** chooses a legal move m_{Agt} from s ; the next state is (s, P, m_{Agt}) ; then **Adam** chooses some state s' in States , and the new configuration is $(s', P \cap \text{Susp}((s, s'), m_{\text{Agt}}))$. In particular, when the state s' chosen by **Adam** is such that $s' = \text{Tab}(s, m_{\text{Agt}})$ (we say that **Adam obeys Eve** when this is the case), then the new configuration is (s', P) .

We define projections proj_1 and proj_2 from V_{\exists} on States and 2^{Agt} , resp., by $\text{proj}_1(s, P) = s$ and $\text{proj}_2(s, P) = P$. We extend these projections to paths in a natural way (but only using **Eve's** states in order to avoid stuttering), letting $\text{proj}_1((s_0, P_0) \cdot (s_0, P_0, m_0) \cdot (s_1, P_1) \cdots) = s_0 \cdot s_1 \cdots$. For any play ρ , $\text{proj}_2(\rho)$ (seen as a sequence of sets of players of \mathcal{G}) is non-increasing, therefore its limit $\lambda(\rho)$ is well defined. We notice that if $\lambda(\rho) \neq \emptyset$, then $\text{proj}_1(\rho)$ is a play in \mathcal{G} . An outcome ρ is *winning for Eve*, if for all $A \in \lambda(\rho)$, it holds $\text{proj}_1(\rho) \preceq_A \pi$. The *winning region* $W(\mathcal{G}, \pi)$ (later simply denoted by W when \mathcal{G} and π are clear from the context) is the set of configurations of $\mathcal{H}(\mathcal{G}, \pi)$ from which **Eve** has a winning strategy. Intuitively **Eve** tries to have the players play a Nash equilibrium, and **Adam** tries to disprove that it is a Nash equilibrium, by finding a possible deviation that improves the payoff of one of the players.

4.2. Correctness of the suspect-game construction. The next lemma establishes a correspondence between winning strategies in $\mathcal{H}(\mathcal{G}, \pi)$ and trigger profiles (and therefore Nash equilibria) in \mathcal{G} .

Lemma 4.4. *Let s be a state of \mathcal{G} and π be a play from s in \mathcal{G} . The following two conditions are equivalent:*

- *Eve* has a winning strategy in $\mathcal{H}(\mathcal{G}, \pi)$ from (s, Agt) , and its outcome ρ' from (s, Agt) when *Adam* obeys *Eve* is such that $\text{proj}_1(\rho') = \rho$;
- there is a trigger profile for π in \mathcal{G} from state s whose outcome from s is ρ .

Proof. Assume there is a winning strategy σ_{\exists} for *Eve* in $\mathcal{H}(\mathcal{G}, \pi)$ from (s, Agt) , whose outcome from (s, Agt) when *Adam* obeys *Eve* is ρ' with $\text{proj}_1(\rho') = \rho$. We define the strategy profile σ_{Agt} according to the actions played by *Eve*. Pick a history $g = s_1 s_2 \cdots s_{k+1}$, with $s_1 = s$. Let h be the outcome of σ_{\exists} from s ending in a state of V_{\exists} and such that $\text{proj}_1(h) = s_1 \cdots s_k$. This history is uniquely defined as follows: the first state of h is (s_1, Agt) , and if its $(2i + 1)$ -st state is (s_i, P_i) , then its $(2i + 2)$ -nd state is $(s_i, P_i, \sigma_{\exists}(h_{\leq 2i+1}))$ and its $(2i + 3)$ -rd state is $(s_{i+1}, P_i \cap \text{Susp}((s_i, s_{i+1}), \sigma_{\exists}(h_{\leq 2i+1})))$. Now, write (s_k, P_k) for the last state of h , and let $h' = h \cdot (s_k, P_k, \sigma_{\exists}(h)) \cdot (s_{k+1}, P_k \cap \text{Susp}((s_k, s_{k+1}), \sigma_{\exists}(h)))$. Then we define $\sigma_{\text{Agt}}(g) = \sigma_{\exists}(h')$. Notice that when $g \cdot s$ is a prefix of $\text{proj}_1(\rho')$, then $g \cdot s \cdot \sigma_{\text{Agt}}(g \cdot s)$ is also a prefix of $\text{proj}_1(\rho')$. In particular, $\text{Out}(s, \sigma_{\text{Agt}}) = \text{proj}_1(\rho') = \rho$.

We now prove that σ_{Agt} is a trigger profile for π . Pick a player $A \in \text{Agt}$, a strategy σ'_A for player A , and let $g = \text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A])$. With a play g , we associate a play h in $\mathcal{H}(\mathcal{G}, \pi)$ in the same way as above. Then player A is a suspect along all the transitions of g , so that she belongs to $\lambda(h)$. Now, as σ_{\exists} is winning, $\text{proj}_1(h) \lesssim_A \pi$, which proves that σ_{Agt} is a trigger profile.

Conversely, assume that σ_{Agt} is a trigger profile for π whose outcome is ρ , and define the strategy σ_{\exists} by $\sigma_{\exists}(h) = \sigma_{\text{Agt}}(\text{proj}_1(h))$. Notice that the outcome ρ' of σ_{\exists} when *Adam* obeys *Eve* satisfies $\text{proj}_1(\rho') = \rho$.

Let η be an outcome of σ_{\exists} from s , and $A \in \lambda(\eta)$. Then A is a suspect for each transition along $\text{proj}_1(\eta)$, which means that for all i , there is a move m_i^A such that

$$\text{proj}_1(\eta_{=i+1}) = \text{Tab}(\text{proj}_1(\eta_{=i}), \sigma_{\text{Agt}}(\text{proj}_1(\eta_{\leq i}))[A \mapsto m_i^A]).$$

Therefore there is a strategy σ'_A such that $\text{proj}_1(\eta) = \text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A])$. Since σ_{Agt} is a trigger profile for π , it holds that $\text{proj}_1(\eta) \lesssim_A \pi$. As this holds for any $A \in \lambda(\eta)$, σ_{\exists} is winning. \square

We now state the correctness theorem for the suspect game construction.

Theorem 4.5. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$ be a concurrent game, s be a state of \mathcal{G} , and π be a play in \mathcal{G} . The following two conditions are equivalent:*

- there is a Nash equilibrium σ_{Agt} from s in \mathcal{G} whose outcome is π .
- there is a play ρ from (s, Agt) in $\mathcal{H}(\mathcal{G}, \pi)$,
 - (1) such that $\text{proj}_1(\rho) = \pi$;
 - (2) along which *Adam* always obeys *Eve*; and
 - (3) such that for all indices i , there is a strategy σ_{\exists}^i for *Eve*, for which any play in $\rho_{\leq i} \cdot \text{Out}(\rho_{=i}, \sigma_{\exists}^i)$ is winning for *Eve*.

Proof. The Nash equilibrium is a trigger profile, and from Lemma 4.4, we get a winning strategy σ_{\exists} in $\mathcal{H}(\mathcal{G}, \pi)$. The outcome ρ of σ_{\exists} from s when *Adam* obeys *Eve* is such that $\pi = \text{proj}_1(\rho)$ is the outcome of the Nash equilibrium. Now for all prefix $\rho_{\leq i}$, the strategy $\sigma_{\exists}^i: h \mapsto \sigma_{\exists}(\rho_{\leq i} \cdot h)$ is such that any play in $\rho_{\leq i} \cdot \text{Out}(\rho_{=i}, \sigma_{\exists}^i)$ is winning for *Eve*.

Conversely, let ρ' be a path in $\mathcal{H}(\mathcal{G}, \pi)$ and assume it satisfies all three conditions. We define a strategy λ_{\exists} that follows ρ' when *Adam* obeys. Along ρ' , this strategy is defined as follows: $\lambda_{\exists}(\rho'_{\leq 2i}) = m_{\text{Agt}}$ such that $\text{Tab}(\text{proj}_1(\rho'_{=i}), m_{\text{Agt}}) = \text{proj}_1(\rho'_{=i+1})$. Such a legal

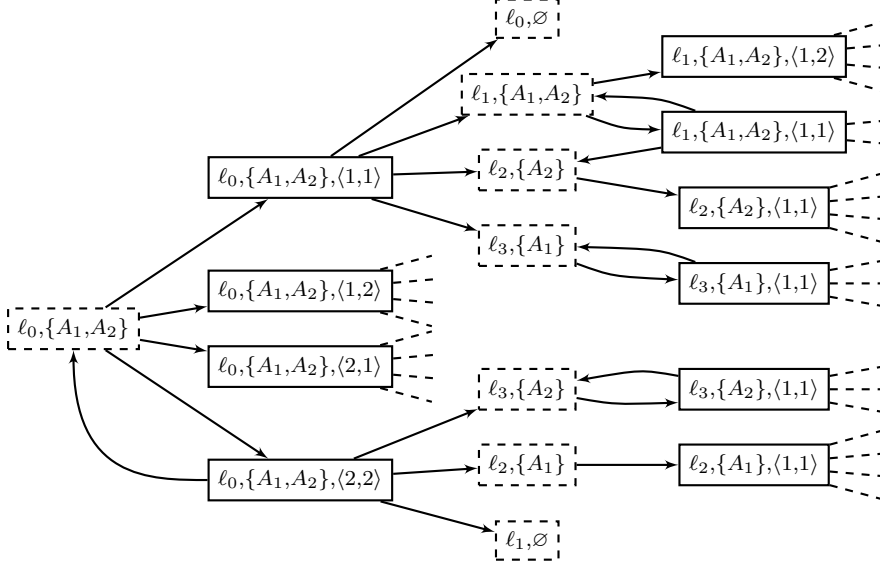


FIGURE 10. A small part of the suspect game for the game of Figure 3

move must exist since **Adam** obeys **Eve** along ρ' by condition 2. Now, if **Adam** deviates from the obeying strategy (at step i), we make λ_{\exists} follow the strategy σ_{\exists}^i (given by condition 3), which will ensure that the outcome is winning for **Eve**.

The outcomes of λ_{\exists} are then either the path ρ' , or a path ρ'' obtained by following a winning strategy after a prefix of ρ' . The path ρ'' is losing for **Adam**, hence for all $A \in \lambda(\rho')$, $\rho'' \lesssim_A \rho'$. This proves that λ_{\exists} is a winning strategy. Applying Lemma 4.4, we obtain a strategy profile σ_{Agt} in \mathcal{G} that is a trigger profile for π . Moreover, the outcome of σ_{Agt} from s is $\text{proj}_1(\rho')$ (using condition 1), so that σ_{Agt} is a Nash equilibrium. \square

Remark 4.6. Assume the preference relations of each player A in \mathcal{G} are prefix-independent, i.e., for all plays ρ and ρ' , $\rho \lesssim_A \rho'$ iff for all indices i and j , $\rho_{\geq i} \lesssim_A \rho'_{\geq j}$. Then the winning condition of **Eve** is also prefix-independent, and condition 3 just states that ρ' has to stay within the winning region of **Eve**. Note that, for prefix-dependent preference relations, condition 3 does not reduce to stay within the winning region of **Eve**: for instance, for safety objectives, if the losing states of all the players have been visited then any prolongation will satisfy the condition, even though it might leave the winning region of **Eve**.

Example 4.7. We depict on Figure 10 part of the suspect game for the game of Figure 3. Note that the structure of $\mathcal{H}(\mathcal{G}, \pi)$ does not depend on π . Only the winning condition is affected by the choice of π .

In the rest of the paper, we use the suspect-game construction to algorithmically solve the NE existence problem and the constrained NE existence problem in finite games for large classes of preference relations. Before that we carefully analyse the size of the suspect game when the original game is finite.

4.3. Size of the suspect games when the original game is finite. We suppose that \mathcal{G} is finite. At first sight, the number of states in $\mathcal{H}(\mathcal{G}, \pi)$ is exponential (in the number of

players of \mathcal{G}). However, there are two cases for which we easily see that the number of states of $\mathcal{H}(\mathcal{G}, \pi)$ is actually only polynomial:

- if there is a state in which all the players have several possible moves, then the transition table (which is part of the input, as discussed in Remark 2.2) is also exponential in the number of players;
- if the game is turn-based, then the transition table is “small”, but there is always at most one suspect player (unless all of them are suspects), so that the number of reachable states in $\mathcal{H}(\mathcal{G}, \pi)$ is also small.

We now prove that, due to the explicit encoding of the set of transitions (recall Remark 2.2, page 7), this can be generalised:

Proposition 4.8. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ be a finite concurrent game and π be a play in \mathcal{G} . The number of reachable configurations from $\text{States} \times \{\text{Agt}\}$ in $\mathcal{H}(\mathcal{G}, \pi)$ is polynomial in the size of \mathcal{G} .*

Proof. The game $\mathcal{H}(\mathcal{G}, \pi)$ contains the state (s, Agt) and the states $(s, \text{Agt}, m_{\text{Agt}})$, where m_{Agt} is a legal move from s ; the number of these states is bounded by $|\text{States}| + |\text{Tab}|$. The successors of those states that are not of the same form, are the $(t, \text{Susp}((s, t), m_{\text{Agt}}))$ with $t \neq \text{Tab}(s, m_{\text{Agt}})$. If some player $A \in \text{Agt}$ is a suspect for transition (s, t) , then besides m_A , she must have at least a second action m' , for which $\text{Tab}(s, m_{\text{Agt}}[A \mapsto m']) = t$. Thus the transition table from state s has size at least $2^{|\text{Susp}((s, t), m_{\text{Agt}})|}$. The successors of $(t, \text{Susp}((s, t), m_{\text{Agt}}))$ are of the form (t', P) or (t', P, m_{Agt}) where P is a subset of $\text{Susp}((s, t), m_{\text{Agt}})$; there can be no more than $(|\text{States}| + |\text{Tab}|) \cdot 2^{|\text{Susp}((s, t), m_{\text{Agt}})|}$ of them, which is bounded by $(|\text{States}| + |\text{Tab}|) \cdot |\text{Tab}|$. The total number of reachable states is then bounded by $(|\text{States}| + |\text{Tab}|) \cdot (1 + (|\text{States}| + |\text{Tab}|) \cdot |\text{Tab}|)$. \square

5. SINGLE-OBJECTIVE PREFERENCE RELATIONS

In this section we will be interested in finite games with single-objective preference relations.

The value problem for finite concurrent games with ω -regular objectives has standard solutions in game theory; they are given in Table 1 (page 3). Let us briefly give some explanations. Most of the basic literature on two-player games focus on turn-based games, and in particular algorithms for solving two-player games with ω -regular objectives only deal with turn-based games (see for instance [23, Chapter 2]). In particular, McNaughton developed an algorithm to solve turn-based parity games in time $O(|\text{States}| \cdot |\text{Edg}|^{p-1})$, where $p - 1$ is the number of priorities [32]. Büchi games and co-Büchi games correspond to parity games with two priorities, hence they are solvable in polynomial time. Similarly reachability games and safety games can be transformed into Büchi games by making the target states absorbing. Hence turn-based game with these types of objectives can be solved in polynomial time.

Note however that we can reuse these algorithms in the concurrent case as follows. Any finite concurrent zero-sum game with objective Ω for player A_1 can be transformed into a turn-based zero-sum game with objective $\tilde{\Omega}$ for player A_1 : the idea is to replace any edge labelled with pair of actions $\langle a_1, a_2 \rangle$ into two consecutive transitions labelled with a_1 (belonging to player A_1) and with a_2 (belonging to player A_2). Furthermore Ω is an ω -regular condition, then so is $\tilde{\Omega}$, and the type of the objective (reachability, Büchi, etc) is preserved (note however that this transformation only preserves Player A_1 objective). Hence

the standard algorithm on the resulting turn-based game can be applied. Lower bounds for reachability/safety and Büchi/co-Büchi games are also folklore results, and can be obtained by encoding the circuit-value problem (we recall the encoding in Section 5.3.3).

We now focus on the NE existence problem and on the constrained NE existence problem when each player has a single (ω -regular) objective using the suspect game construction. The results are summarised in the second column of Table 1.

Streett and Muller objectives are not explicitly mentioned in the rest of the section. The complexity of their respective (constrained) NE existence problems, which is given in Table 1, can easily be inferred from other ones. The $\text{P}_{\parallel}^{\text{NP}}$ -hardness for the NE existence problem with Streett objectives follows from the corresponding hardness for parity objectives (parity objectives can be encoded efficiently as Streett objectives). Hardness for the NE existence problem in Muller games, is deduced from hardness of the value problem (which holds for turn-based games), applying Proposition 3.4. For both objectives, membership in PSPACE follows from PSPACE membership for objectives given as Boolean circuits, since they can efficiently be encoded as Boolean circuits.

We fix for the rest of the section a multi-player finite game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$, and we assume that each \preceq_A is single-objective, given by set Ω_A .

Remark 5.1. Let us come back to Remark 2.2 on our choice of an explicit encoding for the set of transitions. Assuming more compact encodings, the complexity of computing Nash equilibria for qualitative objectives does not allow to distinguish between the intrinsic complexity of the objectives. Indeed, in the formalism of [30], the transition function is given in each state by a finite sequence $((\phi_0, s_0), \dots, (\phi_h, s_h))$, where $s_i \in \text{States}$, and ϕ_i is a boolean combination of propositions ($A = m$) that evaluates to true iff agent A chooses action m . The transition table is then defined as follows: $\text{Tab}(s, m_{\text{Agt}}) = s_j$ iff j is the smallest index such that ϕ_j evaluates to true when, for every player $A \in \text{Agt}$, A chooses action m_A . It is required that the last boolean formula ϕ_h be \top , so that no agent can enforce a deadlock.

We can actually state the following result, whose proof is postponed to the Appendix on page 69.

Proposition 5.2. *For finite concurrent games with compact encoding of transition functions and with reachability/Büchi/safety objectives, the constrained NE existence problems is PSPACE-hard.*

Remark 5.3. It is first interesting to notice that given two plays π and π' the suspect games $\mathcal{H}(\mathcal{G}, \pi)$ and $\mathcal{H}(\mathcal{G}, \pi')$ only differ in their winning conditions. In particular, the structure of the game only depends on \mathcal{G} , and has polynomial size (see Proposition 4.8). We denote it with $\mathcal{J}(\mathcal{G})$. Moreover, as each relation \preceq_A is given by a single objective Ω_A , the winning condition for Eve in $\mathcal{H}(\mathcal{G}, \pi)$ rewrites as: for every $A \in \lambda(\rho) \cap \text{Los}(\pi)$, $\text{proj}_1(\rho)$ is losing (in \mathcal{G}) for player A , where $\text{Los}(\pi)$ is the set of players losing along π in \mathcal{G} . This winning condition only depends on $\text{Los}(\pi)$ (not on the precise value of play π). Therefore in this section, the suspect game is denoted with $\mathcal{H}(\mathcal{G}, L)$, where $L \subseteq \text{Agt}$, and Eve wins play ρ if, for every $A \in \lambda(\rho) \cap L$, A loses along $\text{proj}_1(\rho)$ in \mathcal{G} . In many cases we will be able to simplify this winning condition, and to obtain simple algorithms to the corresponding problems.

We now distinguish between the winning objectives of the players. There are some similarities in some of the cases (for instance safety and co-Büchi objectives), but they nevertheless all require specific techniques and proofs.

5.1. Reachability objectives. The value problem for a reachability winning condition is P-complete. Below, we design a non-deterministic algorithm that runs in polynomial time for solving the constrained NE existence problem. We then end this subsection with a NP-hardness proof of the constrained NE existence problem and NE existence problem. In the end, we prove the following result:

Theorem 5.4. *For finite concurrent games with single reachability objectives, the NE existence problem and the constrained NE existence problem are NP-complete.*

5.1.1. Reduction to a safety game. We assume that for every player A , Ω_A is a single reachability objective given by target set T_A . Given $L \subseteq \text{Agt}$, in the suspect game $\mathcal{H}(\mathcal{G}, L)$, we show that the objective of **Eve** reduces to a safety objective. We define the safety objective Ω_L in $\mathcal{H}(\mathcal{G}, L)$ by the set $T_L = \{(s, P) \mid \exists A \in P \cap L. s \in T_A\}$ of target states.

Lemma 5.5. *Eve has a winning strategy in game $\mathcal{H}(\mathcal{G}, L)$ iff Eve has a winning strategy in game $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L .*

Proof. We first show that any play in Ω_L is winning in $\mathcal{H}(\mathcal{G}, L)$. Let $\rho \in \Omega_L$, and let $A \in \lambda(\rho) \cap L$. Toward a contradiction assume that $\text{Occ}(\text{proj}_1(\rho)) \cap T_A \neq \emptyset$: there is a state (s, P) along ρ with $s \in T_A$. Obviously $\lambda(\rho) \subseteq P$, which implies that $A \in P \cap L$. This contradicts the fact that $\rho \notin \Omega_L$. We have shown so far that any winning strategy for **Eve** in $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L is a winning strategy for **Eve** in $\mathcal{H}(\mathcal{G}, L)$.

Now assume that **Eve** has no winning strategy in game $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L . Turn-based games with safety objectives being determined, **Adam** has a strategy σ_{\forall} which ensures that no outcome of σ_{\forall} is in Ω_L . If $\rho \in \text{Out}(\sigma_{\forall})$, there is a state (s, P) along ρ such that there is $A \in P \cap L$ with $s \in T_A$. We now modify the strategy of **Adam** such that as soon as such a state is reached we switch from σ_{\forall} to the strategy that always obeys **Eve**. This ensures that in every outcome ρ' of the new strategy, we reach a state (s, P) such that there is $A \in P \cap L$ with $s \in T_A$, and $\lambda(\rho') = P$. This **Adam's** strategy thus makes **Eve** lose the game $\mathcal{H}(\mathcal{G}, L)$, and **Eve** has no winning strategy in game $\mathcal{H}(\mathcal{G}, L)$. \square

5.1.2. Algorithm. The algorithm for solving the constrained NE existence problem in a game where each player has a single reachability objective relies on Theorem 4.5 and Proposition 3.1, and on the above analysis:

- (i) guess a lasso-shaped play $\rho = \tau_1 \cdot \tau_2^\omega$ (with $|\tau_i| \leq 2|\text{States}|^2$) in $\mathcal{J}(\mathcal{G})$, such that **Adam** obeys **Eve** along ρ , and $\pi = \text{proj}_1(\rho)$ satisfies the constraint on the payoff;
- (ii) compute the set $W(\mathcal{G}, \text{Los}(\pi))$ of states that are winning for **Eve** in the suspect game $\mathcal{H}(\mathcal{G}, \text{Los}(\pi))$, where $\text{Los}(\pi)$ is the set of losing players along π ;
- (iii) check that ρ stays in $W(\mathcal{G}, \text{Los}(\pi))$.

First notice that this algorithm is non-deterministic and runs in polynomial time: the witness ρ guessed in step (i) has size polynomial; the suspect game $\mathcal{H}(\mathcal{G}, \text{Los}(\pi))$ has also polynomial size (Proposition 4.8); Step (ii) can be done in polynomial time using a standard attractor computation [23, Sect. 2.5.1] as the game under analysis is equivalent to a safety game (Lemma 5.5); finally step (iii) can obviously be performed in polynomial time.

Step (i) ensures that conditions 2 and 1 of Theorem 4.5 hold for ρ and step (iii) ensures condition 3. Correctness of the algorithm then follows from Theorem 4.5 and Proposition 3.1.

5.1.3. *Hardness.* We prove NP-hardness of the constrained NE existence problem by encoding an instance of 3SAT as follows. We assume set of atomic propositions $\text{AP} = \{x_1, \dots, x_k\}$, and we let $\phi = \bigwedge_{i=1}^n c_i$ where $c_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$ where $\ell_{i,j} \in \{x_k, \neg x_k \mid 1 \leq k \leq p\}$. We build the turn-based game \mathcal{G}_ϕ with $n + 1$ players $\text{Agt} = \{A, C_1, \dots, C_n\}$ as follows: for every $1 \leq k \leq p$, player A chooses to visit either location x_k or location $\neg x_k$. Location x_k is winning for player C_i if, and only if, x_k is one of the literals in c_i , and similarly location $\neg x_k$ is winning for C_i if, and only if, $\neg x_k$ is one of the literals of c_i . The construction is illustrated on Figure 11, with the reachability objectives defined as $\Omega_{C_i} = \{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$ for $1 \leq i \leq n$. Now, it is easy to check that this game has a Nash equilibrium with payoff 1 for all players $(C_i)_{1 \leq i \leq n}$ if, and only if, ϕ is satisfiable.

We prove hardness for the NE existence problem by using the transformation described in Section 3.4 once for each player. We define the game \mathcal{G}_0 similar to \mathcal{G} but with an extra player C_{n+1} who does not control any state for now. For $1 \leq i \leq n$, we define $\mathcal{G}_i = E(\mathcal{G}_{i-1}, C_i, C_{n+1}, \rho)$, where ρ is a winning path for C_i . The preference relation can be expressed in any \mathcal{G}_i by a reachability condition, by giving to C_{n+1} a target which is the initial state of \mathcal{G} . According to Proposition 3.5 there is a Nash equilibrium in \mathcal{G}_i if, and only if, there is one in \mathcal{G}_{i-1} where C_i wins. Therefore there is a Nash equilibrium in \mathcal{G}_n if, and only if, ϕ is satisfiable. This entails NP-hardness of the NE existence problem.




