

# The Complexity of LTL Rational Synthesis

Orna Kupferman and Noam Shenwald

School of Computer Science and Engineering, The Hebrew University, Israel

**Abstract.** In *rational synthesis*, we automatically construct a reactive system that satisfies its specification in all rational environments, namely environments that have objectives and act to fulfill them. We complete the study of the complexity of LTL rational synthesis. Our contribution is threefold. First, we tighten the known upper bounds for settings that were left open in earlier work. Second, our complexity analysis is parametric, and we describe tight upper and lower bounds in each of the problem parameters: the game graph, the objectives of the system components, and the objectives of the environment components. Third, we generalize the definition of rational synthesis, combining the cooperative and non-cooperative approaches studied in earlier work, and extend our complexity analysis to the general definition.

## 1 Introduction

*Synthesis* is the automated construction of a system from its specification. The basic idea is simple and appealing: instead of developing a system and verifying that it adheres to its specification, we use an automated procedure that, given a specification, constructs a system that is correct by construction, thus enabling the designers to focus on *what* the system should do rather than *how* to do it. A reactive system interacts with its environment and should satisfy its specification in all environments [9, 26]. Accordingly, synthesis corresponds to a *zero-sum game* between the system and the environment, where they together generate a computation, the system wins if the computation satisfies the specification, and otherwise, the environment wins.

In practice, the requirement to satisfy the specification in all environments is often too strong. Therefore, it is common to add assumptions on the environment. An assumption may be direct, say a specification that restricts the possible behaviors of the environment [6], or less direct, say a bound on the size of the environment or other resources it uses [15]. In [12], the authors suggest a conceptual assumption on the environment, namely its rationality: *Rational synthesis* is based on the idea that the components composing the environment typically have objectives of their own, and they act to achieve their objectives. For example, clients interacting with a server typically have objectives other than to fail the server. As shown in [12], the system can capitalize on the rationality and objectives of components that compose its environment. Adding rationality into the picture makes the corresponding game *non-zero-sum* [23], thus objectives of different players may overlap.

The interesting questions about non-zero-sum games concern *stable outcomes*, in particular *Nash equilibria* (NE) [22]. More formally, each of the players in the game has a *strategy* that directs her which actions to take; a *profile* is a vector of strategies, one for each player; each profile has an *outcome* (in our case, the computation generated when the system and the environment follow their strategies); and a profile is an NE if no player has an incentive to deviate from it (in our case, to change her strategy in a way that would cause the outcome of the new profile to satisfy her objective).

Two approaches to rational synthesis have been studied. In *cooperative* rational synthesis (CRS) [12], the desired output is an NE profile whose outcome satisfies the objective of the system. Thus, in CRS, we assume that we can suggest strategies to the environment players, and once they have no incentive to deviate from these strategies, they follow them. Then, in *non-cooperative* rational synthesis (NRS) [16], the desired output is a strategy for the system player such that the objective of the system is satisfied in the outcome of all NE profiles that include this strategy. Thus, in NRS, the environment players are rational, but we cannot suggest them a strategy.

The cooperative and non-cooperative approaches correspond to different settings in reality, having to do both with the technical ability to communicate a strategy to the environment players, say due to different architectures, as well as the willingness of the environment players to follow a suggested strategy. As shown in [1], the two approaches are related to the two stability-inefficiency measures of *price of stability* [3] and *price of anarchy* [17, 24]. Additional related work includes *rational verification* [28, 13], where we check that a given system satisfies its specification when interacting with a rational environment, and extensions of rational synthesis to richer settings (multi-valued, partial visibility, and more) [4, 14, 19].

The *complexity* of rational synthesis was first studied for the case the input to the problem is the objectives of the players, given by LTL formulas. In this setting, CRS is in 2EXPTIME [12], whereas the best known upper bound for NRS until recently was 3EXPTIME [16] (the paper specifies a 2EXPTIME upper bound, but a careful analysis of the algorithm reveals that it is actually in 3EXPTIME), improved to 2EXPTIME for *turn-based* games with two players [19]. The complexity analysis above suggests that rational synthesis is not harder than traditional synthesis. One may wonder whether this has to do with the doubly-exponential translation of LTL to deterministic automata, which dominates the complexity. To answer this question, [10] studies the complexity of rational synthesis where the objectives of the players are given by  $\omega$ -regular winning conditions in a game graph (e.g., reachability, Büchi, and parity). The analysis in [10] also distinguishes between the case the number of players is fixed and the case it is not. As shown there, in most cases the complexity of the rational variant coincides with the complexity of the zero-sum game. In some cases, however, it does not. For example, while the problem of deciding Rabin games is NP-complete [11], the best algorithm for solving CRS with Rabin objectives

is in  $P^{NP}$ , going up to PSPACE-complete in NRS, and going higher when the number of players is not fixed [10].

In this work, we complete the study of the complexity of LTL rational synthesis. Our contribution is threefold. First, we tighten the known upper bound for NRS for settings with three or more players and for *concurrent* games, which were left open in [10, 19]. Second, our complexity analysis is *parametric*, and we describe tight upper and lower bounds in each of the problem parameters: the game graph, the objectives of the system players, and the objectives of the environment players. Third, we generalize the definition of rational synthesis, combining the cooperative and non-cooperative approaches, and extend our complexity analysis to the general definition. Below we elaborate on each of the contributions.

Let us start with the generalization of the problem. In our general definition, we may suggest a strategy only to a subset of the environment players. Thus, we distinguish between three types of players: controllable, cooperative uncontrollable, and non-cooperative uncontrollable. Then, in the (general) rational-synthesis (RS) problem, we are given a labeled graph and LTL formulas that specify the objectives of the players, and we seek strategies for the controllable and the cooperative-uncontrollable players such that the objectives of the controllable players are satisfied in the outcome of every NE profile that extends these strategies. Note that CRS and NRS can be viewed as special cases of RS where the uncontrollable players are all cooperative or all non-cooperative.

In the tight-complexity front, our algorithms reduce rational synthesis to the nonemptiness problem of tree automata. The automata accept *certified strategy trees*: trees that are labeled by both a strategy for the controllable player<sup>1</sup> and information about uncontrollable players that deviate and the strategies to which they deviate. The most technically-challenging algorithm we describe is for NRS in the concurrent setting. While in the turn-based setting, we need a single player that deviates in order to justify a path in which the objective of the controllable player is not satisfied, in the concurrent setting, where the players choose actions simultaneously and independently, we need to consider sets of uncontrollable players. This makes the certificate much more complex. In particular, it involves labels from an exponential alphabet, which introduces an additional challenge, namely a need to decompose labels along branches in the tree. Also, while in the turn-based setting, an NE always exists, concurrent games with three or more players need not have an NE [8], and so a certified strategy tree should also certify the existence of an NE.

Finally, in the parameterized-complexity front, the fact our algorithms use tree automata (rather than a translation to Strategy Logic [7], which has been the case in [12, 16]), enables us to analyze the complexity in each of the parameters of the problem: the game graph  $G$ , the objective  $\psi_1$  of the controllable player, and the objectives  $\psi_2, \dots, \psi_k$  of the uncontrollable players. For CRS, [19] studies the

---

<sup>1</sup> It is easy to see that several controllable components can be merged to a single one.

parameterized complexity in turn-based games with two players.<sup>2</sup> The algorithm there is based on a distinction between the case the uncontrollable player satisfies her objective and the case she does not. Generalizing this to an arbitrary number of players, we parameterize solutions with the set of the uncontrollable players whose objectives are satisfied, and give a uniform solution to all cases. This also enables us to seek solutions that favor some or all uncontrollable players.

We show that the complexity of CRS is polynomial in  $|G|$ , doubly-exponential in  $|\psi_2|, \dots, |\psi_k|$ , and only exponential in  $|\psi_1|$ . Thus, in terms of the system specification, CRS is in fact easier than traditional synthesis! Once we move to NRS or RS, the complexity becomes doubly exponential in all objectives. We describe tight lower bounds for the different parameters, and we show that they are valid already for the case  $k = 2$  and the game is turn based. Specifically, we prove that CRS is EXPTIME-hard even when  $G$  and  $\psi_2$  are fixed, and is 2EXPTIME-hard even when  $G$  and  $\psi_1$  are fixed. Similarly, NRS is 2EXPTIME-hard even when only one of  $\psi_1$  and  $\psi_2$  is not fixed. In order to see the technical challenge in our lower-bound proofs, consider the current 2EXPTIME lower-bound proof for CRS, where synthesis of an objective  $\psi$  for the system is reduced to CRS with objectives  $\psi$  for the system and  $\neg\psi$  for the environment. The reduction crucially depends on both objectives not being fixed, and just changing either of them to *True* or *False* does not do the trick. In order to get 2EXPTIME-hardness in  $|\psi_2|$ , we need to cleverly manipulate both  $G$  and  $\psi_1$ .

Together, our results complete the complexity picture for a generalized definition of rational synthesis, for both turn-based and concurrent systems, with any number of components, and with the exact dependencies in each of the parameters of the problem.

