

# Concurrent Games with Ordered Objectives

Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels

LSV, CNRS & ENS Cachan, France

{bouyer,brenguier,markey,ummels}@lsv.ens-cachan.fr

**Abstract.** We consider concurrent games played on graphs, in which each player has several qualitative (e.g. reachability or Büchi) objectives, and a preorder on these objectives (for instance the counting order, where the aim is to maximise the number of objectives that are fulfilled).

We study two fundamental problems in that setting: (1) the *value problem*, which aims at deciding the existence of a strategy that ensures a given payoff; (2) the *Nash equilibrium problem*, where we want to decide the existence of a Nash equilibrium (possibly with a condition on the payoffs). We characterise the exact complexities of these problems for several relevant preorders, and several kinds of objectives.

## 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 [15,6]. Until recently, more focus had been put on the study of purely antagonistic games (a.k.a. zero-sum games), useful for modelling systems evolving in a (hostile) environment.

Over the last ten years, non-zero-sum games have come into the picture: they are convenient for modelling complex infrastructures where each individual system tries to fulfill its 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, and several other solution concepts have been defined and studied in the literature, of which Nash equilibrium [11] is the most prominent. A Nash equilibrium is a strategy profile where no player can improve her payoff by unilaterally changing her strategy, resulting in a configuration of the network that is satisfactory to everyone. 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.

*Our contributions.* In this paper, we extend our previous study of pure-strategy Nash equilibria in concurrent games with qualitative objectives [2,4] to a (semi-) quantitative setting: we assume that each player is given a set  $S$  of qualitative objectives (reachability, for instance), together with a preorder on  $2^S$ . This preorder

defines a preference relation (or payoff), and the aim of a player is to maximise her payoff. For instance, the counting order compares the number of objectives which are fulfilled. As another example, we will consider the lexicographic order, defined in an obvious way once we have ordered the simple objectives. More generally, preorders will be defined by Boolean circuits.

We characterise the exact complexity of deciding the existence of a Nash equilibrium, for reachability and Büchi objectives, under arbitrary preorders. Our techniques also provide us with solutions to the value problem, which corresponds to the purely antagonistic setting described above. We prove for instance that both problems are PSPACE-complete for reachability objectives together with a lexicographic order on these objectives (or for the more general class of preorders defined by Boolean circuits). On the other hand, we show that for sets of Büchi objectives (assumed to be indexed) ordered by the maximum index they contain, both problems are solvable in PTIME.

*Related work.* Even though works on concurrent games go back to the fifties, the complexity of computing Nash equilibria in games played on graphs has only recently been addressed [5,16]. Most of the works so far have focused on turn-based games and on qualitative objectives, but have also considered the more general setting of stochastic games or strategies. Our restriction to pure strategies is justified by the undecidability of computing Nash equilibria in concurrent games with qualitative reachability or Büchi objectives, when strategies can be randomised [17]. Although their setting is turn-based, the most relevant related work is [14], where a first step towards quantitative objectives is made: they consider generalised Muller games (with a preference order on the set of states that are visited infinitely often), show that pure Nash equilibria always exist, and give a doubly-exponential algorithm for computing a Nash equilibrium. Generalised Muller conditions can be expressed using Büchi conditions and Boolean circuits (which in the worst-case can be exponential-size): from our results we derive an EXPSPACE upper bound.

For lack of space, the technical proofs are omitted, and can be found in [3].

## 2 Preliminaries

### 2.1 Concurrent Games

**Definition 1** ([1]). *A (finite) concurrent game is a tuple  $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab} \rangle$ , where States is a (finite) set of states, Agt is a finite set of players, Act is a finite set of actions, and*

- 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;
- Tab:  $\text{States} \times \text{Act}^{\text{Agt}} \rightarrow \text{States}$  associates with a given state and a given move<sup>1</sup> of the players the resulting state.

---

<sup>1</sup> A move is an element of  $\text{Act}^{\text{Agt}}$ .

Fig. 1 displays an example of a 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  $m_{\text{Agt}}$ ; 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.

A *path*  $\pi$  in  $\mathcal{G}$  is a sequence  $(s_i)_{0 \leq i < n}$  (where  $n \in \mathbb{N}_{>0} \cup \{\infty\}$ ) of states. The *length* of  $\pi$ , denoted by  $|\pi|$ , is  $n - 1$ . The set of finite paths (also called *histories*) of  $\mathcal{G}$  is denoted by  $\text{Hist}_{\mathcal{G}}$ , the set of infinite paths (also called *plays*) of  $\mathcal{G}$  is denoted by  $\text{Play}_{\mathcal{G}}$ , and  $\text{Path}_{\mathcal{G}} = \text{Hist}_{\mathcal{G}} \cup \text{Play}_{\mathcal{G}}$  is the set of paths of  $\mathcal{G}$ . 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  $\pi$ , denoted by  $\pi_{\leq j}$  (resp.  $\pi_{\geq j}$ ,  $\pi_{=j}$ ), is the finite path  $(s_i)_{0 \leq i < j+1}$  (resp.  $(s_i)_{j \leq i < n}$ , state  $s_j$ ). If  $\pi = (s_i)_{0 \leq i < n}$  is a history, we write  $\text{last}(\pi) = s_{|\pi|}$ . In the sequel, we write  $\text{Hist}_{\mathcal{G}}(s)$ ,  $\text{Play}_{\mathcal{G}}(s)$  and  $\text{Path}_{\mathcal{G}}(s)$  for the respective subsets of paths starting in state  $s$ . If  $\pi$  is a play,  $\text{Occ}(\pi) = \{s \mid \exists j. \pi_{=j} = s\}$  is the sets of states that appear at least once along  $\pi$ , and  $\text{Inf}(\pi) = \{s \mid \forall i. \exists j \geq i. \pi_{=j} = s\}$  is the sets of states that appear infinitely often along  $\pi$ .