FIGURE 11. Reachability game for the reduction of 3SAT

5.2. **Safety objectives.** The value problem for safety objectives is P-complete. We next show that the constrained NE existence problem can be solved in NP, and conclude with NP-hardness of both the constrained NE existence problem and the NE existence problem. We hence prove:

Theorem 5.6. *For finite games with single safety objectives, the NE existence problem and the constrained NE existence problem are NP-complete.*

5.2.1. *Reduction to a conjunction of reachability objectives.* We assume Ω_A is a single safety objective given by set T_A . In the corresponding suspect game, we show that the goal of Eve is equivalent to a conjunction of reachability objectives. Let $L \subseteq \text{Agt}$. In suspect game $\mathcal{H}(\mathcal{G}, L)$, we define several reachability objectives as follows: for each $A \in L$, we define $T'_A = T_A \times \{P \mid P \subseteq \text{Agt}\} \cup \text{States} \times \{P \mid A \notin P\}$, and we write Ω'_A for the corresponding reachability objectives.

Lemma 5.7. *A play ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$ iff $\rho \in \bigcap_{A \in L} \Omega'_A$.*

Proof. Let ρ be a play in $\mathcal{H}(\mathcal{G}, L)$, and assume it is winning for Eve. Then, for each $A \in \lambda(\rho) \cap L$, $\rho \notin \Omega_A$, which means that the target set T_A is visited along $\text{proj}_1(\rho)$, and

therefore T'_A is visited along ρ . If $A \notin \lambda(\rho)$, then a state (s, P) with $A \notin P$ is visited by ρ : the target set T'_A is visited. This implies that $\rho \in \bigcap_{A \in L} \Omega'_A$.

Conversely let $\rho \in \bigcap_{A \in L} \Omega'_A$. For every $A \in L$, T'_A is visited by ρ . Then, either T_A is visited by $\text{proj}_1(\rho)$ (which means that $\rho \notin \Omega_A$) or $A \notin \lambda(\rho)$. In particular, ρ is a winning play for Eve in $\mathcal{H}(\mathcal{G}, L)$. \square

5.2.2. Algorithm for solving finite zero-sum turn-based games with a conjunction of reachability objectives. We now give a simple algorithm for solving zero-sum games with a conjunction of reachability objectives. This algorithm works in exponential time with respect to the size of the conjunction (we will see in Subsection 7.1.6 that the problem is PSPACE-complete). However for computing Nash equilibria in safety games we will only use it for small (logarithmic size) conjunctions.

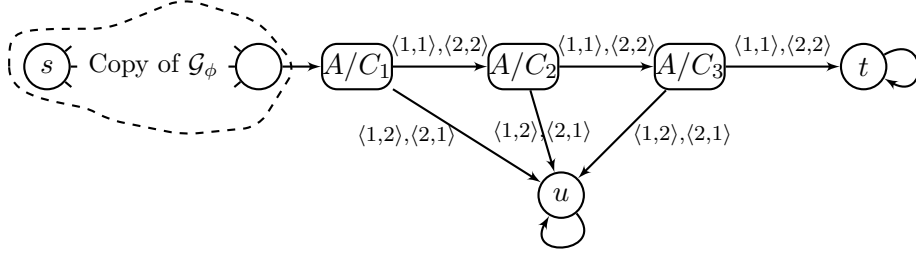
Let $\overline{\mathcal{G}}$ be a two-player turn-based game with a winning objective for Eve given as a conjunction of k reachability objectives $\Omega_1, \dots, \Omega_k$. We assume vertices of Eve and Adam in $\overline{\mathcal{G}}$ are V_{\exists} and V_{\forall} respectively, and that the initial vertex is v_0 . The idea is to construct a new game $\overline{\mathcal{G}}'$ that remembers the objectives that have been visited so far. The vertices of game $\overline{\mathcal{G}}'$ controlled by Eve and Adam are $V'_{\exists} = V_{\exists} \times 2^{\llbracket 1, k \rrbracket}$ and $V'_{\forall} = V_{\forall} \times 2^{\llbracket 1, k \rrbracket}$ respectively. There is a transition from (v, S) to (v', S') iff there is a transition from v to v' in the original game and $S' = S \cup \{i \mid v' \in \Omega_i\}$. The reachability objective Ω for Eve is given by target set $\text{States} \times \llbracket 1, k \rrbracket$. It is clear that there is a winning strategy in $\overline{\mathcal{G}}$ from v_0 for the conjunction of reachability objectives $\Omega_1, \dots, \Omega_k$ iff there is a winning strategy in game $\overline{\mathcal{G}}'$ from $(v_0, \{i \mid v_0 \in \Omega_i\})$ for the reachability objective Ω . The number of vertices of this new game is $|V'_{\exists} \cup V'_{\forall}| = |V_{\exists} \cup V_{\forall}| \cdot 2^k$, and the size of the new transition table Tab' is bounded by $|\text{Tab}| \cdot 2^k$, where Tab is the transition table of $\overline{\mathcal{G}}$. An attractor computation on $\overline{\mathcal{G}}'$ is then done in time $\mathcal{O}(|V'_{\exists} \cup V'_{\forall}| \cdot |\text{Tab}'|)$, we obtain an algorithm for solving zero-sum games with a conjunction of reachability objectives, running in time $\mathcal{O}(2^{2k} \cdot (|V_{\exists} \cup V_{\forall}| \cdot |\text{Tab}|))$.

5.2.3. Algorithm. The algorithm for solving the constrained NE existence problem for single reachability objectives could be copied and would then be correct. It would however not yield an NP upper bound. We therefore propose a refined algorithm:

- (i) guess a lasso-shaped play $\rho = \tau_1 \cdot \tau_2^{\omega}$ (with $|\tau_i| \leq |\text{States}|^2$) in $\mathcal{J}(\mathcal{G})$ such that Adam obeys Eve along ρ , and $\pi = \text{proj}_1(\rho)$ satisfies the constraint on the payoff. Note that if $\text{Los}(\pi)$ is the set of players losing in π , computing $W(\mathcal{G}, \text{Los}(\pi))$ would require exponential time. We will avoid this expensive computation.
- (ii) check that any Adam-deviation along ρ , say at position i (for any i), leads to a state from which Eve has a strategy σ_{\exists}^i to ensure that any play in $\rho_{\leq i} \cdot \text{Out}(\sigma_{\exists}^i)$ is winning for her.

Step (ii) can be done as follows: pick an Adam-state $(s, \text{Agt}, m_{\text{Agt}})$ along ρ and a successor (t, P) such that $t \neq \text{Tab}(s, m_{\text{Agt}})$; we only need to show that $(t, P) \in W(\mathcal{G}, (\text{Los}(\pi) \setminus \text{Los}(\rho_{\leq i})) \cap P)$. We can compute this set efficiently (in polynomial time) using the algorithm of the previous paragraph since $2^{|P|} \leq |\text{Tab}|$ (using the same argument as in Proposition 4.8).

This non-deterministic algorithm, which runs in polynomial time, precisely implements Theorem 4.5, and therefore correctly decides the constrained NE existence problem.


 FIGURE 12. Extending game \mathcal{G}_ϕ with final concurrent modules

5.2.4. *Hardness.* The NP-hardness for the constrained NE existence problem can be proven by encoding an instance of 3SAT using a game similar to that for reachability objectives, see Section 5.1. We only change the constraint which is now that all players C_i should be losing, and we get the same equivalence.

The reduction of Lemma 3.4 cannot be used to deduce the hardness of the NE existence problem, since it assumes a lower bound on the payoff. Here the constraint is an upper bound (“each player should be losing”). We therefore provide an ad-hoc reduction in this special case, which is illustrated on Figure 12. We add some module at the end of the game to enforce that in an equilibrium, all players are losing. We add concurrent states between A and each C_i (named A/C_i). All players C_i are trying to avoid t , and A is trying to avoid u .

Since A has no target in \mathcal{G}_ϕ she cannot lose before seeing u , and then she can always change her strategy in the concurrent states in order to go to t . Therefore an equilibrium always ends in t . A player C_i whose target was not seen during game \mathcal{G}_ϕ , can change her strategy in order to go to u instead of t . That means that if there is an equilibrium, there was one in \mathcal{G}_ϕ where all C_i are losing. Conversely, if there was such an equilibrium in \mathcal{G}_ϕ , we can extend this strategy profile by one whose outcome goes to t and it is an equilibrium in the new game. This concludes the NP-hardness of the NE existence problem.

5.3. **Büchi objectives.** The value problem for Büchi objectives is P-complete. In this subsection we design a polynomial-time algorithm for solving the constrained NE existence problem for Büchi objectives. The P-hardness of the NE existence problem can then be inferred from the P-hardness of the value problem, applying Propositions 3.2 and 3.4. Globally we prove the following result:

Theorem 5.8. *For finite games with single Büchi objectives, the NE existence problem and the constrained NE existence problem are P-complete.*

5.3.1. *Reduction to a co-Büchi game.* We assume that for every player A , Ω_A is a Büchi objective given by target set T_A . Given $L \subseteq \text{Agt}$, in the suspect game $\mathcal{H}(\mathcal{G}, L)$, we show that the objective of Eve is equivalent to a single co-Büchi objective. We define the co-Büchi objective Ω_L in $\mathcal{H}(\mathcal{G}, L)$ given by the target set $T_L = \{(s, P) \mid \exists A \in P \cap L. s \in T_A\}$. Notice that the target set is defined in the same way as for reachability objectives.

Lemma 5.9. *A play ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$ iff $\rho \in \Omega_L$.*

Proof. Assume that ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$. Then for every $A \in \lambda(\rho) \cap L$, it holds $\text{Inf}(\text{proj}_1(\rho)) \cap T_A = \emptyset$. Toward a contradiction, assume that $\text{Inf}(\rho) \cap T_L \neq \emptyset$. There exists

(s, P) such that there is $A \in P \cap L$ with $s \in T_A$, which appears infinitely often along ρ . In particular, $P = \lambda(\rho)$ (otherwise it would not appear infinitely often along ρ). Hence, we have found $A \in \lambda(\rho) \cap L$ such that $\text{Inf}(\text{proj}_1(\rho)) \cap T_A \neq \emptyset$, which is a contradiction. Therefore, $\rho \in \Omega_L$.

Assume $\rho \in \Omega_L$: for every (s, P) such that there exists $A \in P \cap L$ with $s \in T_A$, (s, P) appears finitely often along ρ . Let $A \in \lambda(\rho) \cap L$, and assume towards a contradiction that there is $s \in T_A$ such that s appears infinitely often along $\text{proj}_1(\rho)$. This means that $(s, \lambda(\rho))$ appears infinitely often along ρ , which contradicts the above condition. Therefore, ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$. \square

5.3.2. Algorithm. As for reachability objectives, the winning region for Eve in $\mathcal{H}(\mathcal{G}, L)$ can be computed in polynomial time (since this is the winning region of a co-Büchi game, see Lemma 5.9 above). A non-deterministic algorithm running in polynomial time similar to the one for reachability objectives can therefore be inferred. However we can do better than guessing an appropriate lasso-shaped play ρ by looking at the strongly connected components of the game: a strongly connected component of the game uniquely defines a payoff, which is that of all plays that visit infinitely often all the states of that strongly connected component. Using a clever partitioning of the set of strongly connected components of the game, we obtain a polynomial-time algorithm.

From now on and until the end of Subsection 5.3.2 we relax the hypotheses on the preference relations (that they are all single-objective with a Büchi condition). We present an algorithm in a more general context, since the same techniques will be used in Subsection 6.2.2 (and we chose to only present once the construction). For the rest of this subsection we therefore make the following assumptions on the preference relations $(\lesssim_A)_{A \in \text{Agt}}$. For every player $A \in \text{Agt}$:

- (a) \lesssim_A only depends on the set of states which is visited infinitely often: if ρ and ρ' are two plays such that $\text{Inf}(\rho) = \text{Inf}(\rho')$ then $\rho \lesssim_A \rho'$ and $\rho' \lesssim_A \rho$;
- (b) \lesssim_A is given by an ordered objective ω_A with preorder \lesssim_A , and \lesssim_A is supposed to be monotonic;
- (c) for every threshold w^A , we can compute in polynomial time $S^A \subseteq \text{States}$ such that $\text{Inf}(\rho) \subseteq S^A \Leftrightarrow \rho \lesssim_A w^A$.

Obviously preferences given by single Büchi objectives do satisfy those hypotheses. At every place where it is relevant, we will explain how the particular case of single Büchi objectives is handled. Next we write (\star) for the above assumptions, and $(\star)_a$ (resp. $(\star)_b$, $(\star)_c$) for only the first (resp. second, third) assumption.

We first characterise the ‘good’ plays in $\mathcal{J}(\mathcal{G})$ in terms of the strongly connected components they define: the strongly connected component defined by a play is the set of states that are visited infinitely often by the play. We fix for each player A , equivalence classes of plays u^A and w^A , that represent lower- and upper-bounds for the constrained NE existence problem. Both can be represented as finite sets, representing the set of states which are visited infinitely often. For each $K \subseteq \text{States}$, we write $v^A(K)$ for the equivalence class of all paths π that visits infinitely often exactly K , i.e.: $\text{Inf}(\pi) = K$. We also write $v(K) = (v^A(K))_{A \in \text{Agt}}$. We look for a transition system $\langle K, E \rangle$, with $K \subseteq \text{States}$ and $E \subseteq K \times K$, for which the following properties hold:

- (1) $u^A \lesssim_A v^A(K) \lesssim_A w^A$ for all $A \in \text{Agt}$;

- (2) $\langle K, E \rangle$ is strongly connected;
- (3) $\forall k \in K. (k, \text{Agt}) \in W(\mathcal{G}, v(K))$;
- (4) $\forall (k, k') \in E. \exists (k, \text{Agt}, m_{\text{Agt}}) \in W(\mathcal{G}, v(K)). \text{Tab}(k, m_{\text{Agt}}) = k'$;
- (5) $(K \times \{\text{Agt}\})$ is reachable from (s, Agt) in $W(\mathcal{G}, v(K))$;

where $W(\mathcal{G}, v(K))$ is the winning region of Eve in suspect game $\mathcal{H}(\mathcal{G}, v(K))$.⁵

If one can find one such transition system $\langle K, E \rangle$, then we will be able to build a lasso-play ρ from (s, Agt) in the suspect-game that will satisfy the conditions of Theorem 4.5. Formally, we have the following lemma:

Lemma 5.10. *Under hypothesis $(\star)_a$, there is a transition system $\langle K, E \rangle$ satisfying conditions (1)–(5) if, and only if, there is a path ρ from (s, Agt) in $\mathcal{H}(\mathcal{G}, v(K))$ that never gets out of $W(\mathcal{G}, v(K))$, along which Adam always obeys Eve, $u^A \lesssim_A v^A(K) \lesssim_A w^A$ for all $A \in \text{Agt}$, and $\text{proj}_1(\text{Inf}(\rho) \cap V_{\exists}) = K$ (which implies that $\rho \in v^A(K)$ for all A).*

Proof. The first implication is shown by building a path in $W(\mathcal{G}, v(K))$ that successively visits all the states in $K \times \{\text{Agt}\}$ forever. Thanks to (5), (2) and (4) (and the fact that Adam obeys Eve), such a path exists, and from (3) and (4), this path remains in the winning region. From (1), we have the condition on the preferences. Conversely, consider such a path ρ , and let $K = \text{proj}_1(\text{Inf}(\rho) \cap V_{\exists})$ and $E = \{(k, k') \in K^2 \mid \exists (k, \text{Agt}, m_{\text{Agt}}) \in \text{Inf}(\rho). \text{Tab}(k, m_{\text{Agt}}) = k'\}$. Condition (5) clearly holds. Conditions (1), (3) and (4) are easy consequences of the hypotheses and construction. We prove that $\langle K, E \rangle$ is strongly connected. First, since Adam obeys Eve and ρ starts in (k, Agt) , we have $\lambda(\rho) = \text{Agt}$. Now, take any two states k and k' in K : then ρ visits (k, Agt) and (k', Agt) infinitely often, and there is a subpath of ρ between those two states, all of which states appear infinitely often along ρ . Such a subpath gives rise to a path between k and k' , as required. \square

As a consequence, if $\langle K, E \rangle$ satisfies the five previous conditions, by Theorem 4.5, there is a Nash equilibrium whose outcome lies between the bounds u^A and w^A . Our aim is to compute efficiently all maximal pairs $\langle K, E \rangle$ that satisfy the five conditions.

To that aim we define a recursive function SSG (standing for “solve sub-game”), working on transition systems, that will decompose efficiently any transition system that does not satisfy the five conditions above into polynomially many disjoint sub-transition systems *via* a decomposition into strongly connected components.

- if $K \times \{\text{Agt}\} \subseteq W(\mathcal{G}, v(K))$, and if for all $(k, k') \in E$ there is a $(k, \text{Agt}, m_{\text{Agt}})$ in $W(\mathcal{G}, v(K))$ s.t. $\text{Tab}(k, m_{\text{Agt}}) = k'$, and finally if $\langle K, E \rangle$ is strongly connected, then we set $\text{SSG}(\langle K, E \rangle) = \{\langle K, E \rangle\}$. This means that conditions (2)–(4) are satisfied by $\langle K, E \rangle$.
- otherwise, we let

$$\text{SSG}(\langle K, E \rangle) = \bigcup_{\langle K', E' \rangle \in \text{SCC}(\langle K, E \rangle)} \text{SSG}(T(\langle K', E' \rangle))$$

where $\text{SCC}(\langle K, E \rangle)$ is the set of strongly connected components of $\langle K, E \rangle$ (which can be computed in linear time), and where $T(\langle K', E' \rangle)$ is the transition system whose set of states is $\{k \in K' \mid (k, \text{Agt}) \in W(\mathcal{G}, v(K'))\}$ and whose set of edges is

$$\{(k, k') \in E' \mid \exists (k, \text{Agt}, m_{\text{Agt}}) \in W(\mathcal{G}, v(K')). \text{Tab}(k, m_{\text{Agt}}) = k'\}.$$

⁵Formally the suspect game has been defined with a play as reference, and not a equivalence class. However, in this subsection, if π and π' are equivalent, the games $\mathcal{H}(\mathcal{G}, \pi)$ and $\mathcal{H}(\mathcal{G}, \pi')$ are identical.

Notice that this set of edges is never empty, but $T(\langle K', E' \rangle)$ might not be strongly connected anymore, so that this is really a recursive definition.

The recursive function SSG decomposes any (sub-)transition system of the game into a list of disjoint transition systems which all satisfy conditions (2)-(4) above.

So far the computation does not take into account the bounds for the payoffs of the players (lower bound u^A and upper bound w^A for player A). For each upper bound w^A , we assume condition $(\star)_c$ holds. In the particular case of a single Büchi objective for each player define by target T_A , this is simply done by setting $S^A = \text{States} \setminus T_A$, if this player has to be losing (that is, if w^A does not satisfy the Büchi objective). Now assuming we have found the appropriate set S^A , we define

$$\text{Sol} = \text{SSG}\left(\bigcap_{A \in \text{Agt}} S^A, \text{Edg}'\right) \cap \{\langle K, E \rangle \mid \forall A \in \text{Agt}. u^A \lesssim v^A(K)\}$$

where Edg' restricts Edg to $\bigcap_{A \in \text{Agt}} S^A$.

We now show that the set Sol computes (in a sense that we make clear) the transition systems that are mentioned in Lemma 5.10).

Lemma 5.11. *We suppose condition (\star) holds. If $\langle K, E \rangle \in \text{Sol}$ then it satisfies conditions (1) to (4). Conversely, if $\langle K, E \rangle$ satisfies conditions (1) to (4), then there exists $\langle K', E' \rangle \in \text{Sol}$ such that $\langle K, E \rangle \subseteq \langle K', E' \rangle$.*

Proof. Let $\langle K, E \rangle \in \text{Sol}$. By definition of SSG , all (k, Agt) for $k \in K$ are in $W(\mathcal{G}, v(K))$, and for all $(k, k') \in E$, there is a state $(k, \text{Agt}, m_{\text{Agt}})$ in $W(\mathcal{G}, v(K))$ such that $\text{Tab}(k, m_{\text{Agt}}) = k'$, and $\langle K, E \rangle$ is strongly connected. Also, for all A , $u^A \lesssim v^A(K)$ because $\text{Sol} \subseteq \{\langle K, E \rangle \mid u^A \lesssim v^A(K)\}$. Finally, for any $A \in \text{Agt}$, $v^A(K) \lesssim w^A$ because the set K is included in S^A .

Conversely, assume that $\langle K, E \rangle$ satisfies the conditions. We show that if $\langle K, E \rangle \subseteq \langle K', E' \rangle$ then there is $\langle K'', E'' \rangle$ in $\text{SSG}(\langle K', E' \rangle)$ such that $\langle K, E \rangle \subseteq \langle K'', E'' \rangle$. The proof is by induction on the size of $\langle K', E' \rangle$.

The basic case is when $\langle K', E' \rangle$ satisfies the conditions (2), (3), and (4): in that case, $\text{SSG}(\langle K', E' \rangle) = \{\langle K', E' \rangle\}$, and by letting $\langle K'', E'' \rangle = \langle K', E' \rangle$ we get the expected result.

We now analyze the other case. There is a strongly connected component of $\langle K', E' \rangle$, say $\langle K'', E'' \rangle$, which contains $\langle K, E \rangle$, because $\langle K, E \rangle$ satisfies condition (2). We have $v^A(K) \lesssim_A v^A(K'')$ (because $K \subseteq K''$ and \lesssim_A is monotonic) for every A , and thus $W(\mathcal{G}, v(K)) \subseteq W(\mathcal{G}, v(K''))$. This ensures that $T(\langle K'', E'' \rangle)$ contains $\langle K, E \rangle$ as a subgraph. Since $\langle K'', E'' \rangle$ is a subgraph of $\langle K', E' \rangle$, the graph $T(\langle K'', E'' \rangle)$ also is. We show that they are not equal, so that we can apply the induction hypothesis to $T(\langle K'', E'' \rangle)$. For this, we exploit the fact that $\langle K', E' \rangle$ does not satisfy one of conditions (2) to (4):

- first, if $\langle K', E' \rangle$ is not strongly connected while $\langle K'', E'' \rangle$ is, they cannot be equal;
- if there is some $k \in K'$ such that (k, Agt) is not in $W(\mathcal{G}, v(K'))$, then k is not a vertex of $T(\langle K'', E'' \rangle)$;
- if there some edge (k, k') in E' such that there is no state $(k, \text{Agt}, m_{\text{Agt}})$ in $W(\mathcal{G}, v(K'))$ such that $\text{Tab}(k, m_{\text{Agt}}) = k'$, then the edge (k, k') is not in $T(\langle K'', E'' \rangle)$.

We then apply the induction hypothesis to $T(\langle K'', E'' \rangle)$, and get the expected result. Now, because of condition (1), $u^A \lesssim v^A(K) \lesssim w^A$. Hence, due to the previous analysis, there exists $\langle K', E' \rangle \in \text{SSG}\left(\langle \bigcap_{A \in \text{Agt}} S^A, \text{Edg}' \rangle\right)$ such that $\langle K, E \rangle \subseteq \langle K', E' \rangle$. This concludes the proof of the lemma. \square

Lemma 5.12. *Under assumptions (\star) , if for every K , the set $W(\mathcal{G}, v(K))$ can be computed in polynomial time, then the set Sol can also be computed in polynomial time.*

Proof. Each recursive call to SSG applies to a decomposition in strongly connected components of the current transition system under consideration. Hence the number of recursive calls is bounded by $|\text{States}|^2$. Computing the decomposition in SCCs can be done in linear time. By assumption, each set $W(\mathcal{G}, v(K))$ can be computed in polynomial time. S^A is obtained by removing the target of the losers (for w^A) from States. Hence globally we can compute Sol in polynomial time. \square

To conclude the algorithm, we need to check that condition (5) holds for one of the solutions $\langle K, E \rangle$ in Sol. It can be done in polynomial time by looking for a path in the winning region of Eve in $\mathcal{H}(\mathcal{G}, v(K))$ that reaches $K \times \{\text{Agt}\}$ from (s, Agt) . The correctness of the algorithm is ensured by the fact that if some $\langle K, E \rangle$ satisfies the five conditions, there is a $\langle K', E' \rangle$ in Sol with $K \subseteq K'$ and $E \subseteq E'$. Since $K \subseteq K'$ implies $v^A(K) \lesssim_A v^A(K')$, the winning region of Eve in $\mathcal{H}(\mathcal{G}, v(K'))$ is larger than that $\mathcal{H}(\mathcal{G}, v(K))$, which implies that the path from (s, Agt) to $K \times \{\text{Agt}\}$ is also a path from (s, Agt) to $K' \times \{\text{Agt}\}$. Hence, $\langle K', E' \rangle$ also satisfies condition (5), and therefore the five expected conditions.