Due to the lack of space, some proofs are omitted and can be found in the full version, in the authors' URLs.

## 2 Preliminaries

### 2.1 LTL, trees, and automata

The logic LTL is used for specifying on-going behaviors of reactive systems [25]. Formulas of LTL are constructed from a set  $AP$  of atomic propositions using the usual Boolean operators and the temporal operators  $G$  (“always”) and  $F$  (“eventually”),  $X$  (“next time”) and  $U$  (“until”). The semantics of LTL is defined with respect to infinite computations in  $(2^{AP})^\omega$ . We are going to use LTL for specifying the objectives of the system and the components composing the environment.

Given a set  $D$  of directions, a  $D$ -tree is a set  $T \subseteq D^*$  such that if  $x \cdot d \in T$ , where  $x \in D^*$  and  $d \in D$ , then also  $x \in T$ . The elements of  $T$  are called *nodes*, and the empty word  $\varepsilon$  is the *root* of  $T$ . For every  $x \in T$ , the nodes  $x \cdot d$ , for

<sup>2</sup> The study in [19] considers *perspective games* [20], which adds the challenge of *partial visibility* on top of rational synthesis, but the results there imply the desired bounds for the case of full visibility.

$d \in D$ , are the *successors* of  $x$ , and the *direction* of node  $x \cdot d$  is  $d$ . A *path*  $h$  in a tree  $T$  is a set  $h \subseteq T$  such that  $\varepsilon \in h$  and for every  $x \in h$ , either  $x$  is a leaf or there exists a unique  $d \in D$  such that  $x \cdot d \in h$ . We sometimes refer to paths in  $T$  as words in  $D^*$  or  $D^\omega$ . For a finite path  $h \subseteq D^*$  and a finite or infinite path  $h' \subseteq D^*$ , we use  $h \preceq h'$  to indicate that  $h$  is a prefix of  $h'$ , thus  $h \subseteq h'$ . Given an alphabet  $\Sigma$ , a  $\Sigma$ -labeled  $D$ -tree is a pair  $\langle T, \tau \rangle$  where  $T$  is a tree and  $\tau : T \rightarrow \Sigma$  maps each node of  $T$  to a letter in  $\Sigma$ .

Our algorithms use *automata on infinite words and trees*. We are going to use nondeterministic and universal automata, yet define below alternating automata, which subsume both classes. For a set  $X$ , let  $\mathcal{B}^+(X)$  be the set of positive Boolean formulas over  $X$  (i.e., Boolean formulas built from elements in  $X$  using  $\wedge$  and  $\vee$ ), where we also allow the formulas **true** and **false**. For a set  $Y \subseteq X$  and a formula  $\theta \in \mathcal{B}^+(X)$ , we say that  $Y$  *satisfies*  $\theta$  iff assigning **true** to elements in  $Y$  and assigning **false** to elements in  $X \setminus Y$  makes  $\theta$  true. An *alternating tree automaton* is  $\mathcal{A} = \langle \Sigma, D, Q, q_{in}, \delta, \alpha \rangle$ , where  $\Sigma$  is the input alphabet,  $D$  is a set of directions,  $Q$  is a finite set of states,  $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(D \times Q)$  is a transition function,  $q_{in} \in Q$  is an initial state, and  $\alpha \subseteq Q$  specifies a Büchi or a co-Büchi acceptance condition. For a state  $q \in Q$ , we use  $\mathcal{A}^q$  to denote the automaton obtained from  $\mathcal{A}$  by setting the initial state to be  $q$ . The *size* of  $\mathcal{A}$ , denoted  $|\mathcal{A}|$ , is the sum of lengths of formulas that appear in  $\delta$ .

The alternating automaton  $\mathcal{A}$  runs on  $\Sigma$ -labeled  $D$ -trees. A *run* of  $\mathcal{A}$  over a  $\Sigma$ -labeled  $D$ -tree  $\langle T, \tau \rangle$  is a  $(T \times Q)$ -labeled  $\mathbb{N}$ -tree  $\langle T_r, r \rangle$ . Each node of  $T_r$  corresponds to a node of  $T$ . A node in  $T_r$ , labeled by  $(x, q)$ , describes a copy of the automaton that reads the node  $x$  of  $T$  and visits the state  $q$ . Note that many nodes of  $T_r$  can correspond to the same node of  $T$ . The labels of a node and its successors have to satisfy the transition function. Formally,  $\langle T_r, r \rangle$  satisfies the following:

1. (1)  $\varepsilon \in T_r$  and  $r(\varepsilon) = \langle \varepsilon, q_{in} \rangle$ .
2. (2) Let  $y \in T_r$  with  $r(y) = \langle x, q \rangle$  and  $\delta(q, \tau(x)) = \theta$ . Then there is a (possibly empty) set  $S = \{(c_0, q_0), (c_1, q_1), \dots, (c_{n-1}, q_{n-1})\} \subseteq D \times Q$ , such that  $S$  satisfies  $\theta$ , and for all  $0 \leq i \leq n-1$ , we have  $y \cdot i \in T_r$  and  $r(y \cdot i) = \langle x \cdot c_i, q_i \rangle$ .

For example, if  $\langle T, \tau \rangle$  is a  $\{0, 1\}$ -tree with  $\tau(\varepsilon) = a$  and  $\delta(q_{in}, a) = ((0, q_1) \vee (0, q_2)) \wedge ((0, q_3) \vee (1, q_2))$ , then, at level 1, the run  $\langle T_r, r \rangle$  includes a node labeled  $(0, q_1)$  or a node labeled  $(0, q_2)$ , and includes a node labeled  $(0, q_3)$  or a node labeled  $(1, q_2)$ . Note that if, for some  $y$ , the transition function  $\delta$  has the value **true**, then  $y$  need not have successors. Also,  $\delta$  can never have the value **false** in a run.

A run  $\langle T_r, r \rangle$  is accepting if all its infinite paths satisfy the acceptance condition. Given a run  $\langle T_r, r \rangle$  and an infinite path  $\pi \subseteq T_r$ , let  $inf(\pi) \subseteq Q$  be such that  $q \in inf(\pi)$  if and only if there are infinitely many  $y \in \pi$  for which  $r(y) \in T \times \{q\}$ . That is,  $inf(\pi)$  contains exactly all the states that appear infinitely often in  $\pi$ . A path  $\pi$  satisfies a *Büchi* acceptance condition  $\alpha$  iff  $inf(\pi) \cap \alpha \neq \emptyset$ , and satisfies a *co-Büchi* acceptance condition  $\alpha$  iff  $inf(\pi) \cap \alpha = \emptyset$ . We also consider the *parity* acceptance condition, where  $\alpha : Q \rightarrow \{0, 1, \dots, k\}$  maps each state to a color in  $\{0, 1, \dots, k\}$ , and a path  $\pi$  satisfies  $\alpha$  if the minimal color visited infinitely often

is even, thus  $\min\{i : \text{inf}(\pi) \cap \alpha^{-1}(i) \neq \emptyset\}$  is even. An automaton accepts a tree iff there exists a run that accepts it. We denote by  $L(\mathcal{A})$  the set of all  $\Sigma$ -labeled trees that  $\mathcal{A}$  accepts. The *size* of  $\mathcal{A}$ , denoted  $|\mathcal{A}|$ , is the sum of lengths of the description of its transition function.

The alternating automaton  $\mathcal{A}$  is *nondeterministic* if for all the formulas that appear in  $\delta$ , if  $(c_1, q_1)$  and  $(c_2, q_2)$  are conjunctively related, then  $c_1 \neq c_2$ . (i.e., if the transition is rewritten in disjunctive normal form, there is at most one element of  $\{c\} \times Q$ , for each  $c \in D$ , in each disjunct). Note that then, the run tree  $T_r$  is equal to  $T$ , and the  $\Sigma$ -labels in  $\tau$  are replaced by  $Q$ -labels in  $r$ . The automaton  $\mathcal{A}$  is *universal* if all the formulas that appear in  $\delta$  are conjunctions of atoms in  $D \times Q$ . Note that then, there is only one run tree of  $\mathcal{A}$  on  $\langle T, \tau \rangle$ , yet each node  $x \in T$  may have several nodes  $y \in T_r$  such that  $r(y) = \langle x, q \rangle$  for some  $q \in Q$ . Finally,  $\mathcal{A}$  is *deterministic* if it is both nondeterministic and universal. The automaton  $\mathcal{A}$  is a *word* automaton if  $|D| = 1$ . Then, we can omit  $D$  from the specification of the automaton and denote the transition function of  $\mathcal{A}$  as  $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(Q)$ . If the word automaton is nondeterministic or universal, then  $\delta : Q \times \Sigma \rightarrow 2^Q$ .

We denote different types of automata by three-letter acronyms in  $\{D, N, U\} \times \{F, B, C, P\} \times \{W, T\}$ , where the first letter describes the branching mode of the automaton (deterministic, nondeterministic, or universal), the second letter describes the acceptance condition (finite, Büchi, co-Büchi, or parity), and the third letter describes the object over which the automaton runs (words or trees). For example, UCT stands for a universal co-Büchi tree automaton.