Fig. 1. Example of a two-player concurrent game  $\mathcal{A}$

**Definition 2.** 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  $\sigma_P$  for a coalition  $P \subseteq \text{Agt}$  is a tuple of strategies, one for each player in  $P$ . We write  $\sigma_P = (\sigma_A)_{A \in P}$  for such a strategy. A strategy profile is a strategy for  $\text{Agt}$ . 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 we only consider *pure* (i.e., non-randomised) strategies. Notice also that strategies are based on the sequences of visited states, and not on the sequences of actions played by the players. This is realistic when considering multi-agent systems, where only the global effect of the actions of the agents may be observable. When computing Nash equilibria, this restriction makes it more difficult to detect which players have deviated from their strategies.

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

1.  $m_{\text{Agt}}$  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  $\mathcal{G}$  which are compatible with strategy  $\sigma_P$  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  $\sigma_P$  with initial state  $s$ . Notice that any strategy profile has a single infinite outcome from a given state.

## 2.2 Winning Objectives

*Objectives and preference relations.* An *objective* (or *winning condition*) is an arbitrary set of plays. With a set  $T$  of states, we associate an objective  $\Omega(T)$  in three different ways:

$$\Omega(T) = \{\rho \in \text{Play}_{\mathcal{G}} \mid \text{Occ}(\rho) \cap T \neq \emptyset\} \quad (\text{Reachability})$$

$$\Omega(T) = \{\rho \in \text{Play}_{\mathcal{G}} \mid \text{Occ}(\rho) \cap T = \emptyset\} \quad (\text{Safety})$$

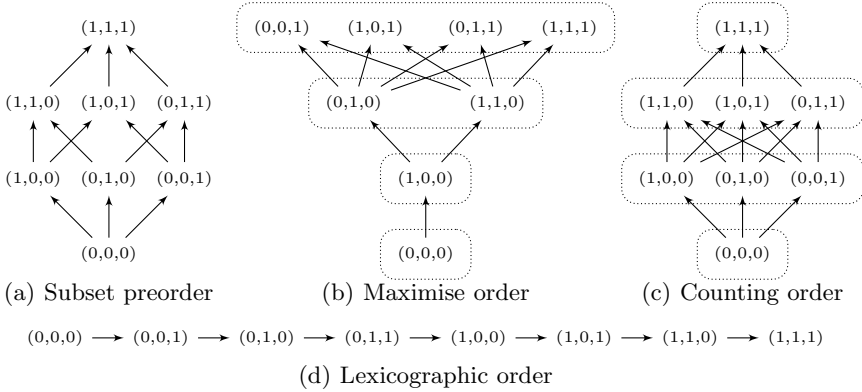
$$\Omega(T) = \{\rho \in \text{Play}_{\mathcal{G}} \mid \text{Inf}(\rho) \cap T \neq \emptyset\} \quad (\text{Büchi})$$

In our setting, each player  $A$  is assigned a tuple of such objectives  $(\Omega_i)_{1 \leq i \leq n}$ , together with a preorder  $\lesssim$  on  $\{0, 1\}^n$ . The *payoff vector* of a play  $\rho$  for player  $A$  is the vector  $\mathbf{1}_{\{i \mid \rho \in \Omega_i\}} \in \{0, 1\}^n$  ( $\mathbf{1}_S$  is the vector  $v$  such that  $v_i = 1 \Leftrightarrow i \in S$ ; we write  $\mathbf{1}$  for  $\mathbf{1}_{[1, n]}$ , and  $\mathbf{0}$  for  $\mathbf{1}_{\emptyset}$ ). The preorder  $\lesssim$  then defines another preorder  $\preceq$  on the set of plays of  $\mathcal{G}$ , called the *preference relation* of  $A$ , by ordering the plays according to their payoffs:  $\rho' \preceq \rho$  if and only if  $\mathbf{1}_{\{i \mid \rho' \in \Omega_i\}} \lesssim \mathbf{1}_{\{i \mid \rho \in \Omega_i\}}$ . Intuitively, each player aims at a play that is preferred to most other plays.

*Examples of preorders.* We now describe some preorders on  $\{0, 1\}^n$  that we consider in the sequel (Fig. 2(a)–2(d) display four of these preorders for  $n = 3$ ). For the purpose of these definitions, we assume that  $\max \emptyset = -\infty$ .

- *Conjunction:*  $v \lesssim w$  iff either  $v_i = 0$  for some  $0 \leq i \leq n$ , or  $w_i = 1$  for all  $0 \leq i \leq n$ . This corresponds to the case where a player wants to achieve all her objectives.
- *Disjunction:*  $v \lesssim w$  iff either  $v_i = 0$  for all  $0 \leq i \leq n$ , or  $w_i = 1$  for some  $0 \leq i \leq n$ . The aim here is to satisfy at least one objective.
- *Counting:*  $v \lesssim w$  iff  $|\{i \mid v_i = 1\}| \leq |\{i \mid w_i = 1\}|$ . The aim is to maximise the number of satisfied conditions.
- *Subset:*  $v \lesssim w$  iff  $\{i \mid v_i = 1\} \subseteq \{i \mid w_i = 1\}$ : in this setting, a player will always struggle to satisfy a larger (for inclusion) set of objectives.
- *Maximise:*  $v \lesssim w$  iff  $\max\{i \mid v_i = 1\} \leq \max\{i \mid w_i = 1\}$ . The aim is to maximise the highest index of the satisfied objectives.
- *Lexicographic:*  $v \lesssim w$  iff either  $v = w$ , or there is an index  $i$  such that  $v_i = 0$ ,  $w_i = 1$  and  $v_j = w_j$  for all  $0 \leq j < i$ .
- *Parity:*  $v \lesssim w$  iff either  $\max\{i \mid w_i = 1\}$  is even, or  $\max\{i \mid v_i = 1\}$  is odd (or  $-\infty$ ). Combined with reachability objectives, this corresponds to a *weak parity condition*; parity objectives as they are classically defined correspond to parity preorders over Büchi objectives.

- *Boolean Circuit*: given a Boolean circuit, with input from  $\{0, 1\}^{2n}$ ,  $v \lesssim w$  if and only if the circuit evaluates to 1 on input  $v_1 \dots v_n w_1 \dots w_n$ .
- *Monotonic Boolean Circuit*: Same as above, with the restriction that the input gates corresponding to  $v$  are negated, and no other negation appears in the circuit.