We have already mentioned that single Büchi objectives do satisfy the hypotheses (\star) . Furthermore, Lemma 5.9 shows that, given $v(K)$, one can compute the set $W(\mathcal{G}, v(K))$ as the winning region of a co-Büchi turn-based game, which can be done in polynomial time (this is argued at the beginning of the section). Therefore Lemma 5.12 and the subsequent analysis apply: this concludes the proof that the constrained NE existence problem for finite games with single Büchi objectives is in P.

5.3.3. Hardness. We recall a possible proof of P-hardness for the value problem, from which we will infer the other lower bounds. The circuit-value problem can be easily encoded into a deterministic turn-based game with Büchi objectives: a circuit (which we assume w.l.o.g. has only AND- and OR-gates) is transformed into a two-player turn-based game, where one player controls the AND-gates and the other player controls the OR-gates. We add self-loops on the leaves. Positive leaves of the circuit are the (Büchi) objective of the OR-player, and negative leaves are the (Büchi) objective of the AND-player. Then obviously, the circuit evaluates to true iff the OR-player has a winning strategy for satisfying his Büchi condition, which in turn is equivalent to the fact that there is an equilibrium with payoff 0 for the AND-player, by Proposition 3.2. We obtain P-hardness for the NE existence problem, using Proposition 3.4: the preference relations in the game constructed in Proposition 3.4 are Büchi objectives.

5.4. Co-Büchi objectives. The value problem for co-Büchi objectives is P-complete. We now prove that the constrained NE existence problem is in NP, and that the constrained NE existence problem and the NE existence problem are NP-hard. We therefore deduce:

Theorem 5.13. *For finite games with single co-Büchi objectives, the NE existence problem and the constrained NE existence problem are NP-complete.*

The proof of this Theorem is very similar to that for safety objectives: instead of conjunction of reachability objectives, we need to deal with conjunction of Büchi objectives. Of course constructions and algorithms need to be adapted. That is what we present now.

5.4.1. *Reduction to a conjunction of Büchi conditions.* We assume that for every player A , Ω_A is a single co-Büchi objective Ω_A given by T_A . In the corresponding suspect game, we show that the goal of player **Eve** is equivalent to a conjunction of Büchi objectives. Let $L \subseteq \text{Agt}$. In suspect game $\mathcal{H}(\mathcal{G}, L)$, we define several Büchi objectives as follows: for each $A \in L$, we define $T'_A = T_A \times \{P \mid P \subseteq \text{Agt}\} \cup \text{States} \times \{P \mid A \notin P\}$, and we write Ω'_A for the corresponding Büchi objective.

Lemma 5.14. *A play ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$ iff $\rho \in \bigcap_{A \in L} \Omega'_A$.*

Proof. Let ρ be a play in $\mathcal{H}(\mathcal{G}, L)$, and assume it is winning for **Eve**. Then, for each $A \in \lambda(\rho) \cap L$, $\rho \notin \Omega_A$, which means that the target set T_A is visited along $\text{proj}_1(\rho)$, and therefore T'_A is visited infinitely often along ρ . If $A \notin \lambda(\rho)$, then a state (s, P) with $A \notin P$ is visited infinitely often by ρ : the target set T'_A is visited infinitely often. This implies that $\rho \in \bigcap_{A \in L} \Omega'_A$.

Conversely let $\rho \in \bigcap_{A \in L} \Omega'_A$. For every $A \in L$, T'_A is visited infinitely often by ρ . Then, either T_A is visited infinitely often by $\text{proj}_1(\rho)$ (which means that $\rho \notin \Omega_A$) or $A \notin \lambda(\rho)$. In particular, ρ is a winning play for **Eve** in $\mathcal{H}(\mathcal{G}, L)$. \square

5.4.2. *Algorithm for solving zero-sum games with a conjunction of Büchi objectives.* We adapt the algorithm for conjunctions of reachability objectives (page 28) to conjunctions of Büchi objectives. Let \mathcal{G} be a two-player turn-based game with a winning objective for **Eve** given as a conjunction of Büchi objectives $\Omega_1, \dots, \Omega_k$. The idea is to construct a new game \mathcal{G}' which checks that each objective Ω_i is visited infinitely often. The vertices of \mathcal{G}' controlled by **Eve** and **Adam** are $V'_\exists = V_\exists \times \llbracket 0, k \rrbracket$ and $V'_\forall = V_\forall \times \llbracket 0, k \rrbracket$ respectively. There is a transition from (v, k) to $(v', 0)$ iff there is a transition from v to v' in the original game and for $0 \leq i < k$, there is a transition from (v, i) to $(v', i + 1)$ iff there is a transition from v to v' in the original game and $v' \in \Omega_{i+1}$. In \mathcal{G}' , the objective for **Eve** is the Büchi objective Ω given by target set $\text{States} \times \{k\}$, where $\text{States} = V_\exists \cup V_\forall$ is the set of vertices of \mathcal{G} . It is clear that there is a winning strategy in \mathcal{G} from v_0 for the conjunction of Büchi objectives $\Omega_1, \dots, \Omega_k$ iff there is a winning strategy in \mathcal{G}' from $(v_0, 0)$ for the Büchi objective Ω . The number of states of game \mathcal{G}' is $|\text{States}'| = |\text{States}| \cdot k$, and the size of the transition table $|\text{Tab}'| = |\text{Tab}| \cdot k$. Using the standard algorithm for turn-based Büchi objectives [13], which works in time $\mathcal{O}(|\text{States}'| \cdot |\text{Tab}'|)$, we obtain an algorithm for solving zero-sum games with a conjunction of Büchi objectives running in time $\mathcal{O}(k^2 \cdot |\text{States}| \cdot |\text{Tab}|)$ (hence in polynomial time).

5.4.3. *Algorithm.* The algorithm is the same as for reachability objectives. Only the computation of the set of winning states in the suspect game is different. Since we just showed that this part can be done in polynomial time, the global algorithm still runs in (non-deterministic) polynomial time.

5.4.4. *Hardness.* The hardness result for the constrained NE existence problem with co-Büchi objectives was already proven in [40]. The idea is to encode an instance of 3SAT into a game with co-Büchi objectives. For completeness we describe the reduction below, and explain how it can be modified for proving NP-hardness of the NE existence problem.

Let us consider an instance $\phi = c_1 \wedge \dots \wedge c_n$ of SAT, where $c_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$, and $\ell_{i,j} \in \{x_k, \neg x_k \mid 1 \leq k \leq p\}$. The game \mathcal{G} is obtained from module $M(\phi)$ depicted on




FIGURE 13. Module $M(\phi)$, where $\phi = c_1 \wedge \dots \wedge c_n$ and $c_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$

Figure 13, by joining the outgoing edge of c_{n+1} to c_1 . Each module $M(\phi)$ involves a set of players B_k , one for each variable x_k , and a player A_1 . Player A_1 controls the clause states. Player B_k control the literal states $\ell_{i,j}$ when $\ell_{i,j} = \neg x_k$, then having the opportunity to go to state \perp . There is no transition to \perp for literals of the form x_k . In $M(\phi)$, assuming that the players B_k will not play to \perp , then A_1 has a strategy that does not visit both x_k and $\neg x_k$ for every k if, and only if, formula ϕ is satisfiable. Finally, the co-Büchi objective of B_k is given by $\{x_k\}$. In other terms, the aim of B_k is to visit x_k only a finite number of times. This way, in a Nash equilibrium, it cannot be the case that both x_k and $\neg x_k$ are visited infinitely often: it would imply that B_k loses but could improve her payoff by going to \perp (actually, $\neg x_k$ should not be visited at all if x_k is visited infinitely often). Therefore setting the objective of A_1 to $\{\perp\}$, there is a Nash equilibrium where she wins iff ϕ is satisfiable. This shows NP-hardness for the constrained NE existence problem.

For the NE existence problem, we use the transformation described in Section 3.4. We add an extra player A_2 to \mathcal{G} and consider the game $\mathcal{G}' = E(\mathcal{G}, A_1, A_2, \rho)$, where ρ is a winning path for A_1 . The objective of the players in \mathcal{G}' can be described by co-Büchi objectives: A_2 has to avoid seeing $T = \{s_1\}$ infinitely often and keep the same target for A_1 . Applying Proposition 3.5, there is a Nash equilibrium in \mathcal{G}' if, and only if, there is one in \mathcal{G} where A_1 wins, this shows NP-hardness for the NE existence problem.

5.5. Objectives given as circuits. The value problem is known to be PSPACE-complete for turn-based games and objectives given as circuits [27]. The transformation presented in the beginning of the section can be used to decide the value problem for finite concurrent games with a single circuit-objective, yielding PSPACE-completeness of the value problem in the case of finite concurrent games as well.

We now show that the (constrained) NE existence problem is also PSPACE-complete in this framework:

Theorem 5.15. *For finite games with single objectives given as circuits, the NE existence problem and the constrained NE existence problem are PSPACE-complete.*

5.5.1. Reduction to a circuit objective. We assume the preference relation of each player $A \in \text{Agt}$ is given by a circuit C_A . Let $L \subseteq \text{Agt}$. We define a Boolean circuit defining the winning condition of Eve in the suspect game $\mathcal{H}(\mathcal{G}, L)$.

We define for each player $A \in \text{Agt}$ and each set P of players (such that $\text{States} \times P$ is reachable in $\mathcal{H}(\mathcal{G}, L)$), a circuit $D_{A,P}$ which outputs **true** for the plays ρ with $\lambda(\rho) = P$ (i.e. whose states that are visited infinitely often are in $\text{States} \times \{P\}$), and whose value by C_A is **true**. We do so by making a copy of the circuit C_A , adding $|\text{States}|$ OR gates $g_1 \cdots g_{|\text{States}|}$ and one AND gate h . There is an edge from (s_i, P) to g_i and from g_{i-1} to g_i if $i < |\text{States}|$ then there is an edge from the output gate of C_A to h and from h to the output gate of the new circuit. Inputs of C_A are now the (s, P) 's (instead of the s 's). The circuit $D_{A,P}$ is given on Figure 14.




FIGURE 14. Circuit $D_{A,P}$

We then define a circuit E_A which outputs **true** for the plays ρ with $A \in \lambda(\rho)$ and whose output by C_A is **true**. We do so by taking the disjunction of the circuits $D_{A,P}$. Formally, for each set of players P such that $\text{States} \times P$ is reachable in the suspect game and $A \in P$, we include the circuit $D_{A,P}$ and writing $o_{A,P}$ for its output gate, we add OR gates so that there is an edge from $o_{A,P}$ to g_i and from g_i to g_{i+1} , and then from g_{n+1} to the output gate.

Finally we define the circuit F_L , which outputs **true** for the plays ρ such that there is no $A \in L$ such that $A \in \lambda(\rho)$ and the output of $\text{proj}_1(\rho)$ by C_A is **true**. This corresponds exactly to the plays that are winning for Eve in suspect game $\mathcal{H}(\mathcal{G}, L)$. We do so by negating the disjunction of all the circuits E_A for $A \in L$.

The next lemma follows from the construction:

Lemma 5.16. *A play ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$ iff ρ evaluates circuit F_L to **true**.*

We should notice that circuit F_L has size polynomial in the size of \mathcal{G} , thanks to Proposition 4.8.

5.5.2. Algorithm and complexity analysis. To solve the constrained NE existence problem we apply the same algorithm as for reachability objectives (see section 5.1). For complexity matters, the only difference stands in the computation of the set of winning states in the suspect game. Thanks to Lemma 5.16, we know it reduces to the computation of the set of winning states in a turn-based game with an objective given as a circuit (of polynomial-size). This can be done in PSPACE [27], which yields a PSPACE upper bound for the constrained NE existence problem (and therefore for the NE existence problem and the value problem – see Proposition 3.2). PSPACE-hardness of all problems follows from that of the value problem in turn-based games [27], and from Propositions 3.2 and 3.4 (we notice that the preference relations in the new games are easily definable by circuits).

5.6. Rabin and parity objectives. The value problem is known to be NP-complete for Rabin conditions [18] and in $\text{UP} \cap \text{co-UP}$ for parity conditions [28].

We then notice that a parity condition is a Rabin condition with half as many pairs as the number of priorities: assume the parity condition is given by p : States $\mapsto \llbracket 0, d \rrbracket$ with $d \in \mathbb{N}$; take for i in $\llbracket 0, \frac{d}{2} \rrbracket$, $Q_i = p^{-1}\{2i\}$ and $R_i = p^{-1}\{2j + 1 \mid j \geq i\}$. Then the Rabin objective $(Q_i, R_i)_{0 \leq i \leq \frac{d}{2}}$ is equivalent to the parity condition given by p .

We design an algorithm that solves the constrained NE existence problem in $\mathsf{P}_{\parallel}^{\text{NP}}$ for Rabin objectives (see footnote 1 on page 3 for an informal definition of $\mathsf{P}_{\parallel}^{\text{NP}}$).

Our algorithm heavily uses non-determinism (via the oracle). We then propose a deterministic algorithm which runs in exponential time, but will be useful in Section 5.7. This subsection ends with proving $\mathsf{P}_{\parallel}^{\text{NP}}$ -hardness of the constrained NE existence problem and NE existence problem for parity objectives. In the end, we will have proven the following theorem:

Theorem 5.17. *For finite games with single objectives given as Rabin or parity conditions, the NE existence problem and the constrained NE existence problem are $\mathsf{P}_{\parallel}^{\text{NP}}$ -complete.*

5.6.1. Reduction to a Streett game. We assume that the preference relation of each player $A \in \text{Agt}$ is given by the Rabin condition $(Q_{i,A}, R_{i,A})_{i \in \llbracket 1, k_A \rrbracket}$. Let $L \subseteq \text{Agt}$. In the suspect game $\mathcal{H}(\mathcal{G}, L)$, we define the Streett objective $(Q'_{i,A}, R'_{i,A})_{i \in \llbracket 1, k_A \rrbracket, A \in L}$, where $Q'_{i,A} = (Q_{i,A} \times \{P \mid A \in P\}) \cup (\text{States} \times \{P \mid A \notin P\})$ and $R'_{i,A} = R_{i,A} \times \{P \mid A \in P\}$, and we write Ω_L for the corresponding set of winning plays.

Lemma 5.18. *A play ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$ iff $\rho \in \Omega_L$.*

Proof. Assume ρ is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$. For all $A \in \lambda(\rho) \cap L$, $\text{proj}_1(\rho)$ does not satisfy the Rabin condition given by $(Q_{i,A}, R_{i,A})_{i \in \llbracket 1, k_A \rrbracket}$. For all $1 \leq i \leq k_A$, $\text{Inf}(\text{proj}_1(\rho)) \cap Q_{i,A} = \emptyset$ or $\text{Inf}(\text{proj}_1(\rho)) \cap R_{i,A} \neq \emptyset$. We infer that for all $1 \leq i \leq k_A$, $\text{Inf}(\rho) \cap Q'_{i,A} = \emptyset$ or $\text{Inf}(\rho) \cap R'_{i,A} \neq \emptyset$. Now, if $A \notin \lambda(\rho)$ then all $Q'_{i,A}$ are seen infinitely often along ρ . Therefore for every $A \in L$, the Streett conditions $(Q'_{i,A}, R'_{i,A})$ is satisfied along ρ (that is, $\rho \in \Omega_L$).

Conversely, if the Streett condition $(Q'_{i,A}, R'_{i,A})_{i \in \llbracket 1, k_A \rrbracket, A \in L}$ is satisfied along ρ , then either the Rabin condition $(Q_{i,A}, R_{i,A})$ is not satisfied along $\text{proj}_1(\rho)$ or $A \notin \lambda(\rho)$. This means that Eve is winning in $\mathcal{H}(\mathcal{G}, L)$. \square

5.6.2. Algorithm. We now describe a $\mathsf{P}_{\parallel}^{\text{NP}}$ algorithm for solving the constrained NE existence problem in games where each player has a single Rabin objective. As in the previous cases, our algorithm relies on the suspect game construction.

Write \mathcal{P} for the set of sets of players of Agt that appear as the second item of a state of $\mathcal{J}(\mathcal{G})$:

$$\mathcal{P} = \{P \subseteq \text{Agt} \mid \exists s \in \text{States}. (s, P) \text{ is a state of } \mathcal{J}(\mathcal{G})\}.$$

Since $\mathcal{J}(\mathcal{G})$ has size polynomial, so has \mathcal{P} . Also, for any path ρ , $\lambda(\rho)$ is a set of \mathcal{P} . Hence, for a fixed L , the number of sets $\lambda(\rho) \cap L$ is polynomial. Now, as recalled on page 25, the winning condition for Eve is that the players in $\lambda(\rho) \cap L$ must be losing along $\text{proj}_1(\rho)$ in \mathcal{G} for their Rabin objective. We have seen that this can be seen as a Streett objective (Lemma 5.18).

Now, deciding whether a state is winning in a turn-based game for a Streett condition can be decided in coNP [18]. Hence, given a state $s \in \text{States}$ and a set L , we can decide in coNP whether s is winning for Eve in $\mathcal{H}(\mathcal{G}, L)$. This will be used as an oracle in our algorithm below.

Now, pick a set $P \subseteq \text{Agt}$ of suspects, i.e., for which there exists $(s, t) \in \text{States}^2$ and m_{Agt} s.t. $P = \text{Susp}((s, t), m_{\text{Agt}})$. Using the same arguments as in the proof of Proposition 4.8, it can be shown that $2^{|P|} \leq |\text{Tab}|$, so that the number of subsets of P is polynomial. Now, for each set P of suspects and each $L \subseteq P$, write $w(L)$ for the size of the winning region of Eve in $\mathcal{H}(\mathcal{G}, L)$. Then the sum $\sum_{P \in \mathcal{P} \setminus \{\text{Agt}\}} \sum_{L \subseteq P} w(L)$ is at most $|\text{States}| \times |\text{Tab}|^2$.

Assume that the exact value M of this sum is known, and consider the following algorithm:

- (1) for each $P \subseteq \mathcal{P} \setminus \{\text{Agt}\}$ and each $L \subseteq P$, guess a set $W(L) \subseteq \text{States}$, which we intend to be the exact winning region for Eve in $\mathcal{H}(\mathcal{G}, L)$.
- (2) check that the sizes of those sets sum up to M ;
- (3) for each $s \notin W(L)$, check that Eve does not have a winning strategy from s in $\mathcal{H}(\mathcal{G}, L)$. This can be checked in non-deterministic polynomial time, as explained above.
- (4) guess a lasso-shaped path $\rho = \pi \cdot \tau^\omega$ in $\mathcal{H}(\mathcal{G}, L)$ starting from (s, Agt) , with $|\pi|$ and $|\tau|$ less than $|\text{States}|^2$ (following Proposition 3.1) visiting only states where the second item is Agt. This path can be seen as the outcome of some strategy of Eve when Adam obeys. For this path, we then check the following:
 - along ρ , the sets of winning and losing players satisfy the original constraint (remember that we aim at solving the constrained NE existence problem);
 - any deviation along ρ leads to a state that is winning for Eve. In other terms, pick a state $h = (s, \text{Agt}, m_{\text{Agt}})$ of Adam along ρ , and pick a successor $h' = (t, P)$ of h such that $t \neq \text{Tab}(s, m_{\text{Agt}})$. Then the algorithm checks that $t \in W(L \cap P)$.

The algorithm accepts the input M if it succeeds in finding the sets W and the path ρ such that all the checks are successful. This algorithm is non-deterministic and runs in polynomial time, and will be used as a second oracle.

We now show that if M is exactly the sum of the $w(L)$, then the algorithm accepts M if, and only if, there is a Nash equilibrium satisfying the constraint, i.e., if, and only if, Eve has a winning strategy from (s, Agt) in $\mathcal{H}(\mathcal{G}, L)$.

First assume that the algorithm accepts M . This means that it is able, for each L , to find sets $W(L)$ of states whose complement does not intersect the winning region of $\mathcal{H}(\mathcal{G}, L)$. Since M is assumed to be the exact sum of $w(L)$ and the size of the sets $W(L)$ sum up to M , we deduce that $W(L)$ is exactly the winning region of Eve in $\mathcal{H}(\mathcal{G}, L)$. Now, since the algorithm accepts, it is also able to find a (lasso-shaped) path ρ only visiting states having Agt as the second component. This path has the additional property that any “deviation” from a state of Adam along this path ends up in a state that is winning for Eve for players in $L \cap P$, where P is the set of suspects for the present deviation. This way, if during ρ , Adam deviates to a state (t, P) , then Eve will have a strategy to ensure that along any subsequent play, the objectives of players in $L \cap P$ (in \mathcal{G}) are not fulfilled, so that along any run ρ' , the players in $L \cap \lambda(\rho')$ are losing for their objectives in \mathcal{G} , so that Eve wins in $\mathcal{H}(\mathcal{G}, L)$.

Conversely, assume that there is a Nash equilibrium satisfying the constraint. Following Proposition 3.1, we assume that the outcome of the corresponding strategy profile has the form $\pi \cdot \tau^\omega$. From Lemma 4.4, there is a winning strategy for Eve in $\mathcal{H}(\mathcal{G}, L)$ whose

outcome when Adam obeys follows the outcome of the Nash equilibrium. As a consequence, the outcome when Adam obeys is a path ρ that the algorithm can guess. Indeed, it must satisfy the constraints, and any deviation from ρ with set of suspects P ends in a state where Eve wins for the winning condition of $\mathcal{H}(\mathcal{G}, L)$, hence also for the winning condition of $\mathcal{H}(\mathcal{G}, L \cap P)$, since any path ρ' visiting (t, P) has $\lambda(\rho') \subseteq P$.

Finally, our global algorithm is as follows: we run the first oracle for all the states and all the sets L that are subsets of a set of suspects (we know that there are polynomially many such inputs). We also run the second algorithm on all the possible values for M , which are also polynomially many. Now, from the answers of the first oracle, we compute the exact value M , and return the value given by the second on that input. This algorithm runs in $\mathbb{P}_{\parallel}^{\text{NP}}$ and decides the constrained NE existence problem.

5.6.3. Deterministic algorithm. In the next section we will need a deterministic algorithm to solve games with objectives given as deterministic Rabin automata. We therefore present it right now. The deterministic algorithm works by successively trying all the possible payoffs, there are $2^{|\text{Agt}|}$ of them. Then it computes the winning strategies of the suspect game for that payoff. In [25] an algorithm for Streett games is given, which works in time $\mathcal{O}(n^k \cdot k!)$, where n is the number of vertices in the game, and k the size of the Streett condition. The algorithm has to find, in the winning region of Eve in $\mathcal{J}(\mathcal{G})$, a lasso that satisfies the Rabin winning conditions of the winners and do not satisfy those of the losers. To do so it tries all the possible choices of elementary Rabin condition that are satisfied to make the players win, there are at most $\prod_{A \in \text{Agt}} k_A$ possible choices. And for the losers, we try the possible choices for whether $Q_{i,A}$ is visited or not, there are $\prod_{A \in \text{Agt}} 2^{k_A}$ such choices. It then looks for a lasso cycle that, when A is a winner, does not visit $Q_{i,A}$ and visits $R_{i,A}$, and when A is a loser, visits $R_{i,A}$ when it has to, or does not visit $Q_{i,A}$. This is equivalent to finding a path satisfying a conjunction of Büchi conditions and can be done in polynomial time $\mathcal{O}(n \times \sum_{A \in \text{Agt}} k_A)$. The global algorithm works in time

$$\mathcal{O} \left(2^{|\text{Agt}|} \cdot \left(|\text{Tab}|^{3 \sum_A k_A} \cdot \left(\sum_A k_A \right)! + \left(\prod_{A \in \text{Agt}} k_A \cdot 2^{k_A} \right) \cdot |\text{Tab}|^3 \cdot \sum_A k_A \right) \right)$$

Notice that the exponential does not come from the size of the graph but from the number of agents and the number of elementary Rabin conditions, this will be important when in the next subsection we will reuse the algorithm on a game structure whose size is exponential.

5.6.4. $\mathbb{P}_{\parallel}^{\text{NP}}$ -hardness. We now prove $\mathbb{P}_{\parallel}^{\text{NP}}$ -hardness of the (constrained) NE existence problem in the case of parity objectives. The main reduction is an encoding of the $\oplus\text{SAT}$ problem, where the aim is to decide whether the number of satisfiable instances among a set of formulas is even. This problem is known to be complete for $\mathbb{P}_{\parallel}^{\text{NP}}$ [22].