## 2.2 Concurrent multiplayer games

For  $k \geq 1$ , let  $[k] = \{1, \dots, k\}$ . A *k-player game graph* is a tuple  $G = \langle AP, V, v_0, \{A_i\}_{i \in [k]}, \{\kappa_i\}_{i \in [k]}, \delta, \tau \rangle$ , where  $AP$  is a set of atomic propositions,  $V$  is a set of vertices,  $v_0 \in V$  is an initial vertex, and for  $i \in [k]$ , the set  $A_i$  is a set of actions of Player  $i$ , and  $\kappa_i : V \rightarrow 2^{A_i}$  specifies the set of actions that Player  $i$  can take at each vertex.

A *move* in  $G$  is a tuple  $\langle a_1, \dots, a_k \rangle \in A_1 \times \dots \times A_k$ , describing possible choices of actions for all  $k$  players. A move  $\langle a_1, \dots, a_k \rangle$  is *possible* for vertex  $v \in V$  if  $a_i \in \kappa_i(v)$  for all  $i \in [k]$ . Then, the transition function  $\delta : V \times A_1 \times \dots \times A_k \rightarrow V$  is a deterministic function that maps each vertex and possible move for it to a successor vertex. Finally, the function  $\tau : V \rightarrow 2^{AP}$  maps each vertex to the set of atomic propositions that hold in it.

A *game* is a tuple  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ , where  $G$  is a  $k$ -player game graph, and  $\psi_i$ , for  $i \in [k]$ , is an LTL formula over  $AP$ , describing the *objective* of Player  $i$ . In a beginning of a play in the game, a token is placed on  $v_0$ . Then, at each round, the players choose actions simultaneously and independently of the other players, and the induced move determines the successor vertex. Repeating this, the players generate a *play*  $\rho = v_0, v_1, \dots$  in  $G$ , which induces the *computation*  $\tau(\rho) = \tau(v_0), \tau(v_1), \dots \in (2^{AP})^\omega$ . For every  $i \in [k]$ , Player  $i$  aims for a play whose computation satisfies  $\psi_i$ . For an LTL formula  $\psi$ , let  $L(\psi) \subseteq (2^{AP})^\omega$  be the set of computations that satisfy  $\psi$ .

A *strategy* for Player  $i$  is a function  $f_i : V^+ \rightarrow A_i$  that maps histories of the game to an action suggested to Player  $i$ . The suggestion has to be consistent with  $\kappa_i$ . Thus, for every  $v_0 v_1 \cdots v_j \in V^+$ , we have that  $f_i(v_0 v_1 \cdots v_j) \in \kappa_i(v_j)$ . A *profile* is a tuple  $\pi = \langle f_1, \dots, f_k \rangle$  of strategies, one for each player. The *outcome* of a profile  $\pi = \langle f_1, \dots, f_k \rangle$  is the play obtained when the players follow their strategies. Formally,  $\text{Outcome}(\pi) = v_0, v_1, \dots$  is such that for all  $j \geq 0$ , we have that  $v_{j+1} = \delta(v_j, \langle f_1(v_0 \cdots v_j), \dots, f_k(v_0 \cdots v_j) \rangle)$ . For a subset  $S \subseteq [k]$  of players, an *S-profile* is a set of strategies, one for each player in  $S$ . We say that a profile  $\pi$  *extends* an *S-profile*  $\pi'$  if the players in  $S$  use in  $\pi$  their strategies in  $\pi'$ .

Consider a profile  $\pi$ . The set of *winners* in  $\pi$ , denoted  $\text{Win}(\pi)$ , is the set of players whose objectives are satisfied in  $\text{Outcome}(\pi)$ . Formally,  $i \in \text{Win}(\pi)$  iff  $\tau(\text{Outcome}(\pi)) \in L(\psi_i)$ . The set of *losers* in  $\pi$ , denoted  $\text{Lose}(\pi)$ , is then  $[k] \setminus \text{Win}(\pi)$ , namely the set of players whose objectives are not satisfied in  $\text{Outcome}(\pi)$ .

A game  $\mathcal{G}$  is *zero-sum* if the objectives of the players form a partition of all possible behaviors. That is, for every  $i \neq j \in [k]$ , we have that  $L(\psi_i) \cap L(\psi_j) = \emptyset$ , and  $\bigcup_{i \in [k]} L(\psi_i) = (2^{AP})^\omega$ . Accordingly, for every profile  $\pi$  in a zero-sum game, we have that  $|\text{Win}(\pi)| = 1$  and  $|\text{Lose}(\pi)| = k - 1$ . We then say that Player  $i$  *wins*  $\mathcal{G}$  if she has a *winning strategy* – a strategy that guarantees the satisfaction of  $\psi_i$  no matter how the other players proceed. Formally,  $f_i$  is a winning strategy if for every profile  $\pi$  with  $f_i$ , we have that  $\text{Win}(\pi) = \{i\}$ .

Games may be *non zero-sum*, thus the objectives of the players may overlap. In such games, we are interested in *stable* profiles. In particular, a profile  $\pi = \langle f_1, \dots, f_k \rangle$  is a *Nash Equilibrium* (NE, for short) [22] if, intuitively, no (single) player can benefit from unilaterally changing her strategy. In our setting, benefiting amounts to moving from the set of losers to the set of winners. Formally, for  $i \in [k]$  and a strategy  $f'_i$  for Player  $i$ , let  $\pi[i \leftarrow f'_i] = \langle f_1, \dots, f_{i-1}, f'_i, f_{i+1}, \dots, f_k \rangle$  be the profile obtained from  $\pi$  by changing the strategy of Player  $i$  to  $f'_i$ . We say that  $\pi$  is an NE if for every  $i \in [k]$ , if  $i \in \text{Lose}(\pi)$ , then for every strategy  $f'_i$ , we have that  $i \in \text{Lose}(\pi[i \leftarrow f'_i])$ . Thus,  $\pi$  is an NE if no player has an incentive to deviate from  $\pi$ . For a subset  $W \subseteq [k]$  of players, we say that  $\pi$  is a *W-NE* if  $\pi$  is an NE with  $W = \text{Win}(\pi)$ .

The game  $\mathcal{G}$  is *turn-based* if the transition function of its graph  $G$  is such that for every vertex  $v \in V$ , there is a single player that “owns”  $v$  and determines the successor vertex whenever the play is in  $v$ . Formally, for every  $v \in V$ , there is  $i \in [k]$  such that for all moves  $\langle a_1, \dots, a_k \rangle$  and  $\langle a'_1, \dots, a'_k \rangle$  that are possible for  $v$ , if  $a_i = a'_i$ , then  $\delta(v, \langle a_1, \dots, a_k \rangle) = \delta(v, \langle a'_1, \dots, a'_k \rangle)$ . Accordingly, we describe the game graph of a turn-based game as  $G = \langle AP, \{V_i\}_{i \in [k]}, v_0, E, \tau \rangle$ , where  $V_1, \dots, V_k$  is a partition of  $V$  to the sets of vertices owned by the different players, and  $E \subseteq V \times V$  is the transition relation, modeling the fact that the set of actions of Player  $i$  in a vertex  $v$  she owned is the set of  $v$ 's successors.

### 3 Rational Synthesis

Consider a  $k$ -player game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . We distinguish between three types of players: A player is *controllable* if she is guaranteed to follow a strategy assigned to her. Otherwise, she is *uncontrollable*. The uncontrollable players are *rational* – they would not deviate from a profile unless they have a beneficial deviation from it. We distinguish between *cooperative uncontrollable* players, to which we can suggest a strategy (which they would follow unless they have a beneficial deviation), and *non-cooperative uncontrollable* players, to which we cannot suggest a strategy. The distinction between the cooperative and non-cooperative uncontrollable players may be induced by the architecture or the nature of the players. We denote by  $C$ ,  $CU$ , and  $NU$  the disjoint partition of  $[k]$  into the classes of controllable, uncontrollable cooperative, and uncontrollable non-cooperative players, respectively.

In rational synthesis, we seek a strategy for each of the players in  $C$  with which their objectives are guaranteed to be satisfied, assuming rationality of the other players. As we have the best interest of the players in  $C$  in mind, we assume that  $C \neq \emptyset$ . We say that a profile  $\pi = \langle f_1, \dots, f_k \rangle$  is a *C-fixed NE*, if no player in  $CU \cup NU$  has a beneficial deviation. Formally, we have the following.

**Definition 1. [Rational Synthesis]** Consider a  $k$ -player game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . The problem of rational synthesis (RS) is to return a  $(C \cup CU)$ -profile  $\pi'$  such that there is a C-fixed NE that extends  $\pi'$ , and for every C-fixed NE  $\pi$  that extends  $\pi'$ , we have that  $C \subseteq \text{Win}(\pi)$ .

Two special cases of rational synthesis have been studied in the literature. The first is *cooperative rational synthesis*, where all uncontrollable players are cooperative [12]. The second is *non-cooperative rational synthesis*, where all uncontrollable players are non-cooperative [16].