**Fig. 2.** Examples of preorders (for  $n = 3$ ): dotted boxes represent equivalence classes for  $\sim$ ; arrows represent the preorder relation  $\lesssim$ , forgetting about  $\sim$ -equivalent elements

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 of our examples can be specified with a circuit of polynomial size.

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 the above examples are monotonic, except the parity preorder and the general Boolean circuits. Moreover, any monotonic preorder can be expressed as a monotonic Boolean circuit.

### 2.3 Nash Equilibria

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

**Definition 3.** Let  $\mathcal{G}$  be a concurrent game with preference relation  $(\lesssim_A)_{A \in \text{Agt}}$ , 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  $\text{Out}(s, \sigma_{\text{Agt}}[B \mapsto \sigma']) \lesssim_B \text{Out}(s, \sigma_{\text{Agt}})$  for all players  $B \in \text{Agt}$  and all strategies  $\sigma' \in \text{Strat}^B$ .

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

*Remark 4.* Another possible way of defining Nash equilibrium would be to require that either  $\text{Out}(s, \sigma_{\text{Agt}}[B \mapsto \sigma']) \lesssim_B \text{Out}(s, \sigma_{\text{Agt}})$  or  $\text{Out}(s, \sigma_{\text{Agt}}) \not\lesssim_B \text{Out}(s, \sigma_{\text{Agt}}[B \mapsto \sigma'])$ . This definition is not equivalent to the one we adopted if the preorder is not total, but both can be meaningful. Notice that with our Definition 3, any Nash equilibrium  $\sigma_{\text{Agt}}$  for the subset preorder is also a Nash equilibrium for any monotonic preorder.

## 2.4 Decision Problems

Given a game  $\mathcal{G} = \langle \text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab} \rangle$ , a type of objective (Reachability, Safety or Büchi), for each player a list  $(T_i^A)_{A \in \text{Agt}, i \in \{1, \dots, n_A\}}$  of targets and a preorder  $\lesssim_A$  on  $\{0, 1\}^{n_A}$ , and a state  $s$ , we consider the following problems:

- *Value:* Given a player  $A$  and a payoff vector  $v$ , can player  $A$  ensure payoff  $v$ , i.e., is there a strategy  $\sigma_A$  for player  $A$  such that any outcome of  $\sigma_A$  in  $\mathcal{G}$  from  $s$  with payoff  $v'$  for  $A$  satisfies  $v \lesssim_A v'$ ?
- *Existence:* Does there exist a Nash equilibrium in  $\mathcal{G}$  from  $s$ ?
- *Constrained existence:* Given two vectors  $u^A$  and  $w^A$  for each player  $A$ , does there exist a Nash equilibrium in  $\mathcal{G}$  from  $s$  with some payoff  $(v^A)_{A \in \text{Agt}}$  satisfying the constraint, i.e.,  $u^A \lesssim_A v^A \lesssim_A w^A$  for all  $A \in \text{Agt}$ ?

## 2.5 Preliminary Lemma

We first characterise outcomes of Nash equilibria as ultimately periodic runs.

**Lemma 5.** *Assume that every player has a preference relation which only depends on the set of states that are visited, and the set of states that are visited infinitely often, i.e. if  $\text{Inf}(\rho) = \text{Inf}(\rho')$  and  $\text{Occ}(\rho) = \text{Occ}(\rho')$ , then  $\rho \sim_A \rho'$  for every player  $A \in \text{Agt}$ . If there is a Nash equilibrium with payoff  $v$ , then there is a Nash equilibrium with payoff  $v$  for which the outcome is of the form  $\pi \cdot \tau^\omega$ , where  $|\pi|$  and  $|\tau|$  are bounded by  $|\text{States}|^2$ .*

## 3 Reachability Objectives

Multiplexer games with one reachability objective per player have been studied in [2], where the existence and constrained existence are shown NP-complete.

We now assume that each player has several reachability objectives. In the general case where the preorders are given as Boolean circuits, we show that the various decision problems are PSPACE-complete, where the hardness result even holds for several simpler preorders. We then improve this result in a number of cases. The results are summarised in Table 1.

**Table 1.** Summary of the results for reachability objectives

	Preorder	Value problem	(Constrained) existence
Disjunction, Maximise	Parity	P-c	NP-c
	Subset	P-c [12]	NP-h and in PSPACE
		PSPACE-c	NP-c
Conjunction, Counting, Lexicographic (Monotonic) Boolean Circuit		PSPACE-c	PSPACE-c
		PSPACE-c	PSPACE-c

### 3.1 General Case

**Theorem 6.** *For reachability objectives with preorders given by Boolean circuits, the value, existence and constrained existence problems are in PSPACE. For preorders having  $\mathbf{1}$  as a unique maximal element, the value problem is PSPACE-complete. If moreover there is an element  $v \in \{0, 1\}^n$  such that  $\mathbf{1} \not\lesssim v' \Leftrightarrow v' \lesssim v$ , then the existence and constrained existence problems are PSPACE-complete (even for two-player games).*

Before proving these results, let us first discuss the above conditions. The conjunction, subset, counting and lexicographic preorders have unique maximal element  $\mathbf{1}$ . The conjunction, counting and lexicographic preorders have an element  $v$  such that  $\mathbf{1} \not\lesssim v' \Leftrightarrow v' \lesssim v$ .

As conjunction (for instance) can easily be encoded using a (monotonic) Boolean circuit in polynomial time, the hardness results are also valid if the order is given by a (monotonic) Boolean circuit. On the other hand, disjunction and maximise preorders do not have a unique maximal element, so we cannot apply the hardness result. In the same way, for the subset preorder there is no  $v$  such that  $\mathbf{1} \not\lesssim v' \Leftrightarrow v' \lesssim v$ , so the hardness result does not apply. We prove later (Section 3.2) that in these special cases, the complexity is actually lower.