Before tackling the whole reduction, we first develop some preliminaries on single instances of SAT, inspired from [12]. Let us consider an instance $\phi = c_1 \wedge \dots \wedge c_n$ of SAT, where $c_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$, and $\ell_{i,j} \in \{x_k, \neg x_k \mid 1 \leq k \leq p\}$. With ϕ , we associate a three-player game $N(\phi)$, depicted on Figure 15 (where the first state of $N(\phi)$ is controlled

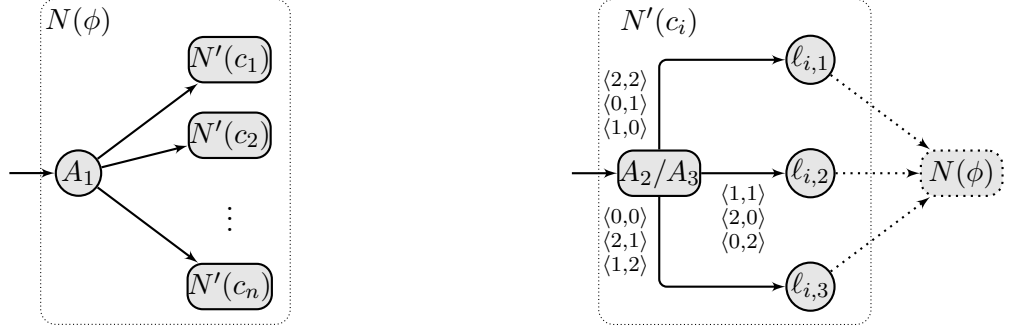


FIGURE 15. The game $N(\phi)$ (left), where $N'(c_i)$ is the module on the right.

by A_1 , and the first state of each $N'(c_j)$ is concurrently controlled by A_2 and A_3). For each variable x_j , players A_2 and A_3 have the following target sets:

$$T_{2j}^{A_2} = \{x_j\} \quad T_{2j+1}^{A_2} = \{\neg x_j\} \quad T_{2j+1}^{A_3} = \{x_j\} \quad T_{2j}^{A_3} = \{\neg x_j\}$$

This construction enjoys interesting properties, given by the following lemma:

Lemma 5.19. *If the formula ϕ is not satisfiable, then there is a strategy for player A_1 in $N(\phi)$ such that players A_2 and A_3 lose. If the formula ϕ is satisfiable, then for any strategy profile σ_{Agt} , one of A_2 and A_3 can change her strategy and win.*

Proof. We begin with the first statement, assuming that ϕ is not satisfiable and defining the strategy for A_1 . With a history h in $N(\phi)$, we associate a valuation $v^h: \{x_k \mid k \in [1, p]\} \rightarrow \{\top, \perp\}$ (where p is the number of distinct variables in ϕ), defined as follows:

$$v^h(x_k) = \top \Leftrightarrow \exists m. h_m = x_k \wedge \forall m' > m. h_{m'} \neq \neg x_k \quad \text{for all } k \in [1, p]$$

We also define $v^h(\neg x_k) = \neg v^h(x_k)$. Under this definition, $v^h(x_k) = \top$ if the last occurrence of x_k or $\neg x_k$ along h was x_k . We then define a strategy σ_1 for player A_1 : after a history h ending in an A_1 -state, we require $\sigma_1(h)$ to go to $N'(c_i)$ for some c_i (with least index, say) that evaluates to false under v^h (such a c_i exists since ϕ is not satisfiable). This strategy enforces that if $h \cdot \sigma_1(h) \cdot \ell_{i,j}$ is a finite outcome of σ_1 , then $v^h(\ell_{i,j}) = \perp$, because A_1 has selected a clause c_i whose literals all evaluate to \perp . Moreover, $v^{h \cdot \sigma_1(h) \cdot \ell_{i,j}}(\ell_{i,j}) = \top$, so that for each j , any outcome of σ_1 will either alternate between x_k and $\neg x_k$ (hence visit both of them infinitely often), or no longer visit any of them after some point. Hence both A_2 and A_3 lose.

We now prove the second statement. Let v be a valuation under which ϕ evaluates to true, and σ_{Agt} be a strategy profile. From σ_{A_2} and σ_{A_3} , we define two strategies σ'_{A_2} and σ'_{A_3} . Consider a finite history h ending in the first state of $N'(c_i)$, for some i . Pick a literal $\ell_{i,j}$ of c_i that is true under v (the one with least index, say). We set

$$\sigma'_{A_2}(h) = [j - \sigma_{A_3}(h) \pmod{3}] \quad \sigma'_{A_3}(h) = [j - \sigma_{A_2}(h) \pmod{3}].$$

It is easily checked that, when σ_{A_2} and σ'_{A_3} (or σ'_{A_2} and σ_{A_3}) are played simultaneously in the first state of some $N'(c_i)$, then the game goes to $\ell_{i,j}$. Thus under those strategies, any visited literal evaluates to true under v , which means that at most one of x_k and $\neg x_k$ is visited (infinitely often). Hence one of A_2 and A_3 is winning, which proves our claim. \square

We now proceed by encoding an instance

$$\exists x_1^1, \dots, x_k^1. \phi^1(x_1^1, \dots, x_k^1)$$

...

$$\exists x_1^m, \dots, x_k^m. \phi^m(x_1^m, \dots, x_k^m)$$

of $\oplus\text{SAT}$ into a parity game. The game involves the three players A_1 , A_2 and A_3 of the game $N(\phi)$ defined above, and it will contain a copy of $N(\phi^r)$ for each $1 \leq r \leq m$. The objectives of A_2 and A_3 are the unions of their objectives in each $N(\phi^r)$, e.g. $p^{A_2}(x_j^1) = p^{A_2}(x_j^2) = \dots = p^{A_m}(x_j^m) = 2j$.

For each such r , the game will also contain a copy of the game $M(\phi^r)$ depicted on Figure 13. Each game $M(\phi^r)$ involves an extra set of players B_k^r , one for each variable x_k^r . As we have seen in Section 5.4, in a Nash equilibrium, it cannot be the case that both x_k^r and $\neg x_k^r$ are visited infinitely often.

In order to test the parity of the number of satisfiable formulas, we then define two families of modules, depicted on Figure 16 to 19. Finally, the whole game \mathcal{G} is depicted on Figure 20. In that game, the objective of A_1 is to visit infinitely often the initial state init .




FIGURE 16. Module $H(\phi^r)$ for $r \geq 2$




FIGURE 17. Module $G(\phi^r)$ for $r \geq 2$




FIG. 18. Module $H(\phi^1)$




FIG. 19. Module $G(\phi^1)$




FIG. 20. The game \mathcal{G}

Lemma 5.20. *There is a Nash equilibrium in the game \mathcal{G} where A_2 and A_3 lose and A_1 wins if, and only if, the number of satisfiable formulas is even.*

Proof. Assume that there is a Nash equilibrium in \mathcal{G} where A_1 wins and both A_2 and A_3 lose. Let ρ be its outcome. As already noted, if ρ visits module $M(\phi^r)$ infinitely often, then it cannot be the case that both x_k^r and $\neg x_k^r$ are visited infinitely often in $M(\phi^r)$, as otherwise B_k^r would be losing and have the opportunity to improve her payoff. This implies that ϕ^r is satisfiable. Similarly, if ρ visits infinitely often the states of $H(\phi^r)$ or $G(\phi^r)$ that is controlled by A_2 and A_3 , then it must be the case that ϕ^r is not satisfiable, since from Lemma 5.19 this would imply that A_2 or A_3 could deviate and improve her payoff by going to $N(\phi^r)$.

We now show by induction on r that if ρ goes infinitely often in module $G(\phi^r)$ then $\#\{j \leq r \mid \phi^j \text{ is satisfiable}\}$ is even, and that (if $n > 1$) this number is odd if ρ goes infinitely in module $H(\phi^r)$.

When $r = 1$, since $H(\phi^1)$ is $M(\phi^1)$, ϕ^1 is satisfiable, as noted above. Similarly, if ρ visits $G(\phi^1)$ infinitely often, it also visits its A_2/A_3 -state infinitely often, so that ϕ^1 is not satisfiable. This proves the base case.

Assume that the result holds up to some $r - 1$, and assume that ρ visits $G(\phi^r)$ infinitely often. Two cases may occur:

- it can be the case that $M(\phi^r)$ is visited infinitely often, as well as $H(\phi^{r-1})$. Then ϕ^r is satisfiable, and the number of satisfiable formulas with index less than or equal to $r - 1$ is odd. Hence the number of satisfiable formulas with index less than or equal to r is even.
- it can also be the case that the state A_2/A_3 of $G(\phi^r)$ is visited infinitely often. Then ϕ^r is not satisfiable. Moreover, since A_1 wins, the play will also visit $G(\phi^{r-1})$ infinitely often, so that the number of satisfiable formulas with index less than or equal to r is even.

If ρ visits $H(\phi^r)$ infinitely often, using similar arguments we prove that the number of satisfiable formulas with index less than or equal to r is odd.

To conclude, since A_1 wins, the play visits $G(\phi^m)$ infinitely often, so that the total number of satisfiable formulas is even.

Conversely, assume that the number of satisfiable formulas is even. We build a strategy profile, which we prove is a Nash equilibrium in which A_1 wins, and A_2 and A_3 lose. The strategy for A_1 in the initial states of $H(\phi^r)$ and $G(\phi^r)$ is to go to $M(\phi^r)$ when ϕ^r is satisfiable, and to state A_2/A_3 otherwise. In $M(\phi^r)$, the strategy is to play according to a valuation satisfying ϕ^r . In $N(\phi^r)$, it follows a strategy along which A_2 and A_3 lose (this exists according to Lemma 5.19). This defines the strategy for A_1 . Then A_2 and A_3 are required to always play the same move, so that the play never goes to some $N(\phi^r)$. In $N(\phi^r)$, they can play any strategy (they lose anyway, whatever they do). Finally, the strategy of B_k^r never goes to \perp .

We now explain why this is the Nash equilibrium we are after. First, as A_1 plays according to fixed valuations for the variables x_k^r , either B_k^r wins or she does not have the opportunity to go to \perp . It remains to prove that A_1 wins, and that A_2 and A_3 lose and cannot improve (individually). To see this, notice that between two consecutive visits to init, exactly one of $G(\phi^r)$ and $H(\phi^r)$ is visited. More precisely, it can be observed that the strategy of A_1 enforces that $G(\phi^r)$ is visited if $\#\{r < r' \leq m \mid \phi^{r'} \text{ is satisfiable}\}$ is even, and that $H(\phi^r)$ is visited otherwise. Then if $H(\phi_1)$ is visited, the number of satisfiable formulas with index between 2 and m is odd, so that ϕ_1 is satisfiable and A_1 can return to init. If $G(\phi^1)$ is visited, an even number of formulas with index between 2 and m is satisfiable, and ϕ^1 is not. Hence A_1 has a strategy in $N(\phi^1)$ to make A_2 and A_3 lose, so that A_2 and A_3 cannot improve their payoffs. \square

This proves hardness for the constrained NE existence problem with parity objectives. For the NE existence problem, we use the construction of Section 3.4, but since it can only be used to get rid of constraint of the type “ A_1 is winning”, we add to the game two players, A_4 and A_5 , whose objectives are opposite to A_2 and A_3 respectively, and one player A_6 that will be playing matching-penny games. The objectives for A_4 and A_5 are definable by parity objectives, by adding 1 to all the priorities. Then, we consider game $\mathcal{G}' = E(E(E(\mathcal{G}, A_1, A_6, \rho_1), A_4, A_6, \rho_4), A_5, A_6, \rho_5)$ where ρ_1 , ρ_4 and ρ_5 are winning paths for A_1 , A_4 and A_5 respectively. Thanks to Proposition 3.5, there is a Nash equilibrium in \mathcal{G}' if,

and only if, there is a Nash equilibrium in \mathcal{G} where A_1 wins and A_2 and A_3 lose. We deduce $P_{\parallel}^{\text{NP}}$ -hardness for the NE existence problem with parity objectives.

5.7. Objectives given as deterministic Rabin automata. In order to find Nash equilibria when objectives are given as deterministic Rabin automata, we first define the notion of *game simulation*, which we show has the property that when \mathcal{G}' game-simulates \mathcal{G} , then a Nash equilibrium in the latter game gives rise to a Nash equilibrium in the former one.

We then define the product of a game with automata (defining the objectives of the players), and show that it game-simulates the original game. This reduces the case of games with objectives defined as Rabin automata to games with Rabin objectives, which we handled at the previous section; the resulting algorithm is in EXPTIME. We then show a PSPACE lower bound for the problem in the case of objectives given as deterministic Büchi automata. This proves the following theorem:

Theorem 5.21. *For finite games with single objectives given as deterministic Rabin automata or deterministic Büchi automata, the NE existence problem and the constrained NE existence problem are in EXPTIME and PSPACE-hard.*

It must be noticed that game simulation can be used in other contexts: in particular, in [4] (where we introduced this notion), it is shown that a region-based abstraction of timed games game simulates its original timed game, which provides a way of computing Nash equilibria in timed games.

5.7.1. Game simulation. We define game simulation and show how that can be used to compute Nash equilibria. We then apply it to objectives given as deterministic Rabin automata.

Definition 5.22. Consider two games $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ and $\mathcal{G}' = \langle \text{States}', \text{Agt}, \text{Act}', \text{Mov}', \text{Tab}', (\preceq'_A)_{A \in \text{Agt}} \rangle$ with the same set Agt of players. A relation $\triangleleft \subseteq \text{States} \times \text{States}'$ is a *game simulation* if $s \triangleleft s'$ implies that for each move m_{Agt} in \mathcal{G} there exists a move m'_{Agt} in \mathcal{G}' such that:

- (1) $\text{Tab}(s, m_{\text{Agt}}) \triangleleft \text{Tab}'(s', m'_{\text{Agt}})$, and
- (2) for each $t' \in \text{States}'$ there exists $t \in \text{States}$ with $t \triangleleft t'$ and $\text{Susp}((s', t'), m'_{\text{Agt}}) \subseteq \text{Susp}((s, t), m_{\text{Agt}})$.

If \triangleleft is a game simulation and $(s_0, s'_0) \in \triangleleft$, we say that \mathcal{G}' *game-simulates* (or simply *simulates*) \mathcal{G} . When there are two paths ρ and ρ' such that $\rho_{=i} \triangleleft \rho'_{=i}$ for all $i \in \mathbb{N}$, we will simply write $\rho \triangleleft \rho'$.

A game simulation \triangleleft is *preference-preserving* from $(s_0, s'_0) \in \text{States} \times \text{States}'$ if for all $\rho_1, \rho_2 \in \text{Play}_{\mathcal{G}}(s_0)$ and $\rho'_1, \rho'_2 \in \text{Play}_{\mathcal{G}'}(s'_0)$ with $\rho_1 \triangleleft \rho'_1$ and $\rho_2 \triangleleft \rho'_2$, for all $A \in \text{Agt}$ it holds that $\rho_1 \preceq_A \rho_2$ iff $\rho'_1 \preceq'_A \rho'_2$.

As we show now, Nash equilibria are preserved by game simulation, in the following sense:

Proposition 5.23. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ and $\mathcal{G}' = \langle \text{States}', \text{Agt}, \text{Act}', \text{Mov}', \text{Tab}', (\preceq'_A)_{A \in \text{Agt}} \rangle$ be two games involving the same set of players. Fix two states s_0 and s'_0 in \mathcal{G} and \mathcal{G}' respectively, and let \triangleleft be a preference-preserving game simulation from (s_0, s'_0) . If there exists a Nash equilibrium σ_{Agt} in \mathcal{G} from s_0 , then there exists a Nash equilibrium σ'_{Agt} in \mathcal{G}' from s'_0 with $\text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}}) \triangleleft \text{Out}_{\mathcal{G}'}(s'_0, \sigma'_{\text{Agt}})$.*

Proof. We fix a strategy profile σ_{Agt} in \mathcal{G} and ρ the outcome of σ_{Agt} from s_0 . We derive a strategy profile σ'_{Agt} in \mathcal{G}' and its outcome ρ' from s'_0 , such that:

- (a) for every $\bar{\rho}' \in \text{Play}_{\mathcal{G}'}(s'_0)$, there exists $\bar{\rho} \in \text{Play}_{\mathcal{G}}(s_0)$ s.t. $\bar{\rho} \triangleleft \bar{\rho}'$ and $\text{Susp}(\bar{\rho}', \sigma'_{\text{Agt}}) \subseteq \text{Susp}(\bar{\rho}, \sigma_{\text{Agt}})$;
- (b) $\rho \triangleleft \rho'$.

Assume we have done the construction, and that σ_{Agt} is a Nash equilibrium in \mathcal{G} . We prove that σ'_{Agt} is a Nash equilibrium in \mathcal{G}' . Towards a contradiction, assume that some player A has a strategy $\bar{\sigma}'_A$ in \mathcal{G}' such that $\bar{\rho}' \not\prec_A \rho'$, where $\bar{\rho}' = \text{Out}_{\mathcal{G}'}(s', \sigma'_{\text{Agt}}[A \mapsto \bar{\sigma}'_A])$. Note that $A \in \text{Susp}(\bar{\rho}', \sigma'_{\text{Agt}})$. Applying ((a)) above, there exists $\bar{\rho} \in \text{Play}_{\mathcal{G}}(s_0)$ such that $\bar{\rho} \triangleleft \bar{\rho}'$ and $\text{Susp}(\bar{\rho}', \sigma'_{\text{Agt}}) \subseteq \text{Susp}(\bar{\rho}, \sigma_{\text{Agt}})$. In particular, $A \in \text{Susp}(\bar{\rho}, \sigma_{\text{Agt}})$, and there exists a strategy $\bar{\sigma}_A$ for A such that $\bar{\rho} = \text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}}[A \mapsto \bar{\sigma}_A])$. As $\rho \triangleleft \rho'$ (by ((b))) and \triangleleft is preference-preserving from (s_0, s'_0) , $\bar{\rho} \not\prec_A \rho$, which contradicts the fact that σ_{Agt} is a Nash equilibrium. Hence, σ'_{Agt} is a Nash equilibrium in \mathcal{G}' from s'_0 .

It remains to show how we construct σ'_{Agt} (and ρ'). We first build ρ' inductively, and define σ'_{Agt} along that path.

- Initially, we let $\rho'_{=0} = s'_0$. Since \triangleleft is a game simulation containing (s_0, s'_0) , we have $s_0 \triangleleft s'_0$, and there is a move m'_{Agt} associated with $\sigma_{\text{Agt}}(s_0)$ satisfying the conditions of Definition 5.22. Then $\rho_{=0} \triangleleft \rho'_{=0}$, and $\text{Susp}(\rho'_{=0}, \sigma'_{\text{Agt}}(\rho'_{=0})) \subseteq \text{Susp}(\rho_{=0}, \sigma_{\text{Agt}}(\rho_{=0}))$.
- Assume we have built $\rho'_{\leq i}$ and σ'_{Agt} on all the prefixes of $\rho'_{\leq i}$, and that they are such that $\rho_{\leq i} \triangleleft \rho'_{\leq i}$ and $\text{Susp}(\rho'_{\leq i}, \sigma'_{\text{Agt}}) \subseteq \text{Susp}(\rho_{\leq i}, \sigma_{\text{Agt}})$ (notice that $\text{Susp}(\rho'_{\leq i}, \sigma'_{\text{Agt}})$ only depends on the value of σ'_{Agt} on all the prefixes of $\rho_{\leq i}$). In particular, we have $\rho_{=i} \triangleleft \rho'_{=i}$, so that with the move $\sigma_{\text{Agt}}(\rho_{\leq i})$, we can associate a move m'_{Agt} (to which we set $\sigma'_{\text{Agt}}(\rho'_{\leq i})$) satisfying both conditions of Definition 5.22. This defines $\rho'_{=i+1}$ in such a way that $\rho_{\leq i+1} \triangleleft \rho'_{\leq i+1}$; moreover, $\text{Susp}(\rho'_{\leq i+1}, \sigma'_{\text{Agt}}) = \text{Susp}(\rho'_{\leq i}, \sigma'_{\text{Agt}}) \cap \text{Susp}((\rho'_{=i}, \rho'_{=i+1}), m'_{\text{Agt}})$ is indeed a subset of $\text{Susp}(\rho_{\leq i+1}, \sigma_{\text{Agt}})$.

It remains to define σ'_{Agt} outside its outcome ρ' . Notice that, for our purposes, it suffices to define σ'_{Agt} on histories starting from s'_0 . We again proceed by induction on the length of the histories, defining σ'_{Agt} in order to satisfy ((a)) on prefixes of plays of \mathcal{G}' from s'_0 . At each step, we also make sure that for every $h' \in \text{Hist}_{\mathcal{G}'}(s'_0)$, there exists $h \in \text{Hist}_{\mathcal{G}}(s)$ such that $h \triangleleft h'$, $\text{Susp}(h', \sigma'_{\text{Agt}}) \subseteq \text{Susp}(h, \sigma_{\text{Agt}})$, and $\sigma_{\text{Agt}}(h)$ and $\sigma'_{\text{Agt}}(h')$ satisfy the conditions of Definition 5.22 in the last states of h and h' , resp.

As we only consider histories from s'_0 , the case of histories of length zero was already handled. Assume we have defined σ'_{Agt} for histories h' of length i , and fix a new history $h' \cdot t' \in \text{Hist}_{\mathcal{G}'}(s'_0)$ of length $i + 1$ (that is not a prefix of ρ'). By induction hypothesis, there is $h \in \text{Hist}_{\mathcal{G}}(s_0)$ such that $h \triangleleft h'$, and $\text{Susp}(h', \sigma'_{\text{Agt}}) \subseteq \text{Susp}(h, \sigma_{\text{Agt}})$, and $\sigma_{\text{Agt}}(h)$ and $\sigma_{\text{Agt}}(h')$ satisfy the required properties. In particular, with t' , we can associate t s.t. $t \triangleleft t'$ and $\text{Susp}((\text{last}(h'), t'), \sigma'_{\text{Agt}}(h')) \subseteq \text{Susp}((\text{last}(h), t), \sigma_{\text{Agt}}(h))$. Then $(h \cdot t) \triangleleft (h' \cdot t')$. Since $t \triangleleft t'$, there is a move m'_{Agt} associated with $\sigma_{\text{Agt}}(h \cdot t)$ and satisfying the conditions of Definition 5.22. Letting $\sigma'_{\text{Agt}}(h' \cdot t') = m'_{\text{Agt}}$, we fulfill all the requirements of our induction hypothesis.

We now need to lift the property from histories to infinite paths. Consider a play $\bar{\rho}' \in \text{Play}_{\mathcal{G}'}(s'_0)$, we will construct a corresponding play $\bar{\rho}$ in \mathcal{G} . Set $\bar{\rho}_0 = s_0$. If $\bar{\rho}$ has been defined up to index i and $\bar{\rho}_i \triangleleft \bar{\rho}'_i$ (this is true for $i = 0$), thanks to the way σ'_{Agt} is constructed, $\sigma_{\text{Agt}}(\bar{\rho}_{\leq i})$ and $\sigma'_{\text{Agt}}(\bar{\rho}'_{\leq i})$ satisfy the conditions of Definition 5.22 in $\bar{\rho}_{\leq i}$ and

$\bar{\rho}'_i$, respectively. We then pick $\bar{\rho}_{i+1}$ such that $\bar{\rho}_{i+1} \triangleleft \bar{\rho}'_{i+1}$ and $\text{Susp}((\bar{\rho}_i, \bar{\rho}_{i+1}), \sigma_{\text{Agt}}(\bar{\rho}_i)) \subseteq \text{Susp}((\bar{\rho}'_i, \bar{\rho}'_{i+1}), \sigma'_{\text{Agt}}(\bar{\rho}'_i))$. This being true at each step, the path $\bar{\rho}$ that is obtained, is such that $\bar{\rho} \triangleleft \bar{\rho}'$ and $\text{Susp}(\bar{\rho}', \sigma'_{\text{Agt}}) \subseteq \text{Susp}(\bar{\rho}, \sigma_{\text{Agt}})$. This is the desired property. \square

5.7.2. Product of a game with deterministic Rabin automata. After this digression on game simulation, we come back to the game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\lesssim_A)_{A \in \text{Agt}} \rangle$, where we assume that some player A has her objective given by a deterministic Rabin automaton $\mathcal{A} = \langle Q, \text{States}, \delta, q_0, (Q_i, R_i)_{i \in [1, n]} \rangle$ (recall that this automaton reads sequences of states of \mathcal{G} , and accepts the paths that are winning for player A). We show how to compute Nash equilibria in \mathcal{G} by building a product \mathcal{G}' of \mathcal{G} with the automaton \mathcal{A} and by computing the Nash equilibria in the resulting game, with a Rabin winning condition for A .

We define the product of the game \mathcal{G} with the automaton \mathcal{A} as the game $\mathcal{G} \times \mathcal{A} = \langle \text{States}', \text{Agt}, \text{Act}, \text{Mov}', \text{Tab}', (\lesssim'_A)_{A \in \text{Agt}} \rangle$, where:

- $\text{States}' = \text{States} \times Q$;
- $\text{Mov}'((s, q), A_j) = \text{Mov}(s, A_j)$ for every $A_j \in \text{Agt}$;
- $\text{Tab}'((s, q), m_{\text{Agt}}) = (s', q')$ where $\text{Tab}(s, m_{\text{Agt}}) = s'$ and $\delta(q, s) = q'$;
- If $B = A$ then \lesssim'_B is given by the internal Rabin condition $Q'_i = \text{States} \times Q_i$ and $R'_i = \text{States} \times R_i$. Otherwise \lesssim'_B is derived from \lesssim_B , defined by $\rho \lesssim'_B \bar{\rho}$ if, and only if, $\text{proj}(\rho) \lesssim_B \text{proj}(\bar{\rho})$ (where proj is the projection of States' on States). Notice that if \lesssim_B is an internal Rabin condition, then so is \lesssim'_B .