**Definition 2. [Cooperative Rational Synthesis]** Consider a  $k$ -player game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$  with  $NU = \emptyset$ . The problem of cooperative rational synthesis (CRS) is to return a C-fixed NE  $\pi$  such that  $C \subseteq \text{Win}(\pi)$ .

**Definition 3. [Non-Cooperative Rational Synthesis]** Consider a  $k$ -player game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$  with  $CU = \emptyset$ . The problem of non-cooperative rational synthesis (NRS) is to return a C-profile  $\pi'$  such that there is a C-fixed NE that extends  $\pi'$ , and for every C-fixed NE  $\pi$  that extends  $\pi'$ , we have that  $C \subseteq \text{Win}(\pi)$ .

*Remark 1.* The original rational synthesis problem does not include a game graph [12]. Instead, the set  $AP$  over which the objectives are defined is partitioned among the players, and at each round of the game, each player chooses an assignment to the subset of  $AP$  she controls. It is easy to see that this setting is a special case of our setting, taking the graph to have vertices in  $2^{AP}$ .  $\square$

*Remark 2.* In previous work, the definition of NRS does not require the existence of a C-fixed NE that extends  $\pi'$  [16, 19]. In some settings (in particular, turn-based games), the existence of such an NE is guaranteed. In others (in particular,



concurrent games) there need not be an NE in games with three or more players [8]. Note, however, that even with the requirement that a  $\mathbf{C}$ -fixed NE that extends  $\pi'$  exists, there is no guarantee that best response dynamics from  $\pi'$  would lead to such a  $\mathbf{C}$ -fixed NE.  $\square$

As in traditional synthesis, one can also define the corresponding decision problems, of *rational realizability*, where we only need to decide whether the desired strategies exist. In order to avoid additional notations, we sometimes refer to RS, CRS, and NRS also as decision problems.

For a set  $W \subseteq [k]$ , we say that a solution to the rational synthesis problem is a *W-solution* iff it is a solution that guarantees the winning of exactly the players in  $W$ . In particular, a  $(\mathbf{C} \cup \mathbf{C} \cup \mathbf{C})$ -profile  $\pi'$  is a *W-RS solution* if it is an RS solution such that for every  $\mathbf{C}$ -fixed NE  $\pi$  that extends  $\pi'$ , we have that  $W = \text{Win}(\pi)$ ; a profile  $\pi$  is a *W-CRS solution* if  $\pi$  is a CRS solution such that  $W = \text{Win}(\pi)$ ; and a  $\mathbf{C}$ -profile  $\pi'$  is a *W-NRS solution* if  $\pi'$  is an NRS solution such that for every  $\mathbf{C}$ -fixed NE  $\pi$  that extends  $\pi'$ , we have that  $W = \text{Win}(\pi)$ .

It is easy to see that since the players in  $\mathbf{C}$  are controllable, we can treat them as a single player with an objective that is the conjunction of the objectives of the players in  $\mathbf{C}$ . Accordingly, in the sequel we assume that  $\mathbf{C} = \{1\}$ .

*Remark 3.* It is easy to add to the setting *uncontrollable hostile* players, namely players that, as in traditional synthesis, do not have an objective. Indeed, an uncontrollable hostile player is equivalent to an uncontrollable (cooperative or non-cooperative) player with objective  $\neg\psi_1$ .  $\square$

## 4 The Complexity of Cooperative Rational Synthesis

In this section we study the complexity of CRS. Consider a  $k$ -player concurrent game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . A strategy for Player  $i$  can be viewed as an  $A_i$ -labeled  $V$ -tree, and a profile can be viewed as an  $(A_1 \times \dots \times A_k)$ -labeled  $V$ -tree. Formally, if  $\pi = \langle f_1, \dots, f_k \rangle$  then for every node  $h \in V^*$  in the *profile tree*  $\langle V^*, \pi \rangle$ , we have  $\pi(h) = \langle f_1(h), \dots, f_k(h) \rangle$ , where  $\langle V^*, f_i \rangle$  is the *strategy tree* that corresponds to  $f_i$ . Note that  $\text{Outcome}(\pi)$  then corresponds to a path in  $\langle V^*, \pi \rangle$ .

Viewing profiles as labeled trees enables us to reduce CRS to the nonemptiness of a tree automaton. Essentially, the automaton accepts all profile trees that are solutions to the CRS problem. We define the automaton by decomposing the solutions according to the set of players that win. Given a set  $W$  of players with  $1 \in W$ , a profile  $\pi$  is a  $W$ -CRS solution iff  $\pi$  is a  $1$ -fixed  $W$ -NE. Thus, iff exactly the players in  $W$  win in  $\text{Outcome}(\pi)$ , and for every  $i \notin W$ , Player  $i$  loses in  $\text{Outcome}(\pi[i \leftarrow f'_i])$ , for every strategy  $f'_i$ . In Theorem 1 below, we construct automata that check these conditions.

**Theorem 1.** *Consider a set of players  $W$  with  $1 \in W$ . We can construct the following tree automata over  $(A_1 \times \dots \times A_k)$ -labeled  $V$ -trees:*

- An NBT  $\mathcal{N}_W$  that accepts a profile tree  $\langle V^*, \pi \rangle$  iff  $\text{Win}(\pi) = W$ . The size of  $\mathcal{N}_W$  is polynomial in  $|G|$  and exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ .

- For every  $i \notin W$ , a UCT  $\mathcal{U}_W^i$  that accepts a profile tree  $\langle V^*, \pi \rangle$  iff  $i \in \text{Lose}(\pi[i \leftarrow f'_i])$ , for every strategy  $f'_i$ . The size of  $\mathcal{U}_W^i$  is polynomial in  $|G|$  and exponential in  $|\psi_i|$ .

*Proof.* We start with the NBT  $\mathcal{N}_W$ . Recall that we want  $\mathcal{N}_W$  to accept a profile tree  $\langle V^*, \pi \rangle$  iff exactly the players in  $W$  win in  $\text{Outcome}(\pi)$ . Let  $\psi = \bigwedge_{i \in W} \psi_i \wedge \bigwedge_{i \notin W} (\neg \psi_i)$ , and let  $\mathcal{A} = \langle 2^{AP}, Q, q_0, \mu, \alpha \rangle$  be an NBW of size exponential in  $|\psi|$  that corresponds to  $\psi$ . The NBT  $\mathcal{N}_W$  follows the outcome in the profile tree, and checks if the players in  $W$  are exactly the winners of the profile. Formally,  $\mathcal{N}_W = \langle A_1 \times \dots \times A_k, V, V \times Q, \langle v_0, q_0 \rangle, \eta, V \times \alpha \rangle$ , where for every  $\langle v, q \rangle \in V \times Q$  and  $\langle a_1, \dots, a_k \rangle \in A_1 \times \dots \times A_k$ , we have that  $\eta(\langle v, q \rangle, \langle a_1, \dots, a_k \rangle) = \bigvee_{q' \in \mu(q, \tau(v))} (\delta(v, \langle a_1, \dots, a_k \rangle), \langle \delta(v, \langle a_1, \dots, a_k \rangle), q' \rangle)$ .

We continue to the UCT  $\mathcal{U}_W^i$ . Recall that we want  $\mathcal{U}_W^i$  to accept a profile tree  $\langle V^*, \pi \rangle$  iff Player  $i$  loses in  $\text{Outcome}(\pi[i \leftarrow f'_i])$ , for every strategy  $f'_i$ . Let  $\mathcal{U}_i = \langle 2^{AP}, Q_i, q_i^0, \delta_i, \alpha_i \rangle$  and  $\neg \mathcal{U}_i = \langle 2^{AP}, S_i, s_i^0, \mu_i, \beta_i \rangle$  be the UCWs corresponding to  $\psi_i$  and  $\neg \psi_i$ , respectively. The UCT  $\mathcal{U}_W^i$  follows every possible deviation for Player  $i$ , and checks that indeed she always loses. Formally,  $\mathcal{U}_W^i = \langle A_1 \times \dots \times A_k, V, V \times S_i, \langle v_0, s_i^0 \rangle, \eta, V \times \beta_i \rangle$ , where for every  $\langle v, s \rangle \in V \times S_i$  and  $\langle a_1, \dots, a_k \rangle \in A_1 \times \dots \times A_k$ , we have that  $\eta(\langle v, s \rangle, \langle a_1, \dots, a_k \rangle) = \bigwedge_{a'_i \in \kappa_i(v)} \bigwedge_{s' \in \mu_i(s, \tau(v))} (\delta(v, \langle a_1, \dots, a'_i, \dots, a_k \rangle), \langle \delta(v, \langle a_1, \dots, a'_i, \dots, a_k \rangle), s' \rangle)$ .  $\square$


**Theorem 2.** *Solving CRS can be done in time polynomial in  $|G|$ , exponential in  $|\psi_1|$ , and doubly-exponential in  $|\psi_2|, \dots, |\psi_k|$ . The problem is EXPTIME-hard in  $|\psi_1|$  and 2EXPTIME-hard in each of  $|\psi_2|, \dots, |\psi_k|$ .*

*Proof.* We start with the upper bound. It is easy to see that for every set  $W \subseteq [k]$  of players with  $1 \in W$ , there is a  $W$ -CRS solution iff the intersection of the automata constructed in Theorem 1 is nonempty. We construct an NBT  $\mathcal{A}$  such that  $L(\mathcal{A}) \neq \emptyset$  iff  $L(\mathcal{N}_W) \cap \bigcap_{i \notin W} L(\mathcal{U}_W^i) \neq \emptyset$ , and  $|\mathcal{A}|$  is polynomial in  $|G|$ , exponential in  $|\psi_1|$ , and doubly-exponential in  $|\psi_2|, \dots, |\psi_k|$ . Since nonemptiness of NBTs can be checked in quadratic time [27], the upper bound follows. Moreover, when  $L(\mathcal{A}) \neq \emptyset$ , the algorithm returns a witness to  $\mathcal{A}$ 's nonemptiness, namely a profile tree that is a solution to the CRS problem.