*Proof of the PSPACE upper bounds.* We first focus on the constrained existence problem, and we fix a game  $\mathcal{G}$  with reachability objectives and a preorder for every player, and a constraint on the payoffs. The algorithm proceeds as follows. For every possible payoff vector that satisfies the given constraint, we will check whether there is an equilibrium with this payoff. Fix such a payoff tuple  $\mathbf{v} = (v^A)_{A \in \text{Agt}}$ . We construct a new game  $\mathcal{G}(\mathbf{v})$ : the structure of  $\mathcal{G}(\mathbf{v})$  is identical to  $\mathcal{G}$ , but each player  $A$  has a single objective given by a 1-weak<sup>2</sup> deterministic Büchi automaton  $\mathcal{A}(v^A)$ . The new game satisfies the following property: there is a Nash equilibrium in  $\mathcal{G}$  with payoff  $\mathbf{v}$  iff there is a Nash equilibrium in  $\mathcal{G}(\mathbf{v})$  with payoff  $\mathbf{0}$  whose outcome has payoff  $\mathbf{v}$  in  $\mathcal{G}$ . Then, applying arguments similar to [4, Thm. 22], we easily design a polynomial-space algorithm for deciding the existence of a Nash equilibrium with a given payoff, and therefore more generally for the constrained existence problem.

<sup>2</sup> That is, each strongly connected component contains exactly one state.

The automata  $\mathcal{A}(v^A)$  are obtained from a common structure  $\mathcal{A}$  (whose construction is illustrated by an example in Fig. 3) by adding the set of accepting states  $F(v^A) = \{S \mid \mathbf{1}_{\{i \mid S \cap T_i^A \neq \emptyset\}} \not\leq v^A\}$ , where  $T_i^A$  is the  $i$ -th target of player  $A$  in  $\mathcal{G}$ . While reading a word  $\rho$  from the initial state, the current state of  $\mathcal{A}(v^A)$  is the set of states that have been seen so far. Hence, if  $v$  is the payoff of  $\rho$  for player  $A$  in game  $\mathcal{G}$ , then  $\mathcal{A}(v^A)$  accepts  $v$  iff  $v \not\leq v^A$ . With this construction, the announced equivalence is straightforward.

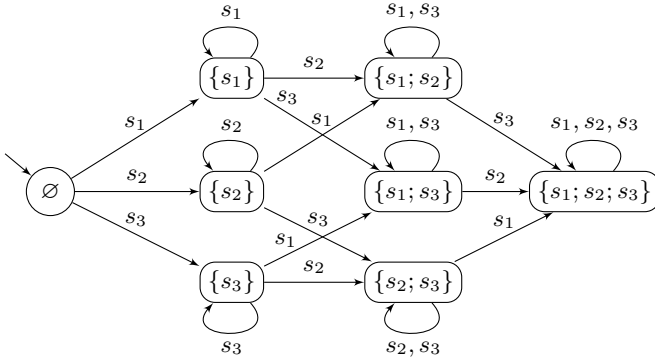


Fig. 3. Automaton  $\mathcal{A}$  for set of states  $\{s_1, s_2, s_3\}$

We now turn to the proof for the value problem. Without loss of generality, we assume that we are given a two-player game  $\mathcal{G}$ , a player  $A$  and a threshold  $v$ . We define a new game (of polynomial size), by only changing the preferences of the players. Player  $A$  has now no objective, and her opponent wins if the payoff is not above  $v$  in the original game, i.e. if the run has payoff  $v'$  with  $v \not\leq v'$ . Then, there is a Nash equilibrium where the opponent loses iff there is a strategy for  $A$  that ensures  $v$  in the original game. We can thereby use the algorithm that decides constrained existence.

*Hardness of the value problem.* The proof is done by encoding an instance of QSAT. Given a formula of QSAT, we construct a two-player turn-based game with several reachability objectives for player  $A$ , such that the formula is valid iff player  $A$  has a strategy that visits all her target sets. We do not give details of the construction but better illustrate it on an example.

*Example 7.* We consider the formula

$$\phi = \forall x_1. \exists x_2. \forall x_3. \exists x_4. (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_1 \vee x_2 \vee x_4) \wedge \neg x_4 \quad (1)$$

The targets for player  $A$  are given by the clauses of  $\phi$ :  $T_1^A = \{x_1, \neg x_2, \neg x_3\}$ ,  $T_2^A = \{x_1, x_2, x_4\}$ , and  $T_3^A = \{\neg x_4\}$ . We fix any preorder with unique maximal element  $(1, 1, 1)$ . The structure of the game is represented in Fig. 4. In this example, player  $B$  has a strategy that falsifies one of the clauses whatever  $A$  does, which means that player  $A$  has no strategy to enforce all its target sets, which means that the formula  $\phi$  is not valid.






Fig. 4. Reachability game associated with the formula (1)

*Hardness of the (constrained) existence problem.* The previous hardness proof applies in particular to conjunctions of reachability objectives. We use a reduction from this problem to prove that the constrained existence problem is PSPACE-hard, under the conditions specified in the statement of Theorem 6. Let  $\mathcal{G}$  be a turn-based game with a conjunction of reachability objectives for player A and  $v$  be a threshold for player A. We construct a new game  $\mathcal{G}'$  as follows. We add to  $\mathcal{G}$  an initial state  $s'_0$ , and a sink state  $z$ . In the initial state  $s'_0$ , the two players A and B play the matching-pennies game, to either go to  $z$  or  $s_0$ .<sup>3</sup> We modify the targets of player A so that, in  $\mathcal{G}'$ , reaching  $z$  exactly gives her payoff  $v$ . The new sink state is the unique target of player B. We can check that if there is no strategy for A ensuring  $v$  in  $\mathcal{G}$ , then there is a Nash equilibrium in game  $\mathcal{G}'$ , which consists in going to  $z$ . And conversely if there is a Nash equilibrium in  $\mathcal{G}'$  then its outcome goes to  $z$ , which means that A cannot ensure  $\mathbf{1}$  in game  $\mathcal{G}$ .