Lemma 5.24. $\mathcal{G} \times \mathcal{A}$ game-simulates \mathcal{G} , with game simulation defined according to the projection: $s \triangleleft (s', q)$ iff $s = s'$. This game simulation is preference-preserving.

Conversely, \mathcal{G} game-simulates $\mathcal{G} \times \mathcal{A}$, with game simulation defined by $(s, q) \triangleleft' s'$ iff $s = s'$, which is also preference-preserving.

Proof. We begin with proving that both relations are preference-preserving. First notice that if $((s_n, q_n))_{n \geq 0}$ is a play in $\mathcal{G} \times \mathcal{A}$, then its proj -projection $(s_n)_{n \geq 0}$ is a play in \mathcal{G} . Conversely, if $\rho = (s_n)_{n \geq 0}$ is a play in \mathcal{G} , then there is a unique path $(q_n)_{n \geq 0}$ from initial state q_0 in \mathcal{A} which reads it, and $((s_n, q_n))_{n \geq 0}$ is then a path in $\mathcal{G} \times \mathcal{A}$ that we write $\text{proj}^{-1}(\rho) = ((s_n, q_n))_{n \geq 0}$. That way, proj defines a one-to-one correspondence between plays in \mathcal{G} and plays in $\mathcal{G} \times \mathcal{A}$ where the second component starts in q_0 . For a player $B \neq A$, the objective is defined so that $\text{proj}(\rho)$ has the same payoff as ρ . Consider now player A , she is winning in \mathcal{G} for $\rho = (s_n)_{n \geq 0}$ iff $(s_n)_{n \geq 0} \in \mathcal{L}(\mathcal{A})$ iff the unique path $(q_n)_{n \geq 0}$ from initial state q_0 that reads $(s_n)_{n \geq 0}$ satisfies the Rabin condition $(Q_i, R_i)_{i \in [1, n]}$ in \mathcal{A} iff $\text{proj}^{-1}(\rho)$ satisfies the internal Rabin condition $(Q'_i, R'_i)_{i \in [1, n]}$ in $\mathcal{G} \times \mathcal{A}$. This proves that \triangleleft is winning-preserving.

It remains to show that both relations are game simulations. Assume $s \triangleleft (s, q)$ and pick a move m_{Agt} in \mathcal{G} . It is also a move in $\mathcal{G} \times \mathcal{A}$, and $\text{Tab}'((s, q), m_{\text{Agt}}) = (\text{Tab}(s, m_{\text{Agt}}), \delta(q, s))$. By definition of \triangleleft it then holds that $\text{Tab}(s, m_{\text{Agt}}) \triangleleft \text{Tab}'((s, q), m_{\text{Agt}})$, which proves condition (1) of the definition of a game simulation. It remains to show condition (2). Pick a state $(s', q') \in \text{States}'$. We distinguish two cases

- If $\delta(q, s) \neq q'$ then $\text{Susp}(((s, q), (s', q')), m_{\text{Agt}}) = \emptyset$, and condition (2) trivially holds.
- Otherwise $\delta(q, s) = q'$. In that case, for any move m'_{Agt} , we have that $\text{Tab}(s, m'_{\text{Agt}}) = s'$ if, and only if, $\text{Tab}'((s, q), m'_{\text{Agt}}) = (s', q')$. It follows that $\text{Susp}(((s, q), (s', q')), m_{\text{Agt}}) = \text{Susp}((s, s'), m_{\text{Agt}})$, which implies condition (2).

This proves that $\mathcal{G} \times \mathcal{A}$ game-simulates \mathcal{G} .

We now assume $(s, q) \triangleleft' s$ and pick a move m_{Agt} in $\mathcal{G} \times \mathcal{A}$. It is also a move in \mathcal{G} , and as previously, condition (1) obviously holds. Pick now $s' \in \text{States}$. We define $q' = \delta(q, s)$, and we have $(s', q') \triangleleft s'$ by definition of \triangleleft' . As before, we get condition (2) because $\text{Susp}(((s, q), (s', q')), m_{\text{Agt}}) = \text{Susp}((s, s'), m_{\text{Agt}})$. \square

We will solve the case where each player's objective is given by a deterministic Rabin automaton by applying the above result inductively. We will obtain a game where each player has an internal Rabin winning condition. Applying Proposition 5.23 each time, we get the following result:

Proposition 5.25. *Let $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ be a finite concurrent game, where for each player A , the preference relation \preceq_A is single-objective given by a deterministic Rabin automaton \mathcal{A} . Write $\text{Agt} = \{A_1, \dots, A_n\}$. There is a Nash equilibrium σ_{Agt} in \mathcal{G} from some state s with outcome ρ iff there is a Nash equilibrium σ'_{Agt} in $\mathcal{G}' = (((\mathcal{G} \times \mathcal{A}_1) \times \mathcal{A}_2) \cdots \times \mathcal{A}_n)$ from $(s, q_{01}, \dots, q_{0n})$ with outcome ρ' , where q_{0i} is the initial state of \mathcal{A}_i and ρ is the projection of ρ' on \mathcal{G} .*

5.7.3. Algorithm. Assume that the objective of player A_i is given by a deterministic Rabin automaton \mathcal{A}_i . The algorithm for solving the constrained NE existence problem starts by computing the product of the game with the automata: $\mathcal{G}' = (((\mathcal{G} \times \mathcal{A}_1) \times \mathcal{A}_2) \cdots \times \mathcal{A}_n)$. The resulting game has size $|\mathcal{G}| \times \prod_{j \in [1, n]} |\mathcal{A}_j|$, which is exponential in the number of players. For each player A_j ($1 \leq j \leq n$), the number of Rabin pairs in the product game is that of the original specification \mathcal{A}_j , say k_j . We then apply the deterministic algorithm that we have designed for Rabin objectives (see Subsection 5.6.3 page 39), which yields an exponential-time algorithm in our framework.

5.7.4. Hardness. We prove PSPACE-hardness in the restricted case of deterministic Büchi automata, by a reduction from (the complement of) the problem of the emptiness of the intersection of several language given by deterministic finite automata. This problem is known to be PSPACE-complete [29, Lemma 3.2.3].

We fix finite automata $\mathcal{A}_1, \dots, \mathcal{A}_n$ over alphabet Σ . Let $\Sigma' = \Sigma \cup \{\text{init}, \text{final}\}$, where init and final are two special symbols not in Σ . For every $j \in [1, n]$, we construct a Büchi automaton \mathcal{A}'_j from \mathcal{A}_j as follows. We add a state F with a self-loop labelled by final and an initial state I with a transition labelled by init to the original initial state. We add transitions labelled by final from every terminal state to F . We set the Büchi condition to $\{F\}$. If \mathcal{L}_j is the language recognised by \mathcal{A}_j , then the language recognised by the Büchi automaton \mathcal{A}'_j is $\mathcal{L}'_j = \text{init} \cdot \mathcal{L}_j \cdot \text{final}^*$. The intersection of the languages recognised by the automata \mathcal{A}_j is empty if, and only if, the intersection of the languages recognised by the automata \mathcal{A}'_j is empty.

We construct the game \mathcal{G} , with $\text{States} = \Sigma'$. For each $j \in [1, n]$, there is a player A_j whose objective is given by \mathcal{A}'_j and one special player A_0 whose objective is States^ω (she is always winning). Player A_0 controls all the states and there are transitions from any state to the states of $\Sigma \cup \{\text{final}\}$. Formally $\text{Act} = \Sigma \cup \{\text{final}\} \cup \perp$, for all state $s \in \text{States}$, $\text{Mov}(s, A_0) = \text{Act}$, and if $j \neq 0$ then $\text{Mov}(s, A_j) = \{\perp\}$ and for all $\alpha \in \Sigma \cup \{\text{final}\}$, $\text{Tab}(s, (\alpha, \perp, \dots, \perp)) = \alpha$.

Lemma 5.26. *There is a Nash equilibrium in game \mathcal{G} from *init* where every player wins if, and only if, the intersection of the languages recognised by the automata \mathcal{A}'_j is not empty.*

Proof. If there is such a Nash equilibrium, let ρ be its outcome. The path ρ forms a word of Σ' , it is accepted by every automata \mathcal{A}'_j since every player wins. Hence the intersection of the languages \mathcal{L}_j is not empty.

Conversely, if a word $w = \textit{init} \cdot w_1 \cdot w_2 \cdots$ is accepted by all the automata, player A_0 can play in a way such that everybody is winning: if at each step j she plays w_j , then the outcome is w which is accepted by all the automata. It is a Nash equilibrium since A_0 controls everything and cannot improve her payoff. \square

Since PSPACE is stable by complementation, this proves that the constrained NE existence problem is PSPACE-hard for objectives described by Büchi automata.

In order to prove hardness for the NE existence problem we use results from Section 3.5. Winning conditions in $E(E(\dots(E(\mathcal{G}, A_n, A_0, \rho_n), \dots, A_2, A_0, \rho_2), A_1, A_0, \rho_1))$, where ρ_j is a winning play for A_i , can be defined by slightly modifying automata $\mathcal{A}'_1, \dots, \mathcal{A}'_n$ to take into account the new states. By Proposition 3.5, there exists a Nash equilibrium in this game if, and only, if there is one in \mathcal{G} where all the players win. Hence PSPACE-hardness also holds for the NE existence problem.

6. ORDERED BÜCHI OBJECTIVES

In this Section we assume that preference relations of the players are given by ordered Büchi objectives (as defined in Section 2.5), and we prove the results listed in Table 2 (page 3). We first consider the general case of preorders given as Boolean circuits, and then exhibit several simpler cases.

For the rest of this section, we fix a game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$, and assume that \preceq_A is given by an ordered Büchi objective $\omega_A = \langle (\Omega_i^A)_{1 \leq i \leq n_A}, (\preceq_A)_{A \in \text{Agt}} \rangle$.

6.1. General case: preorders are given as circuits.

Theorem 6.1. *For finite games with ordered Büchi objectives where preorders are given as Boolean circuits, the value problem, the NE existence problem and the constrained NE existence problem are PSPACE-complete.*

Proof. We explain the algorithm for the constrained NE existence problem. We assume that for each player A , the preorder \preceq_A is given by a Boolean circuit C_A . The algorithm proceeds by trying all the possible payoffs for the players.

Fix such a payoff $(v^A)_{A \in \text{Agt}}$, with $v^A \in \{0, 1\}^{n_A}$ for every player A . We build a circuit D_A which represents a single objective for player A . Inputs to circuit D_A will be states of the game. This circuit is constructed from C_A as follows: We set all input gates $w_1 \cdots w_n$ of circuit C_A to the value given by payoff v^A ; The former input v_i receives the disjunction of all the states in Ω_i ; We negate the output. It is not hard to check that the new circuit D_A is such that for every play ρ , $D_A[\text{Inf}(\rho)]$ evaluates to **true** if, and only if, $\text{payoff}_A(\rho) \not\preceq_A v^A$, i.e. if ρ is an improvement for player A .

Circuit D_A is now viewed as a single objective for player A , we write \mathcal{G}' for the new game. We look for Nash equilibria in this new game, with payoff 0 for each player. Indeed, a Nash equilibrium σ_{Agt} in \mathcal{G} with payoff $(v^A)_{A \in \text{Agt}}$ is a Nash equilibrium in game \mathcal{G}' with

payoff $(0, \dots, 0)$. Conversely a Nash equilibrium σ_{Agt} in game \mathcal{G}' with payoff $(0, \dots, 0)$ is a Nash equilibrium in \mathcal{G} as soon as the payoff of its outcome (in \mathcal{G}) is $(v^A)_{A \in \text{Agt}}$.

We use the algorithm described in Section 5.5. for computing Nash equilibria with single objectives given as Boolean circuits, and we slightly modify it to take into account the constraint that it has payoff v^A for each player A . This can be done in polynomial space, thanks to Proposition 3.1: it is sufficient to look for plays of the form $\pi \cdot \tau^\omega$ with $|\pi| \leq |\text{States}|^2$ and $|\tau| \leq |\text{States}|^2$.

PSPACE-hardness was proven for single objectives given as a Boolean circuit (the circuit evaluates by setting to `true` all states that are visited infinitely often, and to `false` all other states) in Section 5.5. This kind of objective can therefore be seen as an ordered Büchi objective with a preorder given as a Boolean circuit. \square

6.2. When the ordered objective can be (co-)reduced to a single Büchi objective.

For some ordered objectives, the preference relation can (efficiently) be reduced to a single objective. For instance, a disjunction of several Büchi objectives can obviously be reduced to a single Büchi objective, by considering the union of the target sets. Formally, we say that an ordered Büchi objective $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ is *reducible* to a single Büchi objective if, given any payoff vector v , we can construct in polynomial time a target set $\widehat{T}(v)$ such that for all paths ρ , $v \lesssim \text{payoff}_\omega(\rho)$ if, and only if, $\text{Inf}(\rho) \cap \widehat{T}(v) \neq \emptyset$. It means that *securing* payoff v corresponds to ensuring infinitely many visits to the new target set. Similarly, we say that ω is *co-reducible* to a single Büchi objective if for any vector v we can construct in polynomial time a target set $\widehat{T}(v)$ such that $\text{payoff}_\omega(\rho) \not\lesssim v$ if, and only if, $\text{Inf}(\rho) \cap \widehat{T}(v) \neq \emptyset$. It means that *improving* on payoff v corresponds to ensuring infinitely many visits to the new target

We prove the following proposition, which exploits (co-)reducibility for efficiently solving the various problems.

Proposition 6.2. – *For finite games with ordered Büchi objectives which are reducible to single Büchi objectives, and in which the preorders are non-trivial⁶ and monotonic, the value problem is P-complete.*

– *For finite games with ordered Büchi objectives which are co-reducible to single Büchi objectives, and in which the preorders are non-trivial and monotonic the NE existence problem and the constrained NE existence problem are P-complete.*

Note that the hardness results follow from the hardness of the same problems for single Büchi objectives (see Section 5.3). We now prove the two upper bounds.

6.2.1. *Reducibility to single Büchi objectives and the value problem.* We transform the ordered Büchi objectives of the considered player into a single Büchi objective, and use a polynomial-time algorithm [23, Chapter 2] to solve the resulting zero-sum (turn-based) Büchi game.

⁶That is, there is more than one class in the preorder.

6.2.2. Co-reducibility to single Büchi objectives and the (constrained) NE existence problem.

We assume that the ordered objectives $(\omega_A)_{A \in \text{Agts}}$ are all co-reducible to single Büchi objectives. We show that we can use the algorithm presented in Section 5.3.2 to solve the constrained NE existence problem in polynomial time.

We first notice that the preference relations \lesssim_A satisfy the hypotheses (\star) (see page 30): $(\star)_a$ and $(\star)_b$ are obvious, and $(\star)_c$ is by co-reducibility of the ordered objectives. It means that we can apply the results of Lemmas 5.10 and 5.11 to the current framework. To be able to conclude and apply Lemma 5.12, we need to show that for every payoff v , we can compute in polynomial time the set $W(\mathcal{G}, v)$ in the suspect game $\mathcal{H}(\mathcal{G}, v)$.

Lemma 6.3. *Fix a threshold v . The set $W(\mathcal{G}, v)$ can be computed in polynomial time.*

Proof. As the ordered objectives are co-reducible to single Büchi objectives, we can construct in polynomial time target sets $\widehat{T}^A(v)$ for each player A . The objective of **Eve** in the suspect game $\mathcal{H}(\mathcal{G}, K)$ is then equivalent to a co-Büchi objective with target set $\{\widehat{T}^A(v, P) \mid A \in P\}$. The winning region $W(\mathcal{G}, v)$ can then be determined using a polynomial time algorithm of [23, Sect. 2.5.3]. \square

6.2.3. Applications.

We will give preorders to which the above applies, allowing to infer several P-completeness results in Table 2 (those written with reference “Section 6.2”).

We first show that reducibility and co-reducibility coincide when the preorder is total.

Lemma 6.4. *Let $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ be an ordered Büchi objective, and assume that \lesssim is total. Then, ω is reducible to a single Büchi objective if, and only if, ω is co-reducible to a single Büchi objective.*

Proof. Let $u \in \{0, 1\}^n$ be a vector. If u is a maximal element, the new target set is empty, and thus satisfies the property for co-reducibility. Otherwise we pick a vector v among the smallest elements that is strictly larger than u . Since the preorder is reducible to a single Büchi objective, there is a target set \widehat{T} that is reached infinitely often whenever the payoff is greater than v . Since the preorder is total and by choice of v , we have $w \not\lesssim u \Leftrightarrow v \lesssim w$. Thus the target set \widehat{T} is visited infinitely often when u is not larger than the payoff. Hence ω is co-reducible to a single Büchi objective.

The proof of the other direction is similar. \square

Lemma 6.5. *Ordered Büchi objectives with disjunction or maximise preorders are reducible to single Büchi objectives. Ordered Büchi objectives with disjunction, maximise or subset preorders are co-reducible to single Büchi objectives.*

Proof. Let $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ be an ordered Büchi objective. Assume T_i is the target set for Ω_i .

Assume \lesssim is the disjunction preorder. If the payoff v is different from $\mathbf{0}$ then we define $\widehat{T}(v)$ as the union of all the target sets: $\widehat{T}(v) = \bigcup_{i=1}^n T_i$. Then, for every run ρ ,

$$\begin{aligned} v \lesssim \text{payoff}_\omega(\rho) &\Leftrightarrow \text{there is some } i \text{ for which } \text{Inf}(\rho) \cap T_i \neq \emptyset \\ &\Leftrightarrow \text{Inf}(\rho) \cap \widehat{T}(v) \neq \emptyset \end{aligned}$$

If the payoff v is $\mathbf{0}$ then we get the expected result with $\widehat{T}(v) = \text{States}$. Disjunction being a total preorder, it is also co-reducible (from Lemma 6.4).

We assume now that \lesssim is the maximise preorder. Given a payoff v , consider the index $i_0 = \max\{i \mid v_i = 1\}$. We then define $\widehat{T}(v)$ as the union of the target sets that are above i_0 : $\widehat{T}(v) = \bigcup_{i \geq i_0} T_i$. The following four statements are then equivalent, if ρ is a run:

$$\begin{aligned} v \lesssim \text{payoff}_\omega(\rho) &\Leftrightarrow v \lesssim \mathbf{1}_{\{i \mid \text{Inf}(\rho) \cap T_i \neq \emptyset\}} \\ &\Leftrightarrow i_0 \leq \max\{i \mid \text{Inf}(\rho) \cap T_i \neq \emptyset\} \\ &\Leftrightarrow \exists i \geq i_0. \text{Inf}(\rho) \cap T_i \neq \emptyset \end{aligned}$$

Hence ω is reducible, and also co-reducible as it is total, to a single Büchi objective.

Finally, we assume that \lesssim is the subset preorder, and we show that ω is then co-reducible to a single Büchi objective. Given a payoff v , the new target is the union of the target sets that are not reached infinitely often for that payoff: $\widehat{T}(v) = \bigcup_{\{i \mid v_i = 0\}} T_i$. Then the following statements are equivalent, if ρ is a run:

$$\begin{aligned} \text{payoff}_\omega(\rho) \not\lesssim u &\Leftrightarrow \mathbf{1}_{\{i \mid \text{Inf}(\rho) \cap T_i \neq \emptyset\}} \not\lesssim u \\ &\Leftrightarrow \exists i. \text{Inf}(\rho) \cap T_i \neq \emptyset \text{ and } u_i = 0 \\ &\Leftrightarrow \text{Inf}(\rho) \cap \widehat{T}(v) \neq \emptyset \end{aligned}$$

□

As a corollary, we get the following result:

Corollary 6.6. *For finite games with ordered Büchi objectives, with either the disjunction or the maximise preorder, the value problem is P-complete. For finite games with ordered Büchi objectives, with either the disjunction, the maximise or the subset preorder, the NE existence problem and the constrained NE existence problem are P-complete.*

Remark 6.7. Note that we cannot infer P-completeness of the value problem for the subset preorder since the subset preorder is not total, and ordered objectives with subset preorder are not reducible to single Büchi objectives. Such an ordered objective is actually reducible to a generalised Büchi objective (several Büchi objectives should be satisfied).

6.3. When the ordered objective can be reduced to a deterministic Büchi automaton objective. For some ordered objectives, the preference relation can (efficiently) be reduced to the acceptance by a deterministic Büchi automaton. Formally, we say that an ordered objective $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ is *reducible* to a deterministic Büchi automaton whenever, given any payoff vector u , we can construct in polynomial time a deterministic Büchi automaton over States which accepts exactly all plays ρ with $u \lesssim \text{payoff}_\omega(\rho)$. For such preorders, we will see that the value problem can be solved efficiently by constructing the product of the deterministic Büchi automaton and the arena of the game. This construction does however not help for solving the (constrained) NE existence problems since the number of players is a parameter of the problem, and the size of the resulting game will then be exponential.

Proposition 6.8. *For finite games with ordered Büchi objectives which are reducible to deterministic Büchi automata, the value problem is P-complete.*

Proof. Given the payoff v^A for player A , the algorithm proceeds by constructing the automaton that recognises the plays with payoff higher than v^A . By performing the product with the game as described in Section 5.7.2, we obtain a new game, in which there is a winning

strategy if, and only if, there is a strategy in the original game to ensure payoff v^A . In this new game, player A has a single Büchi objective, so that the NE existence of a winning strategy can be decided in polynomial time.

Hardness follows from that of games with single Büchi objectives. \square

Applications. We now give preorders to which the above result applies, that is, which are reducible to deterministic Büchi automata objectives.

Lemma 6.9. *An ordered objective where the preorder is either the conjunction, the subset or the lexicographic preorder is reducible to a deterministic Büchi automaton objective.*

Proof. We first focus on the **conjunction preorder**. Let $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ be an ordered Büchi objective, where \lesssim is the conjunction. For every $1 \leq i \leq n$, let T_i be the target set defining the Büchi condition Ω_i . There are only two possible payoffs: either all objectives are satisfied, or one objective is not satisfied. For the second payoff case, any play has a larger payoff: hence the trivial automaton (which accepts all plays) witnesses the property. For the first payoff case, we construct a deterministic Büchi automaton \mathcal{B} as follows. There is one state for each target set, plus one accepting state: $Q = \{q_0, q_1, \dots, q_n\}$; the initial state is q_0 , and the unique repeated state is q_n . For all $1 \leq i \leq n$, the transitions are $q_{i-1} \xrightarrow{s} q_i$ when $s \in T_i$ and $q_{i-1} \xrightarrow{s} q_{i-1}$ otherwise. There are also transitions $q_n \xrightarrow{s} q_0$ for every $s \in \text{States}$. Automaton \mathcal{B} describes the plays that goes through each set T_i infinitely often, hence witnesses the property. It can furthermore be computed in polynomial time. The construction is illustrated in Figure 21.

We now turn to the **subset preorder**. Let $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ be an ordered Büchi objective, where \lesssim is the subset preorder. For every $1 \leq i \leq n$, let T_i be the target set defining the Büchi condition Ω_i . Fix a payoff u . A play ρ is such that $u \lesssim \text{payoff}_\omega(\rho)$ if, and only if, ρ visits infinitely often all sets T_i with $u_i = 1$. This is then equivalent to the conjunction of all Ω_i 's with $u_i = 1$. We therefore apply the previous construction for the conjunction and get the expected result.

We finish this proof with the **lexicographic preorder**. Let $\omega = \langle (\Omega_i)_{1 \leq i \leq n}, \lesssim \rangle$ be an ordered Büchi objective, where \lesssim is the lexicographic preorder. For every $1 \leq i \leq n$, let T_i be the target set defining the Büchi condition Ω_i . Let $u \in \{0, 1\}^n$ be a payoff vector. We construct the following deterministic Büchi automaton which recognises the runs whose payoff is greater than or equal to u .

In this automaton there is a state q_i for each i such that $u_i = 1$, and a state q_0 that is both initial and repeated: $Q = \{q_0\} \cup \{q_i \mid u_i = 1\}$. We write $I = \{0\} \cup \{i \mid u_i = 1\}$. For every $i \in I$, we write $\text{succ}(i) = \min(I \setminus \{j \mid j \leq i\})$, with the convention that $\min \emptyset = 0$. The transition relation is defined as follows:

- for every $s \in \text{States}$, there is a transition $q_0 \xrightarrow{s} q_{\text{succ}(0)}$;
- for every $i \in I \setminus \{0\}$, we have the following transitions:
 - $q_i \xrightarrow{T_i} q_{\text{succ}(i)}$;
 - $q_i \xrightarrow{T_k \setminus T_i} q_0$ with $k < i$ and $u_k = 0$;
 - $q_i \xrightarrow{s} q_i$ for every $s \in \text{States} \setminus (T_i \cup \bigcup_{k < i, u_k = 0} T_k)$.

An example of the construction is given in Figure 22.