The construction of  $\mathcal{A}$  involves two challenges. First, a naive analysis of the blow-up involved in translating UCTs to NBTs is exponential in the state space of the UCT. In our case, the state space of a UCT  $\mathcal{U}_W^i$  is of the form  $V \times S$ , for some set  $S$  that is independent of  $G$ . Also, the  $V$ -component is updated deterministically: all states sent to the same direction  $v$  of the tree agree on their  $V$ -element. Consequently, the exponential blow up is only in the  $S$ -component, which depends only on  $|\psi_i|$ . Second, the transformation of UCTs to NBTs that is described in [21] preserves nonemptiness, whereas here we need to preserve nonemptiness of an intersection of automata. As detailed in [18], where we coped with a similar challenge, this can be handled by parameterizing the construction in [21] by a rank (essentially, a bound on the size of transducers that generate trees in the language of the automaton) that corresponds to the size of the intersection.

We continue to the lower bounds, and we show they are valid already in the case  $k = 2$ . Proving an EXPTIME lower bound in  $|\psi_1|$ , we describe a reduction from the membership problem for linear-space alternating Turing machines (ATM), defined in the full version. That is, given an ATM  $M$  with space complexity  $s : \mathbb{N} \rightarrow \mathbb{N}$  and a word  $w$ , we construct a 2-player turn-based game  $\mathcal{G} = \langle G, \{\psi_1, \psi_2\} \rangle$ , such that  $G$  and  $\psi_2$  are of a fixed size,  $\psi_1$  is of size linear in  $s(|w|)$ , and there is a CRS solution in  $\mathcal{G}$  iff  $M$  accepts  $w$ .

Essentially, Player 1 and Player 2 generate a branch in the computation tree of  $M$  on  $w$ . Player 1 chooses the letters of the current configuration one by one, and chooses, at the end of each existential configuration, the successor configuration to which the branch continues. Player 2, on the other hand, only chooses successor configurations at the end of each universal configuration (see Figure 1). The objective of Player 1 is to reach an accepting configuration, and the objective of Player 2 is to reach a rejecting configuration.



**Fig. 1.** The game graph  $G$ . The circles are vertices controlled by Player 1, and the square is a vertex controlled by Player 2.

We prove that  $\mathcal{G}$  has a  $\{1\}$ -NE that satisfies  $\psi_1$  iff  $M$  accepts  $w$ . First, if  $M$  accepts  $w$ , then the profile in which Player 1 follows a strategy that generates the configurations in the accepting computation and chooses the appropriate successors to existential configurations, is a  $\{1\}$ -NE that satisfies  $\psi_1$ . Also, if  $M$  rejects  $w$ , then Player 2 can choose successors of universal configurations in a way that leads to a rejecting configuration. Thus, there is no  $\{1\}$ -NE in  $\mathcal{G}$  that satisfies  $\psi_1$ , as either Player 1 loses by not forming a valid branch, or Player 2 can deviate to a strategy where she wins and Player 1 loses. In the full version, we give the details of the reduction.

Proving a 2EXPTIME lower bound in  $|\psi_2|$ , we use a reduction from decidability of 2-player zero-sum games, which is 2EXPTIME-hard already for a game

with a game graph of a fixed size [2]. Given a 2-player zero-sum game  $\mathcal{G} = \langle G, \psi \rangle$ , we construct a 2-player game  $\mathcal{H} = \langle H, \{\psi_1, \psi_2\} \rangle$  such that the size of  $H$  is linear in  $|G|$ ,  $\psi_1$  is of a fixed size,  $\psi_2$  is of size linear in  $|\psi|$ , and there is a CRS solution in  $\mathcal{H}$  iff Player 1 wins  $\mathcal{G}$ . Essentially, the game graph  $H$  contains two copies of  $G$ , and a new initial vertex in which Player 2 chooses between proceeding to the first or the second copy.

Note that Player 1 has no influence in that decision. Then, the objective of Player 1 is for the play to be generated in the first copy, and the objective of Player 2 is for the play to be generated in the second copy and for the computation to not satisfy  $\psi$ .  $\square$

*Remark 4.* Note that our algorithm finds W-CRS solutions for all  $W \subseteq [k]$ , and so it is exponential in  $k$ . As shown in [10], rational synthesis is PSPACE in  $k$  already for rational synthesis with reachability objectives.

## 5 The Complexity of Non-Cooperative Rational Synthesis

In this section we study the complexity of NRS. We start with the turn-based setting, and then proceed the concurrent setting.

### 5.1 Turn-based games

As in the CRS case, we construct a tree automaton that accepts strategy trees that are NRS solutions. Here, however, the trees are labeled not only by a strategy for Player 1, but also by information that certifies that the suggested strategy is indeed a solution. Our construction follows the ideas developed for turn-based games in [10], adding to them a treatment of the LTL objectives (the latter is not too complicated, and our main goal in this section is to set the stage to the concurrent setting, which was left open in [10]). In order to present our solution, we first need some definitions and notations.

Consider a  $k$ -player turn-based game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . Let  $G = \langle AP, \{V_i\}_{i \in [k]}, v_0, E, \tau \rangle$ . Recall that for a subset  $S \subseteq [k]$  of players, an  $S$ -profile is a set of strategies, one for each player in  $S$ , and that a profile  $\pi$  extends an  $S$ -profile  $\pi'$  if the players in  $S$  use in  $\pi$  their strategies in  $\pi'$ . The outcome of an  $S$ -profile  $\pi'$ , denoted  $\text{Outcome}(\pi')$ , is the union of plays that are outcomes of profiles that extend  $\pi'$ . Thus,  $\text{Outcome}(\pi') \subseteq V^\omega$  is the set of plays that are possible outcome of the game when the players in  $S$  follow their strategies in  $\pi'$ .

Consider a profile  $\pi = \langle f_1, \dots, f_k \rangle$  and a prefix  $h \in V^*$  of  $\text{Outcome}(\pi)$ . For a profile  $\pi' = \langle f'_1, \dots, f'_k \rangle$ , we define the profile  $\text{switch}(\pi, \pi', h) = \langle f_1^h, \dots, f_k^h \rangle$  as the profile in which the players first follow  $\pi$  and generate  $h$ , and then switch to following  $\pi'$ . Formally, for every  $x \in V^*$  and Player  $i \in [k]$ , if  $x \preceq h$ , then  $f_i^h(x) = f_i(x)$ , and if  $x = h \cdot y$ , then  $f_i^h(x) = f'_i(y)$ . Note that since the last vertices in  $x$  and  $y$  coincide, then  $\text{switch}(\pi, \pi', h)$  is well defined, in the sense that it returns only allowed actions. Also note that  $f_i^h(h) = f'_i(\varepsilon)$ , thus, switching to following  $\pi'$ , we reset the history of the game so far. The strategies in nodes that are neither a prefix of  $h$  nor an extension of  $h$  are arbitrary and can follow  $\pi$ .

For  $i \in [k] \setminus \{1\}$  and a prefix  $h \in V^* \cdot V_i$  of some play in  $\text{Outcome}(\{f_1\})$ , we say that *Player  $i$  wins from  $h$*  if there exists a strategy  $f'_i$  for Player  $i$  such that for every profile  $\pi = \langle f_1, \dots, f_k \rangle$  with  $h \preceq \text{Outcome}(\pi)$ , we have that  $i \in \text{Win}(\text{switch}(\pi, \pi[i \leftarrow f'_i], h))$ . That is, Player  $i$  wins in every profile in which the players first generate  $h$ , and then Player  $i$  switches to following  $f'_i$ , while the other players adhere to their strategies in the original profile. The strategy  $f'_i$  is then called an  *$h$ -winning strategy for Player  $i$* .