### 3.2 Simple Cases

For some preorders, the preference relation can (efficiently) be reduced to a single reachability objective. For instance, a disjunction of several reachability objectives can obviously be reduced to a single reachability objective, by forming the union of the targets. Formally, we say that a preorder  $\lesssim$  is *reducible* to a single (reachability) objective if, given any payoff vector  $v$ , we can construct in polynomial time a target  $\hat{T}^A$  such that  $v \lesssim \mathbf{1}_{\{i | \text{Occ}(\rho) \cap T_i^A \neq \emptyset\}}$  iff  $\text{Occ}(\rho) \cap \hat{T}^A \neq \emptyset$ . It means that *securing* the payoff corresponds to ensuring a visit to the new target. Similarly, we say that the preorder is *co-reducible* to a single reachability objective, if for any vector  $v$  we can construct  $\hat{T}^A$  such that  $\mathbf{1}_{\{i | \text{Occ}(\rho) \cap T_i^A \neq \emptyset\}} \lesssim v$  if, and only if  $\text{Occ}(\rho) \cap \hat{T}^A \neq \emptyset$ . It means that *improving* the payoff corresponds to reaching the new target. The disjunction and maximise preorders are examples of preorders that are reducible to single reachability objectives. The disjunction, maximise and subset preorders are co-reducible.

**Proposition 8.** *For reachability objectives with a (non-trivial) preorder reducible to a single reachability objective, the value problem is P-complete. For a (non-trivial) preorder co-reducible to a single reachability objective, the existence and constrained existence problems are NP-complete.*

<sup>3</sup> That is, A and B play with two actions 0 and 1, and for instance moves (0, 0) and (1, 1) lead to  $z$  whereas moves (0, 1) and (1, 0) lead to  $s_0$ .

## 4 Safety Objectives

The results for safety objectives are summarised in Table 2. We begin with a polynomial-space algorithm when the preorder is given as a Boolean circuit, and characterise classes of preorders for which PSPACE-hardness holds. We then consider preorders outside those classes and establish the complexity of the associated problems.

**Table 2.** Summary of the results for safety

	Preorder	Value problem	(Constr.) existence
	Conjunction	P-c	NP-c
	Subset	P-c	PSPACE-c
	Disjunction, Parity	PSPACE-c	PSPACE-c
	Counting, Maximise, Lexicographic	PSPACE-c	PSPACE-c
	(Monotonic) Boolean Circuit	PSPACE-c	PSPACE-c

**Theorem 9.** *For safety objectives with preorders given as Boolean circuits, the value, existence and constrained existence problems are in PSPACE. For preorders having  $\mathbf{0}$  as a unique minimal element, the existence and constrained existence problems are PSPACE-complete, even for two players. If additionally there is a vector  $v \in \{0, 1\}^n$  satisfying the equivalence  $v \not\leq v' \Leftrightarrow v' = \mathbf{0}$ , then the value problem is PSPACE-complete.*

*Proof.* In the most general case (Boolean circuits), safety objectives are dual to reachability objectives, hence the PSPACE algorithm.

The hardness proof for the existence problem and preorders with a unique minimal element uses the same arguments as in the proof of Theorem 6. We need to insert a matching-pennies game at the beginning however, because we are interested in Nash equilibria here. Hardness for the value problem is obtained by dualizing the result of Theorem 6 for conjunctions of reachability objectives.

Disjunction, counting, maximise, and lexicographic preorders are examples of preorders that satisfy this condition, and have a unique minimal element. The subset preorder also has a unique element.  $\square$

Note that the hardness results extends to parity, as it can encode disjunction.

We now consider simpler cases. As for reachability, the simple cases are for the preference relations that are *reducible* or *co-reducible* to a single safety objective. For a (non-trivial) preorder reducible to a single safety objective, the value problem retains the same complexity as in the single objective case, namely P-completeness. In the same way, for a (non-trivial) preorder co-reducible to a single safety objective, the existence and constrained existence problems remain NP-complete. The conjunction order is reducible and co-reducible to a single safety objective. The subset preorder is reducible to a single safety objective (but not co-reducible).

## 5 Büchi Objectives

We now turn to Büchi objectives, for which we prove the results listed in Table 3. (For the definition of the class  $P_{\parallel}^{NP}$ , see [13, Chapter 17].)

**Table 3.** Summary of the results for Büchi objectives

	Preorder	Value	Existence	Constr. exist.
Maximise, Disjunction, Subset		P-c	P-c	P-c
Conjunction, Lexicographic		P-c	P-h, in NP	NP-c
Counting		coNP-c	NP-c	NP-c
Monotonic Boolean Circuit		coNP-c	NP-c	NP-c
Parity		$UP \cap coUP$ [8]	coNP-h [4], in $P_{\parallel}^{NP}$	$P_{\parallel}^{NP}$ -c
Boolean Circuit		PSPACE-c	PSPACE-c	PSPACE-c

### 5.1 Reduction to Zero-Sum Games

In this section, we show how, from a multiplayer game  $\mathcal{G}$ , we can construct a two-player game  $\mathcal{H}$ , such that there is a correspondence between Nash equilibria in  $\mathcal{G}$  and certain winning strategies in  $\mathcal{H}$ . This allows us to reuse algorithmic techniques for zero-sum games to solve our problems.

We begin with introducing a few extra definitions. We say that a strategy profile  $\sigma_{Agt}$  is a *trigger strategy* for payoff  $(v^A)_{A \in Agt}$  from state  $s$  if for any strategy  $\sigma'_A$  of any player  $A \in Agt$ , the outcome  $\rho$  of  $\sigma_{Agt}[A \mapsto \sigma'_A]$  from  $s$  satisfies  $\mathbf{1}_{\{i|\rho \in \Omega_i^A\}} \lesssim v^A$ .

*Remark 10.* A Nash equilibrium is a trigger strategy for the payoff of its outcome. Reciprocally, if the outcome of  $\sigma_{Agt}$  has payoff  $(v^A)_{A \in Agt}$  and  $\sigma_{Agt}$  is a trigger strategy for  $(v^A)_{A \in Agt}$ , then  $\sigma_{Agt}$  is a Nash equilibrium.