We now prove correctness of this construction. Consider a path that goes from q_0 to q_0 : if the automaton is currently in state q_i , then since the last occurrence of q_0 , at least one state for each target set T_j with $j < i$ and $u_j = 1$ has been visited. When q_0 is reached again, either it is because we have seen all the T_j with $u_j = 1$, or it is because the run visited some target T_i with $u_i = 0$ and all the T_j such that $u_j = 1$ and $j < i$; in both cases, the set of targets that have been visited between two visits to q_0 describes a payoff greater than u . Assume the play π is accepted by the automaton; then there is a sequence of q_i as above that is taken infinitely often, therefore $\text{payoff}_\omega(\pi)$ is greater than or equal to u for the lexicographic order.

Conversely assume $v = \text{payoff}_\omega(\pi)$ is greater than or equal to u , that we already read a prefix $\pi_{\leq k}$ for some k , and that the current state is q_0 . Reading the first symbol in π after position k , the run goes to the state q_i where i is the least integer such that $u_i = 1$. Either the path visits T_i at some point, or it visits a state in a target T_j , with j smaller than i and $v_j = 0$, in which case the automaton goes back to q_0 . Therefore from q_0 we can again come back to q_0 while reading the following of π , and the automaton accepts. \square




FIGURE 21. The automaton for the conjunction preorder, $n = 3$




FIGURE 22. The automaton for the lexicographic order, $n = 7$ and $u = (0, 1, 0, 0, 1, 1, 0)$

We conclude with the following corollary:

Corollary 6.10. *For finite games with ordered Büchi objectives with either of the conjunction, the lexicographic or the subset preorders, the value problem is P-complete.*

6.4. Preference relations with monotonic preorders. We will see in this part that monotonic preorders lead to more efficient algorithms. More precisely we prove the following result:

Proposition 6.11. – *For finite games with ordered Büchi objectives where the preorders are given by monotonic Boolean circuits, the value problem is in coNP, and the NE existence problem and the constrained NE existence problem are in NP.*

- *Completeness holds in both cases for finite games with ordered Büchi objectives where the preorders are given by monotonic Boolean circuits or with the counting preorder.*
- *NP-completeness also holds for the constrained NE existence problem for finite games with ordered Büchi objectives where the preorders admit an element v such that for every v' , it holds $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$.*⁷

⁷To be fully formal, a preorder \lesssim is in fact a family $(\lesssim_n)_{n \in \mathbb{N}}$ (where \lesssim_n compares two vectors of size n), and this condition should be stated as “for all n , there is an element $v_n \in \{0, 1\}^n$ such that for all $v' \in \{0, 1\}^n$, it holds $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v_n$ ”.

We first show that monotonicity of the preorders imply some memorylessness property in the suspect game. We then give algorithms witnessing the claimed upper bounds, and show the various lower bounds.

6.4.1. *When monotonicity implies memorylessness.* We say that a strategy σ is *memoryless* (resp. memoryless from state s_0) if there exists a function $f: \text{States} \rightarrow \text{Act}$ such that $\sigma(h \cdot s) = f(s)$ for every $h \in \text{Hist}$ (resp. for every $h \in \text{Hist}(s_0)$). A strategy profile is said memoryless whenever all strategies of single players are memoryless. We show that when the preorders used in the ordered Büchi objectives are monotonic, the three problems are also easier than in the general case. This is because we can find memoryless trigger profiles (recall Definition 4.1).

We first show this lemma, that will then be applied to the suspect game.

Lemma 6.12. *Let \mathcal{H} be a turn-based two-player game. Call *Eve* one player, and let σ_{\exists} be a strategy for *Eve*, and s_0 be a state of \mathcal{H} . There is a memoryless strategy σ'_{\exists} such that for every $\rho' \in \text{Out}_{\mathcal{H}}(s_0, \sigma'_{\exists})$, there exists $\rho \in \text{Out}_{\mathcal{H}}(s_0, \sigma_{\exists})$ such that $\text{Inf}(\rho') \subseteq \text{Inf}(\rho)$.*

Proof. This proof is by induction on the size of the set

$$S(\sigma_1) = \{(s, m) \mid \exists h \in \text{Hist}(\sigma_1). \sigma_1(h) = m \text{ and } \text{last}(h) = s\}.$$

If its size is the same as that of $\{s \mid \exists h \in \text{Hist}(\sigma_1). \text{last}(h) = s\}$ then the strategy is memoryless. Otherwise, let s be a state at which σ_1 takes several different actions (i.e., $|(\{s\} \times \text{Act}) \cap S(\sigma_1)| > 1$).

We will define a new strategy σ'_1 that takes fewer different actions in s and such that for every outcome of σ'_1 , there is an outcome of σ_1 that visits (at least) the same states infinitely often.

If σ is a strategy and h is a history, we let $\sigma \circ h: h' \mapsto \sigma(h \cdot h')$ for any history h' . Then for every m such that $(s, m) \in S(\sigma_1)$ we let $H_m = \{h \in \text{Hist}(\sigma_1) \mid \text{last}(h) = s \text{ and } \sigma_1(h) = m\}$, and for every h , $h^{-1} \cdot H_m = \{h' \mid h \cdot h' \in H_m\}$. We pick m such that H_m is not empty.

– Assume that there is $h_0 \in \text{Hist}(\sigma_1)$ with $\text{last}(h_0) = s$, such that $h_0^{-1} \cdot H_m$ is empty.

We define a new strategy σ'_1 as follows. If h is an history which does not visit s , then $\sigma'_1(h) = \sigma_1(h)$. If h is an history which visits s , then decompose h as $h' \cdot h''$ where $\text{last}(h') = s$ is the first visit to s and define $\sigma'_1(h) = \sigma_1(h_0 \cdot h'')$. Then, strategy σ'_1 does not use m at state s , and therefore at least one action has been “removed” from the strategy. More precisely, $|(\{s\} \times \text{Act}) \cap S(\sigma'_1)| \leq |(\{s\} \times \text{Act}) \cap S(\sigma_1)| - 1$. Furthermore the conditions on infinite states which are visited infinitely often by outcomes of σ'_1 is also satisfied.

– Otherwise for any $h \in \text{Hist}(\sigma_1)$ with $\text{last}(h) = s$, $h^{-1} \cdot H_m$ is not empty. We will construct a strategy σ'_1 which plays m at s . Let h be an history, we first define the extension $e(h)$ inductively in that way:

- $e(\varepsilon) = \varepsilon$, where ε is the empty history;
- $e(h \cdot s) = e(h) \cdot h'$ where $h' \in (e(h))^{-1} \cdot H_m$;
- $e(h \cdot s') = e(h) \cdot s'$ if $s' \neq s$.

We extend the definition of e to infinite outcomes in the natural way: $e(\rho)_i = e(\rho_{\leq i})_i$.

We then define the strategy $\sigma'_1: h \mapsto \sigma_1(e(h))$. We show that if ρ is an outcome of σ'_1 , then $e(\rho)$ is an outcome of σ_1 . Indeed assume h is a finite outcome of σ'_1 , that $e(h)$ is an outcome of σ_1 and $\text{last}(h) = \text{last}(e(h))$. If $h \cdot s$ is an outcome of σ'_1 , by construction of e , $e(h \cdot s) = e(h) \cdot h'$, such that $\text{last}(h') = s$, and h' is an outcome of $\sigma_1 \circ e(h)$ and as

$e(h)$ is an outcome of σ_1 by hypothesis, that means that $e(h \cdot s)$ is an outcome of σ_1 . If $h \cdot s'$ with $s' \neq s$ is an outcome of σ'_1 , $e(h \cdot s') = e(h) \cdot s'$, $s' \in \text{Tab}(\text{last}(h), \sigma'_1(h))$, and $\sigma'_1(h) = \sigma_1(e(h))$. Using the hypothesis $\text{last}(h) = \text{last}(e(h))$, and $e(h)$ is an outcome of σ_1 , therefore $e(h \cdot s')$ is an outcome of σ_1 . This shows that if ρ is an outcome of σ'_1 then $e(\rho)$ is an outcome of σ_1 . The property on states visited infinitely often follows. Several moves have been removed from the strategy at s (since the strategy is now memoryless at s , playing m).

In all cases we have $S(\sigma'_1)$ strictly included in $S(\sigma_1)$, and an inductive reasoning entails the result. \square

Lemma 6.13. *If for every player A , \lesssim_A is monotonic, and if there is a trigger profile for some play π from s , then there is a memoryless winning strategy for Eve in $\mathcal{H}(\mathcal{G}, \pi)$ from state (s, Agt) .*

Proof. Assume there is a trigger profile for π . We have seen in Lemma 4.4 that there is then a winning strategy σ_\exists in game $\mathcal{H}(\mathcal{G}, \pi)$ for Eve. Consider the memoryless strategy σ'_\exists constructed as in Lemma 6.12. Let ρ' be an outcome of σ'_\exists , there is an outcome ρ of σ_\exists such that $\text{Inf}(\rho') \subseteq \text{Inf}(\rho)$. As σ_\exists is winning in $\mathcal{H}(\mathcal{G}, \pi)$, for every $A \in \lambda(\rho)$, $\text{proj}_1(\rho) \lesssim_A \pi$. We assume the Büchi conditions are given by the target sets $(T_i^A)_{A,i}$. For each player A , $\{i \mid \text{Inf}(\text{proj}_1(\rho')) \cap T_i^A\} \subseteq \{i \mid \text{Inf}(\text{proj}_1(\rho)) \cap T_i^A\}$. As the preorder is monotonic the payoff of $\text{proj}_1(\rho')$ is smaller than that of $\text{proj}_1(\rho)$: $\text{proj}_1(\rho') \lesssim_A \text{proj}_1(\rho)$. So the play is winning for any player A and σ'_\exists is a memoryless winning strategy in game $\mathcal{H}(\mathcal{G}, \pi)$ for Eve. \square

Lemma 6.14. *If for every player A , \lesssim_A is given by monotonic Boolean circuits, then given a path π , we can decide in polynomial time if a memoryless strategy for Eve in $\mathcal{H}(\mathcal{G}, \pi)$ is winning.*

Proof. Let σ_\exists be a memoryless strategy in $\mathcal{H}(\mathcal{G}, \pi)$ for Eve. By keeping only the edges that are taken by σ_\exists , we define a subgraph of the game. We can compute in polynomial time the strongly connected components of this graph. If one component is reachable and does not satisfy the objective of Eve, then the strategy is not winning. Conversely if all the reachable strongly connected components satisfy the winning condition of Eve, since the preorder is monotonic, σ_\exists is a winning strategy. Notice that since the preorder is given as a Boolean circuit, we can check in polynomial time whether a strongly connected component is winning or not. Globally the algorithm is therefore polynomial-time. \square

We now turn to the proof of the claimed upper bounds.

6.4.2. *Proofs for the upper bounds.* We show that the value problem is in coNP for finite games with ordered Büchi objectives, when preorders are given by monotonic Boolean circuits.

As already mentioned at the beginning of Section 5, for the value problem, we can make the concurrent game turn-based: since player A must win against any strategy of the coalition $P = \text{Agt} \setminus \{A\}$, she must also win in the case where the opponents' strategies can adapt to what A plays. In other terms, we can make A play first, and then the coalition. This turn-based game is determined, so that there is a strategy σ whose outcomes are always better (for A) than v^A if, and only if, for any strategy σ' of coalition P , there is an outcome with payoff (for A) better than v^A . If there is a counterexample to this fact, then thanks to Lemma 6.12 there is one with a memoryless strategy σ' . The coNP algorithm proceeds by

checking that all the memoryless strategies of coalition P have an outcome better than v^A , which is achievable in polynomial time, with a method similar to Lemma 6.14.

We show now that the constrained NE existence problem is in NP for finite games with ordered Büchi objectives, when preorders are given by monotonic Boolean circuits.

The algorithm for the constrained NE existence problem proceeds by guessing:

- the payoff for each player,
- a play of the form $\pi \cdot \tau^\omega$, where $|\pi| \leq |\text{States}|^2$ and $|\tau| \leq |\text{States}|^2$,
- an under-approximation W of the set of winning states in $\mathcal{H}(\mathcal{G}, \pi \cdot \tau^\omega)$
- a memoryless strategy profile σ_{Agt} in $\mathcal{H}(\mathcal{G}, \pi \cdot \tau^\omega)$.

We check that σ_{Agt} is a witness for the fact that the states in W are winning; thanks to Lemma 6.14, this can be done in polynomial time. We also verify that the play $\pi \cdot \tau^\omega$ has the expected payoff, that the payoff satisfies the constraints, and that it never gets out of W . If these conditions are fulfilled, then the play $\pi \cdot \tau^\omega$ meets the conditions of Theorem 4.5, and there is a Nash equilibrium with outcome $\pi \cdot \tau^\omega$. Lemma 6.13 and Proposition 3.1 ensure that if there is a Nash equilibrium, we can find it this way.

6.4.3. Proofs for the hardness results. We first prove the hardness results for the counting preorder.

Lemma 6.15. *For finite games with ordered Büchi objectives that use the counting preorder, the value problem is coNP-hard.*

Proof. We reduce (the complement of) 3SAT into the value problem for two-player turn-based games with Büchi objectives with the counting preorder. Consider an instance

$$\phi = C_1 \wedge \cdots \wedge C_m$$

with $C_j = \ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3}$, over a set of variables $\{x_1, \dots, x_n\}$. With ϕ , we associate a two-player turn-based game \mathcal{G} . Its set of states is made of

- a set containing the unique initial state $V_0 = \{s_0\}$,
- a set of two states $V_k = \{x_k, \neg x_k\}$ for each $1 \leq k \leq n$,
- and a set of three states $V_{n+j} = \{t_{j,1}, t_{j,2}, t_{j,3}\}$ for each $1 \leq j \leq m$.

Then, for each $0 \leq l \leq n + m$, there is a transition between any state of V_l and any state of V_{l+1} (assuming $V_{n+m+1} = V_0$).

The game involves two players: player B owns all the states, but has no objectives (she always loses). Player A has a set of Büchi objectives defined by $T_{2,k}^A = \{x_k\} \cup \{t_{j,p} \mid \ell_{j,p} = x_k\}$, $T_{2,k+1}^A = \{\neg x_k\} \cup \{t_{j,p} \mid \ell_{j,p} = \neg x_k\}$, for $1 \leq k \leq n$. Notice that at least n of these objectives will be visited infinitely often along any infinite play. We prove that if the formula is not satisfiable, then at least $n + 1$ objectives will be fulfilled, and conversely.

Assume the formula is satisfiable, and pick a witnessing valuation v . We define a strategy σ_B for B that “follows” valuation v : from states in V_{k-1} , for any $1 \leq k \leq n$, the strategy plays towards x_k if $v(x_k) = \text{true}$ (and to $\neg x_k$ otherwise). Then, from a state in V_{n+l-1} with $1 \leq l \leq m$, it plays towards one of the $t_{j,p}$ that evaluates to true under v (the one with least index p , say). This way, the number of targets of player A that are visited infinitely often is n .

Conversely, pick a play in \mathcal{G} s.t. at most (hence exactly) n objectives of A are fulfilled. In particular, for any $1 \leq k \leq n$, this play never visits one of x_k and $\neg x_k$, so that it defines a valuation v over $\{x_1, \dots, x_n\}$. Moreover, any state of V_{n+l} , with $1 \leq l \leq m$, that is visited

infinitely often must correspond to a literal that is made true by v , as otherwise this would make one more objective that is fulfilled for A . As a consequence, each clause of ϕ evaluates to true under v , and the result follows. \square




FIGURE 23. The game \mathcal{G} associated with formula ϕ of 6.1

Example 6.16. We illustrate the construction of the previous proof in Figure 23 for the formula

$$\varphi = (x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3). \quad (6.1)$$

The targets for player A are $T_1 = \{x_1, t_{1,1}\}$, $T_2 = \{\neg x_1, t_{2,1}\}$, $T_3 = \{x_2, t_{1,2}, t_{2,2}\}$, $T_4 = \{\neg x_2\}$, $T_5 = \{x_3\}$, $T_6 = \{\neg x_3, t_{1,3}, t_{2,3}\}$. Player A cannot ensure visiting infinitely often four target sets, therefore the formula is satisfiable.

Lemma 6.17. *For finite games with ordered Büchi objectives that use the counting preorder, the NE existence problem is NP-hard.*

Proof. Let \mathcal{G} be the game we constructed for Lemma 6.15. We construct the game \mathcal{G}'' from \mathcal{G} as described in Section 3.3. The preference in \mathcal{G}'' can still be described with ordered Büchi objectives and the counting preorder: the only target set of B is $\{s_1\}$ and we add s_1 to n different targets of A , where n is the number of variables as in Lemma 6.15. From Proposition 3.4 there is a Nash equilibrium in \mathcal{G}'' from s_0 if, and only if, A cannot ensure visiting at least $n + 1$ targets infinitely often. Hence the NE existence problem is NP-hard. \square

This proves also NP-hardness for the constrained NE existence problem for ordered Büchi objectives with the counting preorder. Hardness results for preorders given by monotonic Boolean circuits follow from the above since the counting preorder is a special case of preorder given as a monotonic Boolean circuit (and the counting preorder can be expressed as a polynomial-size monotonic Boolean circuit).

We now show hardness in the special case of preorders with (roughly) at most one maximal element below $\mathbf{1}$.

Lemma 6.18. *For finite turn-based games with ordered Büchi objectives with a monotonic preorder for which there is an element v such that for every v' , $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$, the constrained NE existence problem is NP-hard.*

Proof. Let us consider a formula $\phi = C_1 \wedge \dots \wedge C_m$. For each variable x_i , our game has one player B_i and three states s_i , x_i and $\neg x_i$. The objectives of B_i are the sets $\{x_i\}$ and $\{\neg x_i\}$. Transitions go from each s_i to x_i and $\neg x_i$, and from x_i and $\neg x_i$ to s_{i+1} (with $s_{n+1} = s_0$). Finally, an extra player A has full control of the game (i.e., she owns all the states) and has




FIGURE 24. The Büchi game for a formula with 4 variables

n objectives, defined by $T_i^A = \{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$ for $1 \leq i \leq n$. The construction is illustrated in Figure 24.

We show that formula ϕ is satisfiable if, and only if, there is a Nash equilibrium where each player B_i gets payoff β_i satisfying $\beta_i \lesssim v$ (hence $\beta_i \neq (1, 1)$), and player A gets payoff $\mathbf{1}$.

First assume that the formula is satisfiable, and pick a witnessing valuation u . By playing according to u , player A can satisfy all of her objectives (hence she cannot improve her payoff, since the preorder is monotonic). Since she alone controls all the game, the other players cannot improve their payoff, so that this is a Nash equilibrium. Moreover, since A plays memoryless, only one of x_i and $\neg x_i$ is visited for each i , so that the payoff β_i for B_i satisfies $\beta_i \lesssim v$. Conversely, if there is a Nash equilibrium with the desired payoff, then by hypothesis, exactly one of each x_i and $\neg x_i$ is visited infinitely often (so that the payoff for B_i is not $(1, 1)$), which defines a valuation u . Since in this Nash equilibrium, player A satisfies all its objectives, one state of each target is visited, which means that under valuation u , formula ϕ evaluates to true. \square

6.4.4. *Applications.* We now describe examples of preorders which satisfy the conditions on the existence of an element v such that $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$.

Lemma 6.19. *Conjunction, counting and lexicographic preorders have an element v such that $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$.*

Proof. Consider $v = (1, \dots, 1, 0)$, and $v' \neq \mathbf{1}$. For conjunction, there is i such that $v'_i = 0$, so $v' \lesssim v$. For counting, $|\{i \mid v'_i = 1\}| < n$, so $v' \lesssim v$. For the lexicographic preorder, let i be the smallest index such that $v'_i = 0$, and either $v_i = 1$ and $v_j = v'_j$ for all $j < i$, or for all $j \in \{1, \dots, n\}$, $v_j = v'_j$. In both cases $v' \lesssim v$. \square

As a consequence, the result of Lemma 6.18 applies in particular to the conjunction and lexicographic preorders, for which the constrained NE existence problem is thus NP-complete. Hence we get:

Corollary 6.20. *For finite games with ordered Büchi objectives with either of the conjunction or the lexicographic preorders, the constrained NE existence problem is NP-complete.*

7. ORDERED REACHABILITY OBJECTIVES

In this Section we assume that preference relations of the players are given by ordered reachability objectives (as defined in Section 2.5), and we prove the results listed in Table 3 (page 3). We will first consider the general case when preorders are given by Boolean circuits and we will show that the various decision problems are PSPACE-complete. We will even

notice that the hardness result holds for several simpler preorders. We will finally improve this result in a number of cases.

For the rest of this section, we fix a game $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$, and we assume that \preceq_A is given by an ordered reachability objective $\omega_A = \langle (\Omega_i^A)_{1 \leq i \leq n_A}, (\preceq_A)_{A \in \text{Agt}} \rangle$.

7.1. General case: preorders are given as circuits. We prove the following result:

Proposition 7.1. – *For finite games with ordered reachability objectives where preorders are given by Boolean circuits, the value problem, the NE existence problem and the constrained NE existence problem are in PSPACE.*

- *For finite two-player turn-based games with ordered reachability objectives where preorders have $\mathbf{1}$ as a unique maximal element, the value problem is PSPACE-hard.*
- *For finite two-player games with ordered reachability objectives where preorders have $\mathbf{1}$ as a unique maximal element, and have an element v such that for every v' , $v' \neq \mathbf{1} \Leftrightarrow v' \preceq v$, then the NE existence problem and the constrained NE existence problem are PSPACE-hard.*

The upper bound will be proven by reduction to games with ordered Büchi objectives using game-simulation.

7.1.1. Reduction to a game with ordered Büchi objectives. We show how to transform a game \mathcal{G} with preferences given by Boolean circuits over reachability objectives into a new game \mathcal{G}' , with preferences given by Boolean circuits over Büchi objectives. Although the size of \mathcal{G}' will be exponential, circuit order with Büchi objectives define prefix-independent preference relations and thus checking condition 3 of Theorem 4.5 can be made more efficient.

States of \mathcal{G}' store the set of states of \mathcal{G} that have already been visited. The set of states of \mathcal{G}' is $\text{States}' = \text{States} \times 2^{\text{States}}$. The transitions are as follows: $(s, S) \rightarrow (s', S')$ when there is a transition $s \rightarrow s'$ in \mathcal{G} and $S' = S \cup \{s'\}$. We keep the same circuits to define the preference relations, but the reachability objectives are transformed into Büchi objectives: a target set T is transformed into $T' = \{(s, S) \mid S \cap T \neq \emptyset\}$. Although the game has exponential size, the preference relations only depend on the strongly connected components the path ends in, so that we will be able to use a special algorithm, which we describe after this lemma.

We define the relation $s \triangleleft s'$ over states of \mathcal{G} and \mathcal{G}' if, and only if, $s' = (s, S)$ with $S \subseteq \text{States}$, and prove that it is a game simulation (see Definition 5.22).

Lemma 7.2. *The relation \triangleleft (resp. \triangleleft^{-1}) is a game simulation between \mathcal{G} and \mathcal{G}' , and it is preference-preserving from $(s_0, (s_0, \{s_0\}))$ (resp. $((s_0, \{s_0\}), s_0)$).*

Proof. Let m_{Agt} be a move; writing $t = \text{Tab}(s, m_{\text{Agt}})$, we have $\text{Tab}'((s, S), m_{\text{Agt}}) = (t, S \cup \{t\})$. Therefore $\text{Tab}(s, m_{\text{Agt}}) \triangleleft \text{Tab}'(s', m_{\text{Agt}})$. Let (t, S') be a state of \mathcal{G}' ; then we also have $t \triangleleft (t, S')$. If $S' = S \cup \{t\}$ then $\text{Susp}((s, t), m_{\text{Agt}}) = \text{Susp}((s, S), (t, S'), m_{\text{Agt}})$; otherwise $\text{Susp}((s, S), (t, S'), m_{\text{Agt}}) = \emptyset$. In both cases, condition (2) in the definition of a game simulation is obviously satisfied.

In the other direction, let $(s', S \cup \{s'\}) = \text{Tab}((s, S), m_{\text{Agt}})$; we have that $s' \triangleleft (s', S \cup \{s'\})$. Let $t \in \text{States}$. Then $t \triangleleft (t, S \cup \{t\})$, and $\text{Susp}((s, t), m_{\text{Agt}}) = \text{Susp}((s, S), (t, S \cup \{t\}), m_{\text{Agt}})$. Hence \triangleleft^{-1} is a game simulation.

Let ρ and ρ' be two paths, from s_0 and $(s_0, \{s_0\})$ respectively, and such that $\rho \triangleleft \rho'$. We show preference preservation, by showing that ρ reaches target set T if, and only if,

ρ' visits T' infinitely often. If ρ visits some state $s \in T$, then from that point, states visited by ρ' are of the form (s', S') with $s \in S'$; all these states are in T' , therefore ρ' visits T' infinitely often. Conversely, if ρ' visits T' infinitely often, then some state of T' have been visited by ρ . From this, we easily obtain preference preservation. \square

As a corollary (Proposition 5.23) we get that there is a correspondence between Nash equilibria in \mathcal{G} and Nash equilibria in \mathcal{G}' .