Since turn-based games always have an NE, an NRS solution in  $\mathcal{G}$  is a strategy  $f_1$  for Player 1 such that for every 1-fixed NE  $\pi = \langle f_1, \dots, f_k \rangle$ , we have that  $1 \in \text{Win}(\pi)$ . Equivalently, for every profile  $\pi = \langle f_1, \dots, f_k \rangle$ , we have that either  $1 \in \text{Win}(\pi)$ , or there exists  $i \in \text{Lose}(\pi)$  such that  $i \in \text{Win}(\pi[i \leftarrow f'_i])$  for some strategy  $f'_i$  for Player  $i$ . As detailed in [10], this implies that a strategy  $f_1$  for Player 1 is an NRS solution iff for every path  $\rho$  in  $\text{Outcome}(\{f_1\})$ , either  $\tau(\rho) \in L(\psi_1)$ , or there is  $i \in [k] \setminus \{1\}$  such that  $\tau(\rho) \notin L(\psi_i)$ , and there are a prefix  $h \preceq \rho$  and an  $h$ -winning strategy  $f'_i$  for Player  $i$ . We then say that  $h$  is a *good deviation point* for Player  $i$ , and  $f'_i$  is a *good deviation* for Player  $i$ .

Our goal is to define a tree automaton that accepts a strategy tree for Player 1 iff it is an NRS solution. The tree automaton should check that every path in  $\text{Outcome}(\{f_1\})$  that does not satisfy  $\psi_1$  has a good deviation point for one of the players that lose in it. For that purpose, a strategy  $f_1$  of Player 1 is going to be *certified* by information about deviations: each path in  $\text{Outcome}(\{f_1\})$  that does not satisfy  $\psi_1$  is labeled by a player  $i$  that loses in the path, a good deviation point for Player  $i$ , and a good deviation for Player  $i$ . Note that each deviation may handle only a subset of the paths below the good deviation point, and thus a subtree in the certified strategy tree may be labeled by strategies of different players, each deviating at different points.

Formally, a certified strategy tree is a  $((V \cup \{\circ\}) \times [k])$ -labeled  $V$ -tree  $\langle V^*, g \rangle$ , where each node is labeled by a pair  $\langle v, i \rangle$ , where  $v \in V \cup \{\circ\}$  is a *strategy-label*, and  $i \in [k]$  is a *player-label*. We use  $g_s$  and  $g_p$  to refer to the projection of  $g$  on the strategy and player components. Each path in the tree that corresponds to a play in  $\text{Outcome}(\{f_1\})$  has a suffix all whose nodes are labeled by the same player label. If this label is 1, then the strategy labels describe a strategy of Player 1 and the path should satisfy  $\psi_1$ . If this label is  $i \in [k] \setminus \{1\}$ , then a deviation point of Player  $i$  has been encountered, the strategy labels describe a good deviation for Player  $i$ , and the path should not satisfy  $\psi_i$ . As long as a deviation point has not been encountered, the strategy labels describe a strategy for Player 1 (and so, they are in  $V$  in nodes with a direction in  $V_1$ , and are  $\circ$  in nodes with a direction not in  $V_1$ ). Once a deviation point for Player  $i$  is encountered (which is indicated by the strategy label being changed from  $\circ$  to a vertex in  $V$  in a node with direction  $V_i$ ), the strategy labels describe a strategy for Player  $i$ .

By adjusting Lemma 7 in [10] to the setting with LTL objectives, we get the following.

**Theorem 3.** *A strategy  $f_1$  for Player 1 is an NRS solution iff there is a certified strategy tree  $\langle V^*, g \rangle$  that agrees with  $f_1$ . Thus, for every  $h \in V^*$  and  $v \in V_1$  such that  $h \cdot v \in \text{Outcome}(\{f_1\})$ , we have that  $g_s(h \cdot v) = f_1(h \cdot v)$*

We now define a tree automaton that accepts certified strategy trees, which we then use for solving NRS.

**Theorem 4.** *We can construct a UCT  $\mathcal{U}$  over  $((V \cup \{\emptyset\}) \times [k])$ -labeled  $V$ -trees such that  $\mathcal{U}$  accepts a  $((V \cup \{\emptyset\}) \times [k])$ -labeled  $V$ -tree  $\langle V^*, g \rangle$  iff  $\langle V^*, g \rangle$  is a certified strategy tree. The size of  $\mathcal{U}$  is polynomial in  $|G|$  and exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ .*

*Proof.* The requirements on a certified strategy tree  $\langle V^*, g \rangle$  for Player 1 can be decomposed to the following conditions.

- (**C<sub>1</sub><sup>i</sup>**) For every  $i \in [k] \setminus \{1\}$ , the subtree of every node  $h \in V^* \cdot V_i$  in the tree that is labeled by  $V \times [k]$  is labeled by an  $h$ -winning strategy for Player  $i$ .
- (**C<sub>2</sub>**) The (infinite) suffix of every path in the tree is  $p$ -labeled by a single  $i \in [k]$ .
- (**C<sub>3</sub><sup>i</sup>**) For every  $i \in [k] \setminus \{1\}$ , every path in the tree with a suffix  $p$ -labeled by  $i$  has a good deviation point for Player  $i$ .
- (**C<sub>4</sub><sup>1</sup>**) Player 1 wins in every path in the tree with suffix  $p$ -labeled by 1.
- (**C<sub>4</sub><sup>i</sup>**) for every  $i \in [k] \setminus \{1\}$ , Player  $i$  loses in every path in the tree with suffix  $p$ -labeled by  $i$ .

In the full version, we describe UCTs that check these conditions and whose intersection is of the desired size.  $\square$

**Theorem 5.** *Solving NRS can be done in time polynomial in  $|G|$  and doubly-exponential in  $|\psi_1|, \dots, |\psi_k|$ . The problem is 2EXPTIME-hard in each of  $|\psi_1|, \dots, |\psi_k|$ .*

*Proof.* We start with the upper bound. By Theorems 3 and 4, we can reduce NRS to nonemptiness of a UCT  $\mathcal{U}$  over  $((V \cup \{\emptyset\}) \times [k])$ -labeled  $V$ -trees of size polynomial in  $|G|$  and exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ . Using considerations similar to these used in the proof of Theorem 2 (in particular, the fact  $\mathcal{U}$  is deterministic in its  $V$ -element), we can construct from it an NBT  $\mathcal{N}$  of size polynomial in  $|G|$  and doubly-exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$  that preserves the nonemptiness of  $\mathcal{U}$ . Since the nonemptiness problem for NBT can be solved in quadratic time [27], the desired complexity follows.

We continue to the lower bounds, and we show they are valid already in the case  $k = 2$ . We again use reductions from deciding 2-player zero-sum games. In order to prove 2EXPTIME-hardness in  $|\psi_1|$ , consider a 2-player zero-sum game  $\mathcal{G} = \langle G, \psi \rangle$ , for a fixed-size  $G$ . We claim that the 2-player game  $\mathcal{G}' = \langle G, \{\psi, \mathbf{true}\} \rangle$  is such that  $G$  is of a fixed size and that there is an NRS solution in  $\mathcal{G}'$  iff Player 1 wins  $\mathcal{G}$ . Indeed, since the objective of Player 2 is **true**, every profile  $\pi$  in  $\mathcal{G}'$  is a 1-fixed NE. So, in order for a strategy  $f_1$  to be an NRS solution, it must satisfy that  $1 \in \text{Win}(\langle f_1, f_2 \rangle)$ , for every strategy  $f_2$  for Player 2. Equivalently, it is a winning strategy for Player 1 in  $\mathcal{G}$ .

In order to prove 2EXPTIME-hardness in  $|\psi_2|$ , consider again a 2-player zero-sum game  $\mathcal{G} = \langle G, \psi \rangle$ , for a fixed-size  $G$ . We construct a 2-player game  $\mathcal{G}' = \langle H, \{\psi_1, \psi_2\} \rangle$  such that  $H$  and  $\psi_1$  are of a fixed size, the size of  $\psi_2$  is linear in  $|\psi|$ , and there is an NRS solution in  $\mathcal{G}'$  iff Player 1 wins  $\mathcal{G}$ . The game graph  $H$

is as in the proof of Theorem 2. Thus, it has an initial vertex from which Player 2 chooses between two copies of  $G$ . The states of the first copy are labeled by a fresh atomic proposition  $p$ . Then,  $\psi_1 = XGp$ , and  $\psi_2 = X((\psi \wedge Gp) \vee ((\neg\psi) \wedge G\neg p))$ . Thus, the objective of Player 1 is for the play to be generated in the first copy, and the objective of Player 2 is either to generate a play in the first copy whose computation satisfies  $\psi$ , or to generate a play in the second copy whose computation does not satisfy  $\psi$ .