Given two states  $s$  and  $s'$ , and a move  $m_{Agt}$ , the set of *suspect players* [2] for  $(s, s')$  and  $m_{Agt}$ , denoted with  $Susp((s, s'), m_{Agt})$ , is the set

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

Intuitively, player  $A \in Agt$  is a suspect for transition  $(s, s')$  and move  $m_{Agt}$  if she can unilaterally change her action to trigger the transition to  $s'$ . Notice that if  $Tab(s, m_{Agt}) = s'$ , then  $Susp((s, s'), m_{Agt}) = Agt$ . Also notice that, given a strategy profile  $\sigma_{Agt}$ , player  $A$  is a suspect along all the transitions of a play  $\rho$  (i.e., for all index  $i$ , player  $A$  is in  $Susp((\rho_i, \rho_{i+1}), \sigma_{Agt}(\rho_{\leq i}))$ ) iff there is a strategy  $\sigma'_A$  such that  $Out(s, \sigma_{Agt}[A \mapsto \sigma'_A]) = \rho$ .

With a game  $\mathcal{G}$  and a payoff  $(v^A)_{A \in Agt}$ , we associate a two-player turn-based game  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in Agt})$ . The set  $V_1$  of states (*configurations*) of  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in Agt})$

controlled by player  $A_1$  is (a subset of)  $\text{States} \times 2^{\text{Agt}}$ , and the set  $V_2$  of configurations controlled by player  $A_2$  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  $V_1$ , player  $A_1$  chooses a legal move  $m_{\text{Agt}}$  from  $s$ ; the next configuration is  $(s, P, m_{\text{Agt}})$ , in which  $A_2$  choose some state  $s' \in \text{States}$ , and the new configuration is  $(s', P \cap \text{Susp}((s, s'), m_{\text{Agt}}))$ . In particular, when the state  $s'$  chosen by player  $A_2$  satisfies  $s' = \text{Tab}(s, m_{\text{Agt}})$  (we say that  $A_2$  obeys  $A_1$ ), then the new configuration is  $(s', P)$ .

We define projections  $\pi_1$  and  $\pi_2$  from  $V_1$  on  $\text{States}$  and  $2^{\text{Agt}}$ , respectively, in the natural way. We extend these projections to plays, but only using player  $A_1$  states to avoid stutter, by setting  $\pi_1((s_0, P_0)(s_0, P_0, m_0)(s_1, P_1) \dots) = s_0 s_1 \dots$ . For any run  $\rho$ ,  $\pi_2(\rho)$  (seen as a sequence of sets of players) is decreasing, therefore its limit  $L(\rho)$  is well defined. An outcome  $\rho$  is *winning for player  $A_1$*  if, for all  $A \in L(\rho)$ ,  $\mathbf{1}_{\{i | \pi(\rho) \in \Omega_i^A\}} \lesssim v^A$ . In general, since each  $\Omega_i^A$  is a Büchi objective, the winning condition for  $A_1$  can be represented using a (possibly exponential-size) Muller condition. The *winning region* is the set of configurations  $(s, P)$  from which  $A_1$  has a winning strategy. Intuitively, player  $A_1$  tries to have the players play a Nash equilibrium, and player  $A_2$  tries to disprove that the played strategy profile is a Nash equilibrium, by finding a possible deviation that improves the payoff for one of the original players.

At first sight, the number of states in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  is exponential (in the number of players). However, since the transition table  $\text{Tab}$  is given explicitly [9], the size of  $\mathcal{G}$  is  $\sum_{s \in \text{States}} \prod_{A \in \text{Agt}} |\text{Mov}(s, A)|$ , and we have the following result:

**Lemma 11.** *The number of reachable configurations from  $\text{States} \times \{\text{Agt}\}$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  is polynomial in the size of  $\mathcal{G}$ .*

The next two lemmas state the correctness of our construction, establishing a correspondence between winning strategies in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  and Nash equilibria in  $\mathcal{G}$ .

**Lemma 12.** *Let  $(v^A)_{A \in \text{Agt}}$  be a payoff vector, and  $\rho$  be an infinite path in  $\mathcal{G}$ . The following two conditions are equivalent:*

- player  $A_1$  has a winning strategy in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  from  $(s, \text{Agt})$ , and its outcome  $\rho'$  from  $(s, \text{Agt})$  when  $A_2$  obeys  $A_1$  is such that  $\pi_1(\rho') = \rho$ ;
- there is a trigger strategy for  $(v^A)_{A \in \text{Agt}}$  in  $\mathcal{G}$  from state  $s$  whose outcome from  $s$  is  $\rho$ .

*Proof.* Assume there is a winning strategy  $\sigma^1$  for player  $A_1$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  from  $(s, \text{Agt})$ . We define the strategy profile  $\sigma_{\text{Agt}}$  according to the actions played by  $A_1$ . Pick a history  $g = s_1 s_2 \dots s_{k+1}$  with  $s_1 = s$ . Let  $h$  be the outcome of  $\sigma^1$  from  $s$  ending in a player  $A_1$  state and such that  $\pi_1(h) = s_1 \dots s_k$ . This history is uniquely defined as follows: the first state of  $h$  is  $(s_1, \text{Agt})$ , and if its  $(2i + 1)$ -th state is  $(s_i, P_i)$ , then its  $(2i + 2)$ -th state is  $(s_i, P_i, \sigma^1(h_{\leq 2i+1}))$  and its  $(2i + 3)$ -th state is  $(s_{i+1}, P_i \cap \text{Susp}((s_i, s_{i+1}), \sigma^1(h_{\leq 2i+1})))$ . Now, write  $(s_k, P_k)$  for the last state of  $h$ , and let  $h' = h \cdot (s_k, P_k, \sigma^1(h)) \cdot (s_{k+1}, P_k \cap \text{Susp}((s_k, s_{k+1}), \sigma^1(h)))$ . Then we define  $\sigma_{\text{Agt}}(g) = \sigma^1(h')$ . Notice that when  $g \cdot s$  is a prefix of  $\pi_1(\rho')$

(where  $\rho'$  is the outcome of  $\sigma^1$  from  $s$  when  $A_2$  obeys  $A_1$ ), then  $g \cdot s \cdot \sigma_{\text{Agt}}(g \cdot s)$  is also a prefix of  $\pi_1(\rho')$ .