Lemma 7.3. *If there is a Nash equilibrium σ_{Agt} in \mathcal{G} from s_0 , then there is a Nash equilibrium σ'_{Agt} in \mathcal{G}' from $(s_0, \{s_0\})$ such that $\text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}}) \triangleleft \text{Out}_{\mathcal{G}'}((s_0, \{s_0\}), \sigma'_{\text{Agt}})$. And vice-versa: if there is a Nash equilibrium σ'_{Agt} in \mathcal{G}' from $(s_0, \{s_0\})$, then there is a Nash equilibrium σ_{Agt} in \mathcal{G} from s_0 such that $\text{Out}_{\mathcal{G}'}((s_0, \{s_0\}), \sigma'_{\text{Agt}}) \triangleleft^{-1} \text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}})$.*

Note that, if $\text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}}) \triangleleft \text{Out}_{\mathcal{G}'}((s_0, \{s_0\}), \sigma'_{\text{Agt}})$, then $\text{Out}_{\mathcal{G}}(s_0, \sigma_{\text{Agt}})$ satisfies the reachability objective with target set T if, and only if, $\text{Out}_{\mathcal{G}'}((s_0, \{s_0\}), \sigma'_{\text{Agt}})$ satisfies the Büchi objective with target set $T' = \{(s, S) \mid S \cap T \neq \emptyset\}$. From this strong correspondence between \mathcal{G} and \mathcal{G}' , we get that it is sufficient to look for Nash equilibria in game \mathcal{G}' .

7.1.2. How to efficiently solve the suspect game of \mathcal{G}' . In game \mathcal{G}' , preference relations are prefix-independent. Applying Remark 4.6 the preference relation in the suspect game is then also prefix-independent, and the payoff of a play only depends on which strongly-connected component the path ends in. We now give an alternating algorithm which runs in polynomial time and solves the game $\mathcal{H}(\mathcal{G}', \pi')$, where π' is an infinite path in \mathcal{G}' .

Lemma 7.4. *The winner of $\mathcal{H}(\mathcal{G}', \pi')$ can be decided by an alternating algorithm which runs in time polynomial in the size of \mathcal{G} .*

Proof. Let C^A be the circuit defining the preference relation of player A . Let $\rho = (s_i, S_i)_{i \geq 0}$ be a path in \mathcal{G}' , the sequence $(S_i)_{i \geq 0}$ is non-decreasing and converges to a limit $S(\rho)$. We have $\text{payoff}_A(\rho) = \mathbf{1}_{\{i \mid T_A^i \cap S(\rho) = \emptyset\}}$. Therefore the winning condition of Eve in $\mathcal{H}(\mathcal{G}', \pi')$ for a play ρ only depends on the limits $\lambda(\rho)$ and $S(\text{proj}_1(\rho))$. It can be described as a single Büchi condition with target set $T = \{((s, S), P) \mid \forall A \in P. C^A[v^A(S), w^A] \text{ evaluates to true}\}$ where $v^A(S) = \mathbf{1}_{\{i \mid T_A^i \cap S = \emptyset\}}$ and $w^A = \text{payoff}_A(\pi')$. We now describe the algorithm.

Initially the current state is set to $((s_0, \{s_0\}), \text{Agt})$. We also keep a list of the states which have been visited, and we initialise it with $\text{Occ} \leftarrow \{(s_0, \{s_0\}), \text{Agt}\}$. Then,

- if the current state is $((s, S), P)$, the algorithm existentially guesses a move m_{Agt} of Eve and we set $t = ((s, S), P, m_{\text{Agt}})$;
- otherwise if the current state is of the form $((s, S), P, m_{\text{Agt}})$, it universally guesses a state s' which corresponds to a move of Adam and we set $t = ((s', S \cup \{s'\}), P \cap \text{Susp}((s, s'), m_{\text{Agt}}))$.

If t was already seen (that is, if $t \in \text{Occ}$), the algorithm returns **true** when $t \in T$ and **false** when $t \notin T$, otherwise the current state is set to t , and we add t to the list of visited states: $\text{Occ} \leftarrow \text{Occ} \cup \{t\}$, and we repeat this step. Because we stop when the same state is seen, the algorithm stops after at most $\ell + 1$ steps, where ℓ is the length of the longest acyclic path. Since the size of S can only increase and the size of P only decrease, we bound ℓ with $|\text{States}|^2 \cdot |\text{Agt}|$.

We now prove the correctness of the algorithm. First, $\mathcal{H}(\mathcal{G}', \pi')$ is a turn-based Büchi game, which is a special case of parity game. Parity games are known to be determined with memoryless strategies [34, 19], hence $\mathcal{H}(\mathcal{G}', \pi')$ is determined with memoryless strategies.

If the algorithm returns **true**, then there exist a strategy σ_{\exists} of **Eve** such that for all the strategies σ_{\forall} of **Adam**, any outcome ρ of $\text{Out}(\sigma_{\exists}, \sigma_{\forall})$ is such that there exist $i < j \leq \ell + 1$ with $\rho_i = \rho_j \in T$ and all ρ_k with $k < j$ are different. We extend this strategy σ_{\exists} to a winning strategy σ'_{\exists} for **Eve**. We do so by ignoring the loops we see in the history, formally we inductively define a reduction r of histories by:

- $r(\varepsilon) = \varepsilon$;
- if $((s, S), P)$ does not appear in $r(h)$ then $r(h \cdot ((s, S), P)) = r(h) \cdot ((s, S), P)$;
- otherwise $r(h \cdot ((s, S), P)) = r(h)_{\leq i}$ where i is the smallest index such that $r(h)_i = ((s, S), P)$.

We then define σ'_{\exists} for any history h by $\sigma'_{\exists}(h) = \sigma_{\exists}(r(h))$.

We show by induction that if h is a history compatible with σ'_{\exists} from $((s_0, \{s_0\}), \text{Agt})$ then $r(h)$ is compatible with σ_{\exists} from $((s_0, \{s_0\}), \text{Agt})$. It is true when $h = ((s_0, \{s_0\}), \text{Agt})$, now assuming it holds for all history of length $\leq k$, we show it for history of length $k + 1$. Let $h \cdot s$ be a history of length $k + 1$ compatible with σ'_{\exists} . By hypothesis $r(h)$ is compatible with h and since $\sigma'_{\exists}(h) = \sigma_{\exists}(r(h))$, $r(h) \cdot s$ is compatible with σ_{\exists} . If $r(h \cdot s) = r(h) \cdot s$ then $r(h \cdot s)$ is compatible with σ_{\exists} . Otherwise $r(h \cdot s)$ is a prefix of $r(h)$ and therefore of length $\leq k$, we can apply the induction hypothesis to conclude that $r(h \cdot s)$ is compatible with σ_{\exists} .

We now show that the strategy σ'_{\exists} that we defined, is winning. Let ρ be a possible outcome of σ'_{\exists} , let $i < j$ be the first indexes such that $\rho_i, \rho_j \in (\text{States} \times S(\rho)) \times \lambda(\rho)$ and $\rho_i = \rho_j$. Because there is no repetition between i and $j - 1$: $r(\rho_{\leq j-1}) = r(\rho_{\leq i-1})\rho_i \cdots \rho_{j-1}$. We have that $\sigma_{\exists}(r(\rho_{\leq i-1})\rho_i \cdots \rho_{j-1}) = \sigma'_{\exists}(\rho_{j-1})$. From this move, ρ_j is a possible next state, so $r(\rho_{\leq i-1})\rho_i \cdots \rho_j$ is a possible outcome of σ_{\exists} . As $\rho_i = \rho_j$ and all other states are different, by the hypothesis on σ_{\exists} we have that $\rho_j \in T$. This shows that ρ ultimately loops in states of T and therefore ρ is a winning run for **Eve**.

Reciprocally, if **Eve** has a winning strategy, she has a memoryless one σ_{\exists} since this is a Büchi game. We can see this strategy as an oracle for the various existential choices in the algorithm. Consider some universal choices in the algorithm, it corresponds to a strategy σ_{\forall} for **Adam**. The branch corresponding to $(\sigma_{\exists}, \sigma_{\forall})$ ends the first time we encounter a loop, we write this history $h \cdot h'$ with $\text{last}(h') = \text{last}(h)$. Since the strategy σ_{\exists} is memoryless, $h \cdot h'^{\omega}$ is a possible outcome. Since it is winning, $\text{last}(h')$ is in T and therefore the branch is accepting. This being true for all the branches given by the choices of σ_{\exists} , the algorithm answers **true**. \square

7.1.3. Proof of the PSPACE upper bounds in Proposition 7.1. We describe a PSPACE algorithm for solving the constrained NE existence problem. The algorithm proceeds by trying all plays π in \mathcal{G} of the form described in Proposition 3.1. This corresponds to a (unique) play π' in \mathcal{G}' . We check that π' has a payoff satisfying the constraints, and that there is a path ρ in $\mathcal{H}(\mathcal{G}', \pi')$, whose projection is π' , along which **Adam** obeys **Eve**, and which stays in the winning region of **Eve**. This last step is done by using the algorithm of Lemma 7.4 on each state ρ goes through. All these conditions are satisfied exactly when the conditions of Theorem 4.5 are satisfied, in which case there is a Nash equilibrium within the given bounds.

The PSPACE upper bound for the value problem can be inferred from Proposition 3.2.

7.1.4. Proof of PSPACE-hardness for the value problem. We show PSPACE-hardness of the value problem when the preorder has $\mathbf{1}$ as a unique maximal element.

We reduce QSAT to the value problem, where QSAT is the satisfiability problem for quantified Boolean formulas. For an instance of QSAT, we assume without loss of generality that the Boolean formula is a conjunction of disjunctive clauses⁸.

Let $\phi = Q_1x_1 \dots Q_px_p \cdot \phi'$, where $Q_i \in \{\forall, \exists\}$ and $\phi' = c_1 \wedge \dots \wedge c_n$ with $c_i = \bigvee_{1 \leq j \leq 3} \ell_{i,j}$ and $\ell_{i,j} \in \{x_k, \neg x_k \mid 1 \leq k \leq p\} \cup \{\top, \perp\}$. We define a turn-based game $\mathcal{G}(\phi)$ in the following way (illustrated in Example 7.6 below). There is one state for each quantifier, one for each literal, and two additional states \top and \perp :

$$\text{States} = \{Q_k \mid 1 \leq k \leq p\} \cup \{x_k, \neg x_k \mid 1 \leq k \leq p\} \cup \{\top, \perp\}.$$

The game involves two players, A and B . The states \top , and \perp , the existential-quantifier states and the literal states are all controlled by A , while the universal-quantifier states belong to player B . For all $1 \leq k \leq p$, the state corresponding to quantifier Q_k has two outgoing transitions, going to x_k and $\neg x_k$ respectively. Those two literal states only have one transition to the next quantifier state Q_{k+1} , or to the final state \top if $k = p$. Finally, states \top and \perp carries a self-loop (notice that \perp is not reachable, while \top will always be visited).

Player A has one target set for each clause: if $c_i = \bigwedge_{1 \leq j \leq 3} \ell_{i,j}$ then $T_i^A = \{\ell_{i,j} \mid 1 \leq j \leq 3\}$. The i -th objective Ω_i^A is to reach target set T_i^A . The following result is then straightforward:

Lemma 7.5. *Formula ϕ is valid if, and only if, player A has a strategy whose outcomes from state Q_1 all visit each target set T_i^A .*

Proof. We begin with the direct implication, by induction on p . For the base case, $\phi = Q_1x_1 \cdot \bigwedge_i c_i$ where c_i only involves x_1 and $\neg x_1$. We consider two cases:

- $Q_1 = \exists$: since we assume ϕ be true, there must exist a value for x_1 which makes all clauses true. If this value is \top , consider the strategy σ_\top of Player A such that $\sigma_\top(Q_1) = x_1$. Then each clause c_i must have x_1 as one of its literals, so that the objective Ω_i^A is satisfied with this strategy. The same argument applies if the value for x_1 were \perp .
- $Q_1 = \forall$: in that case, Player A has only one strategy. For both x_1 and $\neg x_1$ all the clauses are satisfied. It follows that each clause c_i must contain x_1 and $\neg x_1$, so that objective Ω_i^A is satisfied for any strategy of player B .

Now, assume that the result holds for all QSAT instances with at most $p - 1$ quantifiers.

- if $Q_1 = \exists$, then one of $Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \top]$ and $Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \perp]$ is valid. We handle the first case, the second one being symmetric. For a literal $\lambda_k \in \{x_k, \neg x_k\}$, we write T_{λ_k} for the set of target sets T_i^A such that the clause c_i contains the literal λ_k .

Assume $Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \top]$ is valid; by induction we know that there exists a strategy σ^{x_1} such that all the targets in T_{λ_k} are visited along any outcome from state Q_2 (because $\mathcal{G}(Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \top])$ is the same game as $\mathcal{G}(\phi)$, but with Q_2 as the initial state, and with the targets in T_{x_1} containing $\{\top\}$ in place of x_1). We define the strategy σ by $\sigma(Q_1) = x_1$ and $\sigma(Q_1 \cdot x_1 \cdot \rho) = \sigma^{x_1}(\rho)$. An outcome of σ will necessarily visit x_1 , hence visiting all the targets in T_{x_1} ; because σ follows σ^{x_1} , all the objectives not in T_{x_1} are met as well.

- if $Q_1 = \forall$, then $Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \top]$ is valid. Using the induction hypothesis we know that from Q_2 there is a strategy σ^{x_1} that enforces a visit to all the targets in T_{x_1} . Similarly, $Q_2x_2 \dots Q_px_p\phi'[x_1 \leftarrow \perp]$ is valid, and there is a strategy $\sigma^{\neg x_1}$ that visits all the

⁸With the convention that an empty disjunction is equivalent to \perp .

objectives not in $T_{\neg x_1}$. We define a new strategy σ as follows: $\sigma(Q_1 \cdot x_1 \cdot \rho) = \sigma^{x_1}(\rho)$ and $\sigma(Q_1 \cdot \neg x_1 \cdot \rho) = \sigma^{\neg x_1}(\rho)$. Consider an outcome of σ : if it visits x_1 , then all the objectives in T_{x_1} are visited, and because the path follows σ^{x_1} , the objectives not in T_{x_1} are also visited. The other case is similar.

We now turn to the converse implication. Assume the formula is not valid. We prove that for any strategy σ of player A , there is an outcome ρ of this strategy such that some objective Ω_i^A is not satisfied. We again proceed by induction, beginning with the case where $n = 1$.

- if $Q_1 = \exists$, then both $\phi'[x_1 \leftarrow \top]$ and $\phi'[x_1 \leftarrow \perp]$ are false. This entails that one of the clauses only involves \perp (no other disjunction involving x_1 and/or $\neg x_1$ is always false), and the corresponding reachability condition is \perp , which is not reachable.
- if $Q_1 = \forall$, then one of $\phi'[x_1 \leftarrow \top]$ and $\phi'[x_1 \leftarrow \perp]$ is false. In the former case, one of the clauses c_i contains $\neg x_1$, or only contains \perp . Then along the run $Q_1 \cdot x_1 \cdot \top^\omega$, the objective T_i^A is not visited. The other case is similar.

Now, assuming that the result holds for formulas with $n - 1$ quantifiers, we prove the result with n quantifiers.

- if $Q_1 = \exists$, then both $Q_2 x_2 \dots Q_p x_p \phi'[x_1 \leftarrow \top]$ and $Q_2 x_2 \dots Q_p x_p \phi'[x_1 \leftarrow \perp]$ are false. Using the induction hypothesis, any run from Q_2 fails to visit some objective not in $T_{x_1} \cup T_{\neg x_1}$. Hence no strategy from Q_1 can enforce a visit to all the objectives.
- if $Q_1 = \forall$, then one of $Q_2 x_2 \dots Q_p x_p \phi'[x_1 \leftarrow \top]$ and $Q_2 x_2 \dots Q_p x_p \phi'[x_1 \leftarrow \perp]$ is false. We handle the first case, the second one being symmetric. By induction hypothesis, for any strategy σ of player A in the game $\mathcal{G}(\phi'[x_1 \leftarrow \top])$, one of the outcome fails to visit all the objective not in T_{x_1} . Then along the path $\rho = Q_1 \cdot x_1 \cdot \rho'$, some objectives not in T_{x_1} are not visited. \square

We can directly conclude from this lemma that the value of the game for A is **1** (the unique maximal payoff for our preorder) if, and only if, the formula ϕ is valid, which proves that the former problem is PSPACE-hard.

Example 7.6. As an example of the construction, let us consider the formula

$$\phi = \forall x_1. \exists x_2. \forall x_3. \exists x_4. (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_1 \vee x_2 \vee x_4) \wedge (\neg x_4 \vee \perp \vee \perp) \quad (7.1)$$

The target sets for player A are given by $T_1^A = \{x_1; \neg x_2; \neg x_3\}$, $T_2^A = \{x_1; x_2; x_4\}$, and $T_3^A = \{\neg x_4; \perp\}$. The structure of the game is represented in Figure 25. B has a strategy that falsifies one of the clauses whatever A does, which means that the formula is not valid.




FIGURE 25. Reachability game associated with the formula (7.1)

7.1.5. *Proof of PSPACE-hardness for the (constrained) NE existence problem.* We will now prove PSPACE-hardness for the NE existence problem, under the conditions specified in the statement of Proposition 7.1, using Proposition 3.4. We specify the new preference relation for the construction of Section 3.3. We give B one objective, which is to reach s_1 (s_1 is the sink state introduced by the construction). In terms of preferences for A , going s_1 should be just below visiting all targets. For this we use the statement in Proposition 7.1, that there is v such that for every v' , $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$, and add s_1 as a target to each T_i^A such that $v_i = 1$. This defines a preference relation equivalent to the one in the game constructed in Section 3.3, therefore we deduce with Proposition 3.4 that the NE existence problem is PSPACE-hard.

7.1.6. *Applications.* We should now notice that conjunction, counting and lexicographic preorders (thanks to the fact that $\mathbf{1}$ is the unique maximal element for these orders and to Lemma 6.19). As conjunction (for instance) can easily be encoded using a (monotonic) Boolean circuit in polynomial time, the hardness results are also valid if the preorder is given by a (monotonic) Boolean circuit. Finally the subset preorder can be expressed as a polynomial-size Boolean circuit and has a maximal element. We therefore get the following summary of results:

Corollary 7.7. – *For finite games with ordered reachability objectives, with either the conjunction, the counting or the lexicographic preorder, the value problem, the NE existence problem and the constrained NE existence problem are PSPACE-complete.*

- *For finite games with ordered reachability objectives, where the preorders are given by (monotonic) Boolean circuits, the value problem, the NE existence problem and the constrained NE existence problem are PSPACE-complete.*
- *For finite games with ordered reachability objectives, with the subset preorder, the value problem is PSPACE-complete.*

On the other hand, the disjunction and maximise preorders do not have a unique maximal element, so the hardness result does not carry over to these preorders. In the same way, for the subset preorder, there is no v such that $v' \neq \mathbf{1} \Leftrightarrow v' \lesssim v$, so the hardness result does not apply. We prove later (in Section 7.2) that in these special cases, the complexity is actually lower.

7.2. **Simple cases.** As for ordered Büchi objectives, for some ordered reachability objectives, the preference relation can be (efficiently) (co-)reduced to a single reachability objective. We do not give the formal definitions, they can easily be inferred from that for Büchi objectives on page 48.

Proposition 7.8. – *For finite games with ordered reachability objectives which are reducible to single reachability objectives and in which the preorders are non-trivial, the value problem is P-complete.*

- *For finite games with ordered reachability objectives which are co-reducible to single reachability objectives, and in which the preorders are non-trivial, the NE existence problem and the constrained NE existence problem are NP-complete.*

Proof. Since P-hardness (resp. NP-hardness) already holds for the value (resp. NE existence) problem with a single reachability objective (see [23, Sect. 2.5.1]), we only focus on the upper bounds.

We begin with the value problem: given a payoff vector u for player A , we build the new target set \widehat{T} in polynomial time, and then use a classical algorithm for deciding whether A has a winning strategy (see [23, Sect. 2.5.1]). If she does, then she can secure payoff u .

Consider now the constrained NE existence problem, and assume that the preference relation for each player A is given by target sets $(T_i^A)_{1 \leq i \leq n_A}$. The NP-algorithm consists in guessing the payoff vector $(v_A)_{A \in \text{Agt}}$ and an ultimately periodic play $\rho = \pi \cdot \tau^\omega$ with $|\pi|, |\tau| \leq |\text{States}|^2$, which, for each A , visits T_i^A if, and only if, $v_i^A = 1$. We then co-reduce the payoff to a new target set $\widehat{T}^A(v^A)$ for each player A .

The run ρ is the outcome of a Nash equilibrium with payoff $(v_A)_{A \in \text{Agt}}$ for the original preference relation if, and only if, ρ is the outcome of a Nash equilibrium with payoff 0 with the single reachability objective $\widehat{T}^A(v^A)$ for each $A \in \text{Agt}$. Indeed, in both cases, this is equivalent to the property that no player A can enforce a payoff greater than v^A . Applying the algorithm presented in Section 5.1. this condition can be checked in polynomial time. \square

We now see to which ordered objectives this result applies. It is not difficult to realise that the same transformations as those made in the proof of Lemma 6.5 can be made as well for reachability objectives. We therefore get the following lemma, from which we get the remaining results in Table 3.

Lemma 7.9. *Ordered reachability objectives with disjunction or maximise preorders are reducible to single reachability objectives. Ordered reachability objectives with disjunction, maximise or subset preorders are co-reducible to single reachability objectives.*

We conclude with stating the following corollary:

Corollary 7.10. – *For finite games with ordered reachability objectives, with either the disjunction or the maximise preorder, the value problem is P-complete.*

– *For finite games with ordered reachability objectives, with either the disjunction, the maximise or the subset preorder, the NE existence problem and the constrained NE existence problem are NP-complete.*

8. CONCLUSION

Summary and impact of the results. Concurrent games are a natural class of games, extending classical turn-based games with more complex interactions. We have developed a complete methodology, involving new techniques, for computing pure Nash equilibria in this class of games. We were able to characterise the complexity of finding Nash equilibria (possibly with constraints on the payoff) for simple qualitative objectives first (Section 5), and then for semi-quantitative objectives (Section 6 and 7). We would like to point out that the algorithm for Büchi objectives with maximise preorder (see Section 6.2) has been implemented in tool Praline⁹ [8]

We believe the methodology we have developed in this paper can be used in many other contexts, and the suspect game is a very powerful tool that will allow to analyze various properties of multi-agent systems. Indeed, the correspondence between pure Nash equilibria in the original game and winning strategies in the suspect game holds with no assumption on the structure of the game. In particular it can be applied to games given as

⁹Available on <http://www.lsv.ens-cachan.fr/Software/praline/>

pushdown systems, counter systems, etc. Also it does not assume anything on the preference relations, only the resulting winning condition in the suspect game can become very complex if the preference relations are complex. Now the matter is just algorithmics, in that we have to solve a two-player turn-based game in a potentially complex arena (if the original game structure is complex) with a potentially complex winning condition (if the preference relations are complex).

The suspect game construction can also be adapted to compute many other kinds of equilibria; this is for instance applied to robust equilibria in [9]. We believe this can be used in many other contexts.

We have also developed in this paper another tool that might have its own interest and be useful in some other contexts: the game-simulation (see Section 5.7.1). We used this tool several times (for handling objectives given by deterministic Rabin automata, but also for handling ordered reachability objectives). This tool can also be used to handle more complex game structures, like we did in [5] for timed games, when we originally introduced this notion. In particular, the construction done in [5] shows that we can compute Nash equilibria for timed games with all kinds of objectives studied in the current paper.

Our future researches will include extending the use of the suspect game abstraction for other families of games, and to push it further to also handle truly quantitative objectives.

Discussion on the various hypotheses made in this paper. We have assumed strategies are pure, and game structures are deterministic. This is indeed a restriction, and allowing for randomised strategies would be of great interest. Note however that pure Nash equilibria are resistant to malicious randomised players (that is, to deviations by randomised strategies). There is no obvious way to modify the suspect game construction to handle either stochastic game structures or randomised strategies. Indeed, given a history, it is hard to detect strategy deviations if they can be randomised, and therefore the set of suspects is hard to compute (and actually even define). This difficulty is non-surprising, since the existence of a Nash equilibrium in pure or randomised strategies is undecidable for stochastic games with reachability or Büchi objectives [42], and the existence of a Nash equilibrium in randomised strategies is undecidable for deterministic games [41]. However we would like to exhibit subclasses of stochastic games for which we can synthesize randomised Nash equilibria, this is part of our research programme.