If Player 1 has a winning strategy  $f_1$  in  $\mathcal{G}$ , then there is an NRS solution  $f'_1$  in  $\mathcal{G}'$ , where  $f'_1$  follows  $f_1$  in the copy Player 2 chooses. Indeed, as  $f'_1$  guarantees the satisfaction of  $\psi$  for all possible behaviors of Player 2, a profile  $\pi$  is a 1-fixed NE only if Player 2 chooses the first copy. If Player 1 loses  $\mathcal{G}$ , then for every strategy  $f_1$  for Player 1, there is a strategy  $f_2$  for Player 2 such that  $\text{Outcome}(\langle f_1, f_2 \rangle)$  does not satisfy  $\psi$ . So, for every strategy  $f_1$  for Player 1 in  $\mathcal{G}'$ , we have that there is a 1-fixed NE  $\pi = \langle f_1, f_2 \rangle$  such that  $1 \in \text{Lose}(\pi)$ , where  $f_2$  is the strategy that chooses the second copy, and ensures that  $\psi$  is not satisfied. Hence, there is no NRS solution in  $\mathcal{G}'$ .  $\square$

## 5.2 Concurrent games

Consider a  $k$ -player concurrent game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . Let  $G = \langle AP, V, v_0, \{A_i\}_{i \in [k]}, \{\kappa_i\}_{i \in [k]}, \delta, \tau \rangle$ . As our constructions in this section are loaded with notations, we simplify the setting and assume that there is one set  $A$  of actions, available to all players in all vertices. That is,  $A_1 = A_2 = \dots = A_k = A$ , and for every  $v \in V$  and  $i \in [k]$ , we have that  $\kappa_i(v) = A$ . All our constructions and results can be easily extended to the general case.

As in the turn-based setting, we define a UCT that accepts certified strategy trees for Player 1. In the concurrent setting, however, certification is much more complicated. Below we explain the challenges in the concurrent setting and how we overcome them. For  $i \in [k] \setminus \{1\}$  and a prefix  $h \in V^*$  of some path in  $\text{Outcome}(\{f_1\})$ , we say that *Player  $i$  wins from  $h$*  if for every profile  $\pi = \langle f_1, \dots, f_k \rangle$  with  $h \preceq \text{Outcome}(\pi)$ , we have that there exists a strategy  $f'_i$  for Player  $i$  such that  $i \in \text{Win}(\text{switch}(\pi, \pi[i \leftarrow f'_i], h))$ . Thus, Player  $i$  wins from  $h$  if she has a beneficial deviation to switch to from  $h$ , for every profile  $\pi$  with  $h \preceq \text{Outcome}(\pi)$ . Note that for different profiles, Player  $i$  might have different such beneficial deviations. Here, however, the prefix  $h$  need not end in  $V_i$  (in fact, there is no  $V_i$  in the concurrent setting). We say that  $h$  is a *winning point* for Player  $i$ . Also, we say that  $(h, \langle f_1(h), \dots, f_k(h) \rangle)$  is a *good deviation pair* for Player  $i$  iff there exists  $a'_i \in A$  such that  $h \cdot \delta(h, \langle f_1(h), \dots, a'_i, \dots, f_k(h) \rangle)$  is a winning point for Player  $i$ .

In order to understand better the difference between NRS solutions in the turn-based and concurrent settings, recall that a strategy  $f_1$  for Player 1 is not an NRS solution iff there is a 1-fixed NE  $\pi = \langle f_1, \dots, f_k \rangle$  whose outcome  $\rho$  does not satisfy  $\psi_1$ . Note that  $\pi$  being a 1-fixed NE means that for every prefix  $h \cdot v \cdot u$  of  $\rho$ , there exist actions  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$  such that  $\delta(v, \langle f_1(h \cdot v), a_2, \dots, a_k \rangle) = u$  and for every  $i \in \text{Lose}(\rho)$ , we have that  $(h \cdot v, \langle f_1(h \cdot v), a_2, \dots, a_k \rangle)$  is not a good deviation pair for Player  $i$ . In particular, we can choose  $a_i = f_i(h \cdot v)$ . Hence,

if  $f_1$  is an NRS solution, and there exists a path  $\rho \in \text{Outcome}(\{f_1\})$  that does not satisfy  $\psi_1$ , then there must be a prefix  $h \cdot v \cdot u \preceq \rho$  such that for every  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$  with  $\delta(v, \langle f_1(h \cdot v), a_2, \dots, a_k \rangle) = u$ , there exists  $i \in \text{Lose}(\rho)$  such that  $(h \cdot v, \langle f_1(h \cdot v), a_2, \dots, a_k \rangle)$  is a good deviation point for Player  $i$ . We then say that  $h \cdot v \cdot u$  is a *good deviation transition for*  $\text{Lose}(\rho)$ . Thus, while in the turn-based settings it is sufficient to find in every path in which Player 1 loses a good deviation point for one of the players that lose in it, in the concurrent setting the definition of good deviation depends on the transition induced by the specific profile being used, and so we have to consider deviating transitions, and there may be several players in  $\text{Lose}(\rho)$  that deviate. Accordingly, in order to certify a strategy for Player 1, we should describe a mapping from every vector of actions to a set of players, along with their deviations.

Another difference between the turn-based setting and the concurrent setting is that only in the first, the existence of some 1-fixed NE is guaranteed [8]. Hence, we have to add to the algorithm such a check (which is in fact easy).

We can now define certified strategy trees for the concurrent setting. Every node in a certified strategy tree is labeled by the following components:

1. An action  $a_1 \in A$ , which is the strategy for Player 1.
2. A *deviation function*  $d : A^{k-1} \rightarrow (A \cup \{\perp\})^{k-1}$ , which maps a vector of actions of players  $2, \dots, k$  to the set of players that deviate from it, along with their deviations. Specifically,  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$  being mapped to  $\langle a'_2, \dots, a'_k \rangle \in (A \cup \{\perp\})^{k-1}$  indicates that for every  $i \in [k] \setminus \{1\}$ , if  $a'_i \in A$ , then  $a'_i$  is the deviation for Player  $i$  from  $\langle a_2, \dots, a_k \rangle$ , and if  $a'_i = \perp$  then no deviation from  $\langle a_2, \dots, a_k \rangle$  is specified for Player  $i$ . Let  $\mathcal{D}$  denote the set of all possible deviation functions.
3. A set  $L \subseteq \{2, \dots, k\}$  of players, which describes the set of players that lose in a given path and which are therefore expected to have a good deviation transition. That is, if a suffix of a path is labeled by  $\emptyset$ , then Player 1 should win in this path, and if a suffix of a path is labeled by  $L \neq \emptyset$ , then all the players in  $L$  lose in this path.
4. A vector of actions  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$ , which describes the strategies for the other players in the required 1-fixed NE.

Formally, a certified strategy tree is a  $(A \times \mathcal{D} \times 2^{\{2, \dots, k\}} \times A^{k-1})$ -labeled  $V$ -tree  $\langle V^*, g \rangle$ , where each node is labeled by both a *strategy-label*  $a_1 \in A$ , a *deviation-label*  $d \in \mathcal{D}$ , a *player-label*  $L \in 2^{\{2, \dots, k\}}$ , and an *NE-label*  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$ . We use  $g_s, g_d, g_p$ , and  $g_{NE}$  to refer to the projection of  $g$  on its different components.

For a node  $h \cdot v$  that is  $s$ -labeled by  $a_1$  and  $d$ -labeled by  $d$ , a possible successor  $u$  of  $v$ , and a set of losers  $L$ , we say that  $h \cdot v \cdot u$  is *marked as a good deviation transition for*  $L$  iff the following hold:

1. For every  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$  such that  $\delta(v, \langle a_1, a_2, \dots, a_k \rangle) = u$ , there exists  $i \in L$  such that  $(d(\langle a_2, \dots, a_k \rangle))_i \in A$ . That is, for every vector of actions  $\langle a_2, \dots, a_k \rangle$  that leads to  $u$ , there is  $i \in L$  such that  $d$  assigns a deviation for Player  $i$  from  $\langle a_2, \dots, a_k \rangle$ .



2. For every  $i \in L$ , there is a vector of actions  $\langle a_2, \dots, a_k \rangle \in A^{k-1}$  such that  $\delta(v, \langle a_1, a_2, \dots, a_k \rangle) = u$ , and  $(d(\langle a_2, \dots, a_k \rangle))_i \in A$ . That is, for every  $i \in L$  there exists a vector of actions that leads to  $u$ , from which  $d$  assigns a deviation for Player  $i$ .

Now, an  $(A \times \mathcal{D} \times 2^{\{2, \dots, k\}} \times A^{k-1})$ -labeled  $V$ -tree  $\langle V^*, g \rangle$  is a certified strategy tree iff it satisfies the following conditions:

- (**C<sub>1</sub>**) If  $h \cdot v \cdot u \in V^*$  is marked as a good deviation transition for a set of players  $L \subseteq \{2, \dots, k\}$ ,  $L \neq \emptyset$ , then it is indeed a good deviation transition for  $L$ .
- (**C<sub>2</sub>**) Every path  $\rho$  has a set  $L$  such that  $\rho$  is eventually always  $p$ -labeled by  $L$ .
- (**C<sub>3</sub><sup>L</sup>**) For every  $L \subseteq \{2, \dots, k\}$ ,  $L \neq \emptyset$ , every path in the tree with a suffix  $p$ -labeled by  $L$  has a good deviation transition for  $L$ .
- (**C<sub>4</sub><sup>1</sup>**) Player 1 wins in every path in the tree with suffix  $p$ -labeled by  $\emptyset$ .
- (**C<sub>4</sub><sup>i</sup>**) For every  $i \in [k] \setminus \{1\}$ , Player  $i$  loses in every path in the tree with suffix  $p$ -labeled by  $L$  such that  $i \in L$ .
- (**C<sub>5</sub>**) The  $s$  and  $NE$ -labeling of the tree specifies a 1-fixed NE.