We now prove that  $\sigma_{\text{Agt}}$  is a trigger strategy for  $(v^A)_{A \in \text{Agt}}$ . Pick a player  $A \in \text{Agt}$ , a strategy  $\sigma'_A$  for  $A$ , and an infinite play  $g$  in  $\text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A])$ . With  $g$ , we associate an infinite play  $h$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  in the same way as above. Then player  $A$  is a suspect along all the transitions of  $g$ , so that she belongs to  $L(h)$ . Now, as  $\sigma^1$  is winning, the payoff for  $A$  of  $g = \pi_1(h)$  is less than  $v^A$ , which proves that  $\sigma_{\text{Agt}}$  is a trigger strategy.

Conversely, assume that  $\sigma_{\text{Agt}}$  is a trigger strategy for  $(v^A)_{A \in \text{Agt}}$ , and define the strategy  $\sigma^1$  by  $\sigma^1(h) = \sigma_{\text{Agt}}(\pi_1(h))$ . Notice that the outcome  $\rho'$  of  $\sigma^1$  when  $A_2$  obeys  $A_1$  satisfies  $\pi_1(\rho') = \rho$ .

Let  $\eta$  be an outcome of  $\sigma^1$  from  $s$ , and  $A \in L(\eta)$ . Then  $A$  is a suspect for each transition along  $\pi_1(\eta)$ , which means that for all  $i$  there is a move  $m_i^A$  such that  $\pi_1(\eta)_{=i+1} = \text{Tab}(\pi_1(\eta)_{=i}, \sigma_{\text{Agt}}(\pi_1(\eta)_{\leq i})[A \mapsto m_i^A])$ . Therefore there is a strategy  $\sigma'_A$  such that  $\pi_1(\eta) = \text{Out}(s, \sigma_{\text{Agt}}[A \mapsto \sigma'_A])$ . Since  $\sigma_{\text{Agt}}$  is a trigger strategy for  $(v^A)_{A \in \text{Agt}}$ , the payoff for player  $A$  of  $\pi_1(\eta)$  is less than  $v_A$ . As this holds for any  $A \in L(\eta)$ ,  $\sigma^1$  is winning.  $\square$

**Lemma 13.** *Let  $\rho$  be an infinite path in  $\mathcal{G}$  with payoff  $(v^A)_{A \in \text{Agt}}$ . The following two conditions are equivalent:*

- *there is a path  $\rho'$  from  $(s, \text{Agt})$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  that never leaves the winning region of  $A_1$  and along which  $A_2$  obeys  $A_1$ , such that  $\pi_1(\rho') = \rho$ ;*
- *there is a Nash equilibrium  $\sigma_{\text{Agt}}$  from  $s$  in  $\mathcal{G}$  whose outcome is  $\rho$ .*

*Proof.* Let  $\rho$  be a path in the winning region of  $A_1$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$ . We define a strategy  $\sigma^1$  that follows  $\rho$  when  $A_2$  obeys. Along  $\rho$ , this strategy is defined as follows:  $\sigma^1(\rho_{\leq 2i}) = m_{\text{Agt}}$  such that  $\text{Tab}(\pi_1(\rho)_{=i}, m_{\text{Agt}}) = \pi_1(\rho)_{=i+1}$ . Such a legal move must exist since  $A_2$  obeys  $A_1$  along  $\rho$ . Now, if player  $A_2$  deviates from the obeying strategy, we make  $\sigma^1$  follow a winning strategy of  $A_1$ : given a finite outcome  $\eta$  of  $\sigma^1$  that is not a prefix of  $\rho$ , we let  $j$  be the largest index such that  $\eta_{\leq j}$  is a prefix of  $\rho$ . In particular,  $\eta_{=j}$  belongs to the winning region  $W$  of  $A_1$ , and belongs to player  $A_2$  (otherwise  $\eta_{\leq j+1}$  would also be a prefix of  $\rho$ ). Hence, all the successors of  $\eta_{=j}$  are in  $W$ . Thus player  $A_1$  has a winning strategy  $\hat{\sigma}^1$  from  $\eta_{=j+1}$ . We then define  $\sigma^1(\eta_{\leq j} \cdot \eta') = \hat{\sigma}^1(\eta')$  for any outcome  $\eta'$  of  $\hat{\sigma}^1$  from  $\eta_{=j+1}$ .

Each outcome of  $\sigma^1$  is either the path  $\rho$  or a path that, from some point onwards, is compatible with a winning strategy. Since  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$  has a Muller winning condition, it follows that  $\sigma^1$  is winning. Applying Lemma 12, we obtain a strategy profile  $\sigma_{\text{Agt}}$  in  $\mathcal{G}$  that is a trigger strategy for  $(v^A)_{A \in \text{Agt}}$ . Moreover, the outcome of  $\sigma_{\text{Agt}}$  from  $s$  equals  $\pi_1(\rho)$ , so that  $\sigma_{\text{Agt}}$  is a Nash equilibrium.

Conversely, the Nash equilibrium is a trigger strategy, and from Lemma 12, we get a winning strategy  $\sigma^1$  in  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$ . The outcome  $\rho$  of  $\sigma^1$  from  $s$  when  $A_2$  obeys  $A_1$  is such that  $\pi_1(\rho)$  is the outcome of the Nash equilibrium, so that its payoff is  $(v^A)_{A \in \text{Agt}}$ . Since  $\sigma^1$  is winning,  $\rho$  never leaves the winning region, which concludes the proof.  $\square$

## 5.2 Applications of the Reduction

**General Case.** As noticed in [7], the algorithm from [10] to find the winning states in a game can be adapted to the case where the winning conditions are given as a Boolean circuit (the circuit has as many input gates as the number of states, and a path is declared winning if the circuit evaluates to 1 when setting the input gates to 1 for the states that are visited infinitely often). It uses polynomial space. Using such an algorithm we get the following result.