We have assumed that strategies are based on histories that only record states which have been visited, and not actions which have been played. We believe this is more relevant in the context of distributed systems, where only the effect of an action might be visible to other players. Furthermore, this framework is more general than the one where every player could see the actions of the other players, since the latter can easily be encoded in the former. In the context of complete information (precise view of the actions), computing Nash equilibria is rather easy since, once a player has deviated from the equilibrium, all the other players know it and can make a coalition against that player. To illustrate that simplification, we only mention that the constrained NE existence problem falls in NP for finite games with single parity objectives (we can obtain this bound based on the suspect game construction), if we assume that strategies can observe actions, whereas the problem is $P_{\parallel}^{\text{NP}}$ -hard if strategies do not observe the actions.

Finally we have chosen the framework of concurrent games, and not that of turn-based games as is often the case in the literature. Concurrent games naturally appear when studying timed games [5] (the semantics of a timed game is that of a concurrent game, and

the abstraction based on regions that is correct for timed games is concurrent), and in the context of distributed systems, concurrent moves are also very natural. In fact turn-based games are even a simpler case of concurrent games when we assume strategies can see the actions. Of course, the suspect game construction applies to turn-based games, but becomes quite simple (as is the case if strategies do see actions), since the set of suspect players is either the set Agt of all players (this is the case as long as no player has deviated from the equilibrium), or reduces to a singleton, as soon as a player has deviated. To illustrate this simplification, we notice that in the turn-based finite games, the constrained NE existence problem is NP-complete for single parity objectives [40] (it is $\text{P}_{\parallel}^{\text{NP}}$ -complete in finite concurrent games).

REFERENCES

- [1] L. d. Alfaro, M. Faella, Th. A. Henzinger, R. Majumdar, and M. Stoelinga. The element of surprise in timed games. In *Proc. 14th International Conference on Concurrency Theory (CONCUR'03)*, LNCS 2761, p. 142–156. Springer, 2003.
- [2] R. Alur, Th. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. of the ACM*, 49:672–713, 2002.
- [3] E. Asarin, O. Maler, A. Pnueli, and J. Sifakis. Controller synthesis for timed automata. In *Proc. IFAC Symposium on System Structure and Control*, p. 469–474. Elsevier Science, 1998.
- [4] P. Bouyer, R. Brenguier, and N. Markey. Nash equilibria for reachability objectives in multi-player timed games. In *Proc. 21th International Conference on Concurrency Theory (CONCUR'10)*, LNCS 6269, p. 192–206. Springer, 2010.
- [5] P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Nash equilibria in concurrent games with Büchi objectives. In *Proc. 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'11)*, LIPIcs 13, p. 375–386. Leibniz-Zentrum für Informatik, 2011.
- [6] P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Concurrent games with ordered objectives. In *Proc. 15th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'12)*, LNCS 7253, p. 301–315. Springer, 2012.
- [7] R. Brenguier. *Nash Equilibria in Concurrent Games – Application to Timed Games*. PhD thesis, ENS Cachan, 2012.
- [8] R. Brenguier. PRALINE: A tool for computing Nash equilibria in concurrent games. In *Proc. 25th International Conference on Computer Aided Verification (CAV'13)*, LNCS 8044, p. 890–895. Springer, 2013.
- [9] R. Brenguier. Robust equilibria in concurrent games. Technical Report abs/1311.7683, CoRR, 2013. <http://arxiv.org/abs/1311.7683>.
- [10] K. Chatterjee. Two-player nonzero-sum ω -regular games. In *Proc. 16th International Conference on Concurrency Theory (CONCUR'05)*, LNCS 3653, p. 413–427. Springer, 2005.
- [11] K. Chatterjee, T. A. Henzinger, and M. Jurdziński. Games with secure equilibria. In *Proc. 3rd International Symposium on Formal Methods for Components and Objects (FMCO'04)*, p. 141–161. Springer, 2005.
- [12] K. Chatterjee, T. A. Henzinger, and N. Piterman. Generalized parity games. In *Proc. 10th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'07)*, LNCS 4423, p. 153–167. Springer, 2007.
- [13] K. Chatterjee, T. A. Henzinger, and N. Piterman. Algorithms for Büchi games. Research Report 0805.2620, Computing Research Repository, 2008.
- [14] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. & Comp.*, 208(6):677 – 693, 2010. Special Issue: 18th International Conference on Concurrency Theory (CONCUR 2007).
- [15] K. Chatterjee, R. Majumdar, and M. Jurdziński. On Nash equilibria in stochastic games. In *Proc. 18th International Workshop on Computer Science Logic (CSL'04)*, LNCS 3210, p. 26–40. Springer, 2004.
- [16] A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts: Expressiveness and model checking. In *FSTTCS'10*, LIPIcs 8, p. 120–132. Leibniz-Zentrum für Informatik, 2010.

- [17] J. De Pril, J. Flesch, J. Kuipers, G. Schoenmakers, and K. Vrieze. Existence of secure equilibrium in multi-player games with perfect information. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science (MFCS'14) – Part II*, LNCS 8635, p. 213–225, Budapest, Hungary, 2014. Springer.
- [18] E. A. Emerson and C. S. Jutla. The complexity of tree automata and logics of programs. In *Proc. 29th Annual Symposium on Foundations of Computer Science (FOCS'88)*, p. 328–337. IEEE Computer Society Press, 1988.
- [19] E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *Proc. 32nd Annual Symposium on Foundations of Computer Science (FOCS'91)*, p. 368–377. IEEE Computer Society Press, 1991.
- [20] A. M. Fink. Equilibrium in a stochastic n -person game. *Journal of Science in Hiroshima University*, 28(1):89–93, 1964.
- [21] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'10)*, LNCS 6015, p. 190–201. Springer, 2010.
- [22] G. Gottlob. NP trees and Carnap's modal logic. *J. of the ACM*, 42(2):421–457, 1995.
- [23] E. Grädel, W. Thomas, and Th. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.
- [24] T. A. Henzinger. Games in system design and verification. In *Proc. 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK'05)*, p. 1–4, 2005.
- [25] F. Horn. Streett games on finite graphs. In *Proc. 2nd Workshop on Games in Design and Verification (GDV'04)*, 2005.
- [26] F. Horn. Explicit Muller games are PTIME. In *Proc. 28th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'08)*, LIPIcs 2, p. 235–243. Leibniz-Zentrum für Informatik, 2008.
- [27] P. W. Hunter. *Complexity and infinite games on finite graphs*. PhD thesis, University of Cambridge, 2007.
- [28] M. Jurdziński. Deciding the winner in parity games is in $\text{UP} \cap \text{coUP}$. *Information Processing Letters*, 68(3):119–124, 1998.
- [29] D. Kozen. Lower bounds for natural proof systems. In *Proc. 18th Annual Symposium on Foundations of Computer Science (FOCS'77)*, p. 254–266. IEEE Computer Society Press, 1977.
- [30] F. Laroussinie, N. Markey, and G. Oreiby. On the expressiveness and complexity of ATL. *Logical Methods in Computer Science*, 4(2), 2008.
- [31] A. B. MacKenzie and S. B. Wicker. Stability of multipacket slotted Aloha with selfish users and perfect information. In *Proc. 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'03)*. IEEE Computer Society Press, 2003.
- [32] R. McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.
- [33] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In *Proc. 29th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, LIPIcs 8, p. 133–144. Schloss Dagstuhl - LZI, 2010.
- [34] A. W. Mostowski. Games with forbidden positions. Technical Report 78, University of Gdansk, Poland, 1991.
- [35] J. F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1):48–49, 1950.
- [36] Ch. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [37] C. U. Saraydar, N. B. Mandayam, and D. J. Goodman. Pareto efficiency of pricing-based power control in wireless data networks. In *Proc. IEEE Wireless Communications and Networking Conference (WCNC'99)*, p. 231–235. IEEE Computer Society Press, 1999.
- [38] L. S. Shapley. A value for n -person games. *Contributions to the Theory of Games*, 2:307–317, 1953.
- [39] W. Thomas. Infinite games and verification. In *Proc. 14th International Conference on Computer Aided Verification (CAV'02)*, LNCS 2404, p. 58–64. Springer, 2002. Invited Tutorial.
- [40] M. Ummels. The complexity of Nash equilibria in infinite multiplayer games. In *Proc. 11th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS'08)*, LNCS 4962, p. 20–34. Springer, 2008.

- [41] M. Ummels and D. Wojtczak. The complexity of Nash equilibria in limit-average games. In *Proc. 22nd International Conference on Concurrency Theory (CONCUR'11)*, LNCS 6901, p. 482–496. Springer, 2011.
- [42] M. Ummels and D. Wojtczak. The complexity of Nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 7(3), 2011.
- [43] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [44] K. W. Wagner. On restricting the access to an NP-oracle. In *Proc. 15th International Colloquium on Automata, Languages and Programming (ICALP'88)*, LNCS 317, p. 682–696. Springer, 1988.

APPENDIX: PROOF OF PROPOSITION 5.2

We show PSPACE-hardness of the constrained existence of a Nash equilibrium for various kinds of qualitative objectives, using an encoding of the satisfiability of a QSAT formula $\psi = \forall x_1. \exists x_2. \dots \forall x_{p-1}. \exists x_p. \bigwedge_{1 \leq i \leq n} c_i$, where each c_i is of the form $\ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$ and $\ell_{i,j} \in \{x_k, \neg x_k \mid 1 \leq k \leq p\}$.

We construct a game $\mathcal{G}_\psi = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, (\preceq_A)_{A \in \text{Agt}} \rangle$ as follows: $\text{States} = \{u, w\} \cup \bigcup_{k \in [1,p]} \{s_k, t_k, f_k, d_k, e_k\} \cup \bigcup_{i \in [1,n]} \{b_i, c_i\}$; $\text{Agt} = \{\mathbf{Eve}\} \cup \bigcup_{k \in [1,p]} \{A_k, B_k\}$; $\text{Act} = \{0, 1, 2\}$. We now define the transition table (the structure of the game is represented in Figure 26).

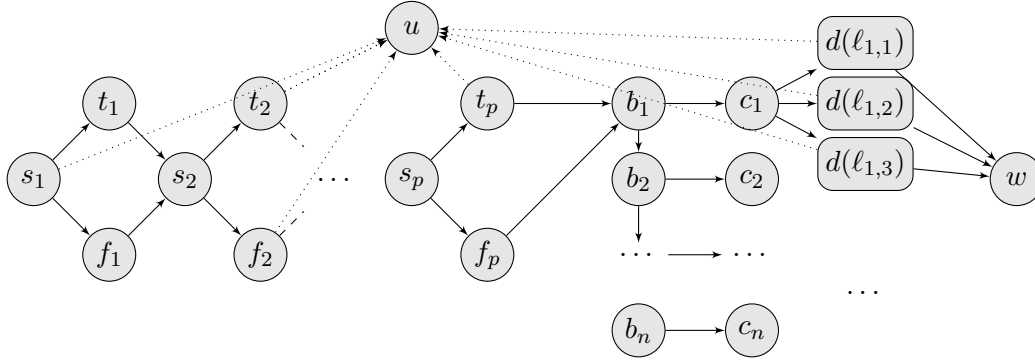


FIGURE 26. Encoding of a QSAT formula into a game with succinct representation of the transition formula. Dotted edges correspond to the strategy profile that in each states selects action 0 for every player.

- If $k \leq p$ is odd, then in state s_k , the transition function is given by:¹⁰

$$\left(\left(\bigotimes_{1 \leq k' \leq p} (A_{k'} = 1) \right) \otimes \left(\bigotimes_{1 \leq k' \leq p, k' \neq k} (B_{k'} = 1), t_k \right), \right. \\ \left. \left(\bigotimes_{1 \leq k' \leq p, k' \neq k} (A_{k'} = 0) \right) \otimes \left(\bigotimes_{1 \leq k' \leq p} (B_{k'} = 0), f_k \right), (\top, u) \right)$$

In the first part, the coalition of all the players except \mathbf{Eve} and B_k takes the decision to go to t_k , and any of those players can switch her action and enforce state t_k (meaning that x_k is set to **true**); if the move to state t_k is not chosen, then the coalition of all players except \mathbf{Eve} and A_k takes the decision to go to f_k , and any of those players can switch her action and enforce state f_k (meaning that x_k is set to **false**); otherwise the game goes to state u .

In states t_k and f_k , the transition function is (\top, s_{k+1}) .

- If $k \leq p$ is even, then in state s_k the transition function is given by $((\mathbf{Eve} = 1, t_k), (\top, f_k))$: \mathbf{Eve} decides the value of variable x_k (state t_k corresponds to setting x_k to **true**, and state f_k corresponds to setting variable x_k to **false**).

¹⁰The operator \bigotimes evaluates the parity of the number of subformulas that are true: $\bigotimes_{h=1}^g \alpha_h$ is **true** iff $|\{\alpha_h \mid \alpha_h \text{ evaluates to true}\}|$ is odd.

In state t_k , the transition function is given by

$$\left(\left(\bigvee_{1 \leq k' \leq p} (A_{k'} = 1) \vee \bigvee_{1 \leq k' \leq p, k' \neq k} (B_{k'} = 1), s_{k+1} \right), (\top, u) \right)$$

with $s_{p+1} = b_1$: any player except **Eve** and B_k can decide to go to state s_{k+1} by playing action 1; otherwise the game proceeds to state u . Intuitively, any of the above players can “validate” the previous choice of **Eve** having set x_k to **true**.

In state f_k , the transition function is given by

$$\left(\left(\bigvee_{1 \leq k' \leq p, k' \neq k} (A_{k'} = 0) \vee \bigvee_{1 \leq k' \leq p} (B_{k'} = 0), s_{k+1} \right), (\top, u) \right)$$

with $s_{p+1} = b_1$: any player except **Eve** and A_k can decide to go to state s_{k+1} by playing action 1; otherwise the game proceeds to state u . Intuitively, any of the above players can “validate” the previous choice of **Eve** having set x_k to **false**.

– If $i \leq n$, in b_i , the transition function is given by

$$\left(\left(\bigotimes_{1 \leq k \leq p} ((A_k = 1) \otimes (B_k = 1)), c_i \right), (\top, b_{i+1}) \right)$$

with $b_{n+1} = u$. Intuitively the coalition of all players except **Eve** can decide to go to state c_i , which will mean that they want to check the truth of clause c_i . Moreover, any of those players can switch her action and decide by her own to check this clause.

– If $i \leq n$, in c_i , the transition function is given by

$$\left((\mathbf{Eve} = 1, d(\ell_{i,1})), (\mathbf{Eve} = 2, d(\ell_{i,2})), (\top, d(\ell_{i,3})) \right)$$

where for all $1 \leq k \leq p$, $d(x_k) = d_k$ and $d(\neg x_k) = e_k$. Intuitively **Eve** proves the current clause is satisfied by pointing to the literal which is set to **true**.

In state d_k ($1 \leq k \leq p$), the transition function is given by $((B_k = 1, w), (\top, u))$. Player B_k decides to go to u or w .

In state e_k ($1 \leq k \leq p$), the transition function is given by $((A_k = 1, w), (\top, u))$.

Intuitively, in the game we have just defined, **Eve** will be in charge of properly choosing the value of the existentially quantified variables in ψ . The value of the variables will be given by the history (visiting t_k means variable x_k is set to **true**, whereas visiting f_k means variable x_k is set to **false**). Then, player A_k will be in charge of witnessing that variable x_k is set to **true**, whereas player B_k will be in charge of witnessing that variable x_k is set to **false**. Their role will be clearer in the proof.

The objective for each player A_k, B_k is to reach state w , and for **Eve** to reach state u . This is naturally a reachability objectives but can also be encoded as a Büchi objective or a safety objective where the goal is to avoid state u for A_k and B_k , and avoid v for **Eve**.

We will show that there is a Nash equilibrium in \mathcal{G}_ψ where **Eve** wins if, and only if, ψ is valid.

Before switching to the proof of this equivalence, we define a correspondence between (partial) valuations and histories in the game. with a partial valuation $v: \{x_1, \dots, x_k\} \rightarrow \{\mathbf{true}, \mathbf{false}\}$, we associate the history $h(v) = s_1 w_1 s_2 w_2 \dots w_k s_{k+1}$ where for all $1 \leq k' \leq k$, $w_{k'} = t_{k'}$ (resp. $w_{k'} = f_{k'}$) if $v(x_{k'}) = \mathbf{true}$ (resp. $v(x_{k'}) = \mathbf{false}$). Conversely with every history h in \mathcal{G}_ψ , we associate the partial valuation $v_h: \{x_1, \dots, x_k\} \rightarrow \{\mathbf{true}, \mathbf{false}\}$ such that state s_{k+1} (with $s_{p+1} = b_1$) is the latest such state appearing along h , and $v_h(x_{k'}) = \mathbf{true}$ (resp. **false**) if h visits $t_{k'}$ (resp. $f_{k'}$), for all $1 \leq k' \leq k$.

Assume formula ψ is valid. For all players A_k and B_k we set strategies σ_{A_k} and σ_{B_k} to always play action 2. We now turn to the strategy for **Eve**. Consider a history $h = s_1 \cdots w_{k-1} \cdot s_k$ where $k < p$ is odd. Let v' be the valuation where $v'(x_{k'}) = v_h(x_{k'})$ for all $k' < k$, and $v'(x_k) = 1$. We set $\sigma_{\text{Eve}}(h)$ to be 1 if v' makes the formula $\forall x_{k+1}. \dots \exists x_p. \bigwedge_{1 \leq i \leq n} c_i$ valid, and 0 otherwise. Since ψ is valid, one of the two choices makes the rest of the formula true. This ensures that a history that reaches b_1 and that is compatible with σ_{Eve} will define a valuation that makes $\bigwedge_{1 \leq i \leq n} c_i$ true. Fix a history h that is compatible with σ_{Eve} and ends up in some state c_i : the strategy of **Eve** is to go to $d(\ell_{i,j})$ where the literal $\ell_{i,j}$ makes the clause c_i true under valuation v_h . For all other histories, we set the strategy of **Eve** to be 2.

We show that the strategy profile $\sigma_{\text{Agt}} = (\sigma_{\text{Eve}}, (\sigma_{A_k}, \sigma_{B_k})_{1 \leq k \leq p})$ is a Nash equilibrium. First notice that the outcome of σ_{Agt} is $s_1 \cdot u$ (since all players A_k and B_k play action 2): **Eve** wins, and all the other players lose. We now describe interesting deviating strategies for the players A_k or B_k :

- Consider a deviating strategy σ'_{A_k} for player A_k : let $h \in \text{Out}^f(\sigma_{\text{Agt}}[A_k \mapsto \sigma'_{A_k}])$; if $\sigma'_{A_k}(h) = 2$, then $\text{Out}(\sigma_{\text{Agt}}[A_k \mapsto \sigma'_{A_k}])$ ends up in state u ; therefore an interesting deviating strategy should choose some value 0 or 1 after any history. Now if k' is odd with $k' \neq k$, then from $s_{k'}$, player A_k can choose to go to $t_{k'}$ (action 1) or $f_{k'}$ (action 0). If k is odd, then the only way not to end up in u from s_k is to choose action 1 which leads to state t_k . Now at state $t_{k'}$ with k' even, σ'_{A_k} should validate the choice of **Eve** (that is, play action 1 in $t_{k'}$ – meaning that variable $x_{k'}$ has value **true**). At state $f_{k'}$ with k' even, if $k' \neq k$, σ'_{A_k} should validate the choice of **Eve** (that is, play action 0 in $f_{k'}$ – meaning that variable $x_{k'}$ has value **false**). At state f_k if k is even, nothing can be done which could be profitable to player A_k : state u will be reached.
- A similar reasoning can be done for player B_k : the only difference is at state s_k when k is odd, where player B_k can only choose action 0 and go through f_k .
- In the part of the game after b_1 , each player can deviate and choose to go to some state c_i ; this choice will be made for checking the truth of clause c_i under the valuation that has been fixed by the history so far.

Consider a deviation of some player that moves to c_i , and write h for the corresponding history up to state c_i . The strategy of **Eve** after h is to go to $d(\ell_{i,j})$ where $\ell_{i,j}$ sets c_i to true under valuation v_h . If $d(\ell_{i,j}) = x_k$, then (a) this means that $v_h(x_k) = \mathbf{true}$, and (b) the next state is controlled by player B_k . Using the characterization of interesting deviating strategies above, it cannot be the case that player B_k is the deviating player since from t_k (which is visited by h), if only B_k deviates, the game unavoidably goes to state u . Hence, for every strategy σ'_{B_k} for player B_k , history h cannot be an outcome of $\sigma_{\text{Agt}}[B_k \mapsto \sigma'_{B_k}]$. In particular, no deviation of player B_k can lead to state w . Similarly, if $d(\ell_{i,j}) = \neg x_m$, the outcome ends up in u . In other words, each time a player other than **Eve** changes her strategy, the outcome ends up in u , yielding no improvement for the player.

Hence no player can improve her outcome by changing unilaterally her strategy, which shows that the strategy profile σ_{Agt} is a Nash equilibrium where **Eve** wins.

Now assume there is a Nash equilibrium σ_{Agt} in which **Eve** wins. Let v be a valuation such that for every $2 \leq k' \leq p$ even,

$$(\sigma_{\text{Eve}}(h(v|_{\{x_1, \dots, x_{k'-1}\}})) = 1) \Leftrightarrow (v(x_{k'}) = \mathbf{true}) \quad (\#)$$

where $v|_{\{x_1, \dots, x_{k'-1}\}}$ is the valuation v restricted to $\{x_1, \dots, x_{k'-1}\}$. We show the following two properties:

- if $v(x_k) = \mathbf{true}$ then there is a strategy σ'_{A_k} for A_k s.t. $\mathbf{h}(v) \in \text{Out}^f(\sigma_{\text{Agt}}[A_k \mapsto \sigma'_{A_k}])$;
- if $v(x_k) = \mathbf{false}$ then there is a strategy σ'_{B_k} for B_k s.t. $\mathbf{h}(v) \in \text{Out}^f(\sigma_{\text{Agt}}[B_k \mapsto \sigma'_{B_k}])$.

We show the result by induction on the number of atomic propositions. For zero atomic propositions, the result obviously holds. Assume the result holds for atomic propositions $\{x_1, \dots, x_{h-1}\}$ ($h \leq p$). Let v be a valuation over $\{x_1, \dots, x_h\}$, and k such that $v(x_k)$ is true. Define v' as the restriction of v to atomic propositions $\{x_1, \dots, x_{h-1}\}$. By induction hypothesis, $\mathbf{h}(v') = s_1 \cdot w_1 \cdots w_{h-1} \cdot s_h$ is an outcome of some strategy σ'_{A_k} for player A_k .

- If h is odd. Let $m_{\text{Agt}} = \sigma_{\text{Agt}}(\mathbf{h}(v'))$. We set $\sigma'_{A_k}(\mathbf{h}(v'))$ to be 1 if $\bigotimes_{1 \leq k' \leq p, k' \leq k} (m_{A_{k'}} = 1) \otimes \bigotimes_{1 \leq k' \leq p, k' \neq k} (m_{B_{k'}} = 1)$ is different from $v(x_h)$, and to be 0 otherwise. Then we have that the next state is t_h if, and only if, $v(x_h)$ is true.
- If h is even, then the state after s_h (actually after $\mathbf{h}(v')$) is t_h if $v(x_h)$ is true, and f_h otherwise, and this cannot be changed by player A_k . Then in t_h and f_h we set $\sigma'_{A_k}(\mathbf{h}(v')t_h)$ (resp. $\sigma'_{A_k}(\mathbf{h}(v')f_h)$) to be 1. Note that since $v(x_k)$ is true we cannot reach f_k , hence setting the action of A_k in those states to 1 always ensures that the next state is s_{h+1} .

This shows that $\mathbf{h}(v) \in \text{Out}^f(\sigma_{\text{Agt}}[A_k \mapsto \sigma'_{A_k}])$ for some strategy σ'_{A_k} .

The second property can be proven similarly for player B_k .

Let v be a valuation satisfying condition (#). We show that ψ evaluates to **true** under that valuation. Let c_i be a clause of ψ , and let $j = \sigma_{\text{Eve}}(\mathbf{h}(v) \cdot b_1 \cdots b_l \cdot c_l)$. We show that $v(\ell_{i,j}) = \mathbf{true}$, which means that c_i evaluates to **true** under v . This will show that formula ψ is valid since condition (#) defines sufficiently many witness valuations. Assume w.l.o.g. that $\ell_{i,j} = x_k$. Assume towards a contradiction that $v(x_k) = \mathbf{false}$. We have proven that there is a strategy σ'_{B_k} for player B_k such that $\mathbf{h}(v) \cdot b_1 \cdots b_l \cdot c_l \in \text{Out}^f(\sigma_{\text{Agt}}[B_k \mapsto \sigma'_{B_k}])$. Now, the state x_k is controlled by player B_k , so B_k can enforce a visit to w from x_k , so there is a strategy σ''_{B_k} for player B_k such that $\mathbf{h}(v) \cdot b_1 \cdots b_l \cdot c_l \cdot x_k \in \text{Out}^f(\sigma_{\text{Agt}}[B_k \mapsto \sigma''_{B_k}])$. This contradicts the fact that σ_{Agt} is a Nash equilibrium. We conclude that $v(x_k) = \mathbf{true}$, and we conclude that ψ is valid (as explained above). \square