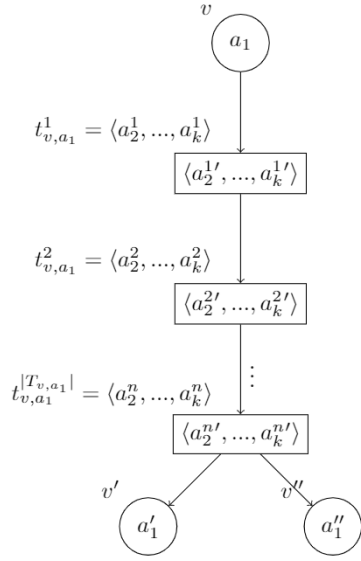
**Theorem 6.** *A strategy  $f_1$  for Player 1 is an NRS solution iff there is a certified strategy tree  $\langle V^*, g \rangle$  that agrees with  $f_1$ . That is, for every  $h \in V^*$ , we have that  $g_s(h) = f_1(h)$ .*

**Theorem 7.** *We can construct a UCT over  $(A \times \mathcal{D} \times 2^{\{2, \dots, k\}} \times A^{k-1})$ -labeled  $V$ -trees that accepts a  $(A \times \mathcal{D} \times 2^{\{2, \dots, k\}} \times A^{k-1})$ -labeled  $V$ -tree  $\langle V^*, g \rangle$  iff  $\langle V^*, g \rangle$  is a certified strategy tree. The size of the UCT is polynomial in  $|G|$  and exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ .*

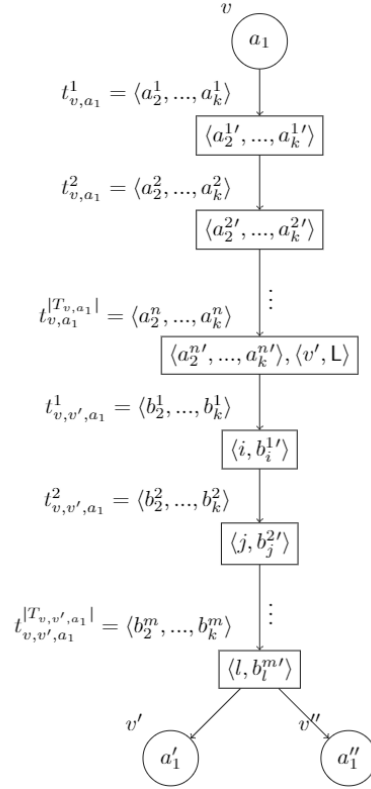
*Proof.* We can construct UCTs for **C<sub>2</sub>**, **C<sub>3</sub><sup>L</sup>**, and **C<sub>4</sub><sup>i</sup>** that are very similar to the UCTs for **C<sub>2</sub>**, **C<sub>3</sub><sup>i</sup>** and **C<sub>4</sub><sup>i</sup>** from the turn-based setting. The UCT for **C<sub>5</sub>** is similar to the UCT for CRS solutions. In the full version, we describe in detail a UCT that checks the satisfaction of **C<sub>1</sub>**. Then, the conjunction of the above UCTs results in a UCT that accepts certified strategy trees.  $\square$

The number of deviation functions  $d : A^{k-1} \rightarrow (A \cup \{\perp\})^{k-1}$  is exponential in  $|A|$ . Consequently, the UCT described in Theorem 7 has an exponential alphabet, which causes the NBT generated in [21] to have exponentially many transitions, making its nonemptiness problem exponential. In order to overcome this problem, we introduce *vertical annotation of certified strategy trees*, which essentially replace a node labeled by  $d \in \mathcal{D}$  by a sequence of nodes, labeled by a smaller alphabet.

Explaining our vertical annotation, we find it clearer to go back to a detailed description of the actions of the different players, thus refer to  $A_i$  and  $\kappa_i$  rather than assuming they are all equal to  $A$ . For every vertex  $v \in V$  and an action  $a_1 \in \kappa_1(v)$  for Player 1, we denote by  $T_{v, a_1}$  the set of possible vectors of actions from  $v$ , given Player 1 chose  $a_1$ . That is,  $T_{v, a_1} = \{a_1\} \times \kappa_2(v) \times \dots \times \kappa_k(v)$ . We



**Fig. 2.** A vertically certified strategy tree. Information about deviations from a given vector of actions is stored in an intermediate node that corresponds to that vector. There is no good deviation transition starting from  $v$ .



**Fig. 3.**  $h \cdot v \cdot v'$  is a good deviation transition for  $L$ , where the root  $v$  is reachable via history  $h$ .

order the vectors in  $T_{v,a_1}$  arbitrarily, and, for  $1 \leq i \leq |T_{v,a_1}|$ , denote by  $t_{v,a_1}^i$  the  $i$ -th item in  $T_{v,a_1}$ . We also denote by  $t_{v,a_1}^i[j \leftarrow a'_j]$  the vector of actions obtained from  $t_{v,a_1}^i$  by replacing the action for Player  $j$  by  $a'_j$ .

For a given transition from  $v$  to  $v'$ , let  $T_{v,v',a_1}$  be the restriction of  $T_{v,a_1}$  to vectors  $\langle a_1, a_2, \dots, a_k \rangle$  such that  $\delta(v, \langle a_1, a_2, \dots, a_k \rangle) = v'$ , and let  $t_{v,v',a_1}^i$  denote the  $i$ -th item in  $T_{v,v',a_1}$ . Also, let  $T = \bigcup_{v \in V} \bigcup_{a_1 \in \kappa_1(v)} T_{v,a_1}$ , and  $\Sigma = A_1 \cup ((A_2 \times \dots \times A_k) \cup \bigcup_{j \in L} (\{j\} \times A_j)) \times (\emptyset \cup (V \times 2^{\{2, \dots, k\}}))$ . We also need an additional  $2^{\{2, \dots, k\}}$  component, for annotating the set of losers in each path in the tree, but we omit it for now, as it is not relevant for the vertical annotations.

A certified strategy tree is now a  $\Sigma$ -labeled  $(V \cup T)$ -tree, where nodes with direction in  $V$  are labeled by the strategy for Player 1, and nodes with direction in  $T$  are labeled with deviation information. Consider a node that corresponds to a vertex  $v$  with history  $h$  and is labeled by  $a_1$ . Following vertically there are nodes corresponding to  $T_{v,a_1}$ , each labeled by a vector of actions (see Figure 2).

This information is for verifying good deviation transitions. That is, after a good deviation transition is announced, we need to verify that the involved players indeed have appropriate beneficial deviations. The last node in the chain, which corresponds to  $t_{v,a_1}^{|T_{v,a_1}|}$ , is either not labeled, or labeled by a vertex  $v'$  and a set  $L$  of losers. If it is not labeled, it means there are no good deviation transitions from  $v$ , in which case we continue to the nodes corresponding to the possible successors of  $v$ . If it is labeled by  $\langle v', L \rangle$ , it means that  $h \cdot v \cdot v'$  is a good deviation transition for  $L$  (see Figure 3). Hence, the following nodes correspond to  $T_{v,v',a_1}$ , each labeled by a single deviation, followed by a chain of good deviation transitions, using the appropriate  $V \times 2^{\{2, \dots, k\}}$  annotations, or by nodes corresponding to successor vertices. In the full version, we describe formally a UCT for verifying good deviation transitions in vertically annotated certified strategy trees.

**Theorem 8.** *Solving NRS in the concurrent setting can be done in time polynomial in  $|G|$  and doubly-exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ . The problem is 2EXPTIME-hard in each of  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ .*

*Proof.* We start with the upper bound. We can easily modify the UCTs for  $\mathbf{C}_2$ ,  $\mathbf{C}_3^L$ ,  $\mathbf{C}_4^i$ , and  $\mathbf{C}_5$  from the proof of Theorem 7 to accommodate the vertical annotation. With conjunction with the UCT for verifying good deviation transitions, described in the full version, we have a UCT  $\mathcal{U}$  over  $\Sigma$ -labeled  $(V \cup T)$ -trees such that  $\mathcal{U}$  accepts a  $\Sigma$ -labeled  $(V \cup T)$ -tree  $\langle V^*, g \rangle$  iff  $\langle V^*, g \rangle$  is a vertically annotated certified strategy tree. By Theorem 6, there is an NRS solution in  $\mathcal{G}$  iff  $\mathcal{U}$  is not empty. The size of  $\mathcal{U}$  is polynomial in  $|G|$  and exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ , and its alphabet is polynomial in  $G$ . Also,  $\mathcal{U}$  is deterministic in its  $V$ -element