**Proposition 14.** *For Büchi objectives with preorders given as Boolean circuits, the value, existence and constrained existence problems are PSPACE-complete.*

**Reduction to a Single Büchi Objective.** The preorders that were reducible to a single reachability objectives in the case of reachability can also be reduced to a single Büchi objective in the Büchi case: just replace Occ with Inf. The same holds for co-reducibility. The algorithm from [4] can then be adapted.

**Proposition 15.** *For Büchi objectives with a monotonic preorder reducible to a single objective, the value problem is P-complete. For Büchi objectives with a preorder co-reducible to a single objective, the existence and constrained existence problems are P-complete.*

**Reduction to a Deterministic Büchi Automaton.** For some preorders, given any payoff  $u$ , it is possible to construct (in polynomial time) a deterministic Büchi automaton that recognises the plays whose payoff  $v$  for player  $A$  is higher than  $u$  (i.e.  $u \lesssim v$ ). When this is the case, we say that the preorder is *reducible to a deterministic Büchi automaton*.

**Proposition 16.** *For Büchi objectives and a preorder reducible to a deterministic Büchi automaton, the value problem is in P. In particular, it is P-complete for conjunction, lexicographic and subset preorders.*

The idea of the algorithm is to compute the product of the game with the Büchi automaton to which the given payoff  $v^A$  reduces. Notice that *reachability* objectives with the parity order are also reducible to a deterministic Büchi automaton; we thus recover the complexity result about weak parity games from [12].

**Monotonic Preorders.** When the preorder is monotonic, our problems are also simpler than in the general case. This is because we can find suspect-based trigger strategies, corresponding to memoryless strategies in the game  $\mathcal{H}(\mathcal{G}, (v^A)_{A \in \text{Agt}})$ .

**Proposition 17.** *For Büchi objectives with preorders given by a monotonic circuits, the value problem is coNP-complete, and the existence and constrained existence problem are NP-complete. For the counting order, the value problem is coNP-complete, and existence and constrained existence are NP-complete. For monotonic preorders with an element  $v$  such that  $u \not\lesssim v \Leftrightarrow u = \mathbf{1}$ , the constrained existence problem is NP-complete.*

**Parity Games.** Finally, for Büchi objectives with the parity preorder, we have:

**Proposition 18.** *For Büchi objectives with the parity preorder, the constrained existence problem is  $\text{P}_{\parallel}^{\text{NP}}$ -complete.*

## 6 Conclusion

We have contributed to the algorithmics of Nash equilibria computation in concurrent games with ordered objectives. We believe the game abstraction proposed in Section 5.1 can be used in other contexts, which we are currently investigating. The algorithms presented in this paper have partly been implemented in the tool PRALINE (<http://www.lsv.ens-cachan.fr/Software/praline/>).

## References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *J. ACM* 49(5), 672–713 (2002)
2. Bouyer, P., Brenguier, R., Markey, N.: Nash Equilibria for Reachability Objectives in Multi-player Timed Games. In: Gastin, P., Laroussinie, F. (eds.) *CONCUR 2010. LNCS*, vol. 6269, pp. 192–206. Springer, Heidelberg (2010)
3. Bouyer, P., Brenguier, R., Markey, N., Ummels, M.: Concurrent games with ordered objectives. Research Report LSV-11-22, LSV, ENS Cachan, France (2011)
4. Bouyer, P., Brenguier, R., Markey, N., Ummels, M.: Nash equilibria in concurrent games with Büchi objectives. In: *FSTTCS 2011. LIPIcs*, vol. 13, pp. 375–386. Leibniz-Zentrum für Informatik, LZI (2011)
5. Chatterjee, K., Majumdar, R., Jurdziński, M.: On Nash Equilibria in Stochastic Games. In: Marcinkowski, J., Tarlecki, A. (eds.) *CSL 2004. LNCS*, vol. 3210, pp. 26–40. Springer, Heidelberg (2004)
6. Henzinger, T.A.: Games in system design and verification. In: *TARK 2005*, pp. 1–4 (2005)
7. Hunter, P., Dawar, A.: Complexity Bounds for Regular Games. In: Jędrzejowicz, J., Szepietowski, A. (eds.) *MFCS 2005. LNCS*, vol. 3618, pp. 495–506. Springer, Heidelberg (2005)
8. Jurdzinski, M.: Deciding the winner in parity games is in  $UP \cap co-UP$ . *Inf. Process. Lett.* 68(3), 119–124 (1998)
9. Laroussinie, F., Markey, N., Oreiby, G.: On the expressiveness and complexity of ATL. *Logical Methods in Computer Science* 4(2) (2008)
10. McNaughton, R.: Infinite games played on finite graphs. *Annals of Pure and Applied Logic* 65(2), 149–184 (1993)
11. Nash Jr., J.F.: Equilibrium points in  $n$ -person games. *Proc. National Academy of Sciences of the USA* 36(1), 48–49 (1950)
12. Neumann, J., Szepietowski, A., Walukiewicz, I.: Complexity of weak acceptance conditions in tree automata. *Inf. Process. Lett.* 84(4), 181–187 (2002)
13. Papadimitriou, C.H.: *Complexity Theory*. Addison-Wesley (1994)
14. Paul, S., Simon, S.: Nash equilibrium in generalised Muller games. In: *FSTTCS 2009. LIPIcs*, vol. 4, pp. 335–346. LZI (2010)
15. Thomas, W.: Infinite Games and Verification (Extended abstract of a tutorial). In: Brinksma, E., Larsen, K.G. (eds.) *CAV 2002. LNCS*, vol. 2404, pp. 58–64. Springer, Heidelberg (2002)
16. Ummels, M.: The Complexity of Nash Equilibria in Infinite Multiplayer Games. In: Amadio, R.M. (ed.) *FOSSACS 2008. LNCS*, vol. 4962, pp. 20–34. Springer, Heidelberg (2008)
17. Ummels, M., Wojtczak, D.: The Complexity of Nash Equilibria in Limit-Average Games. In: Katoen, J.-P., König, B. (eds.) *CONCUR 2011 – Concurrency Theory. LNCS*, vol. 6901, pp. 482–496. Springer, Heidelberg (2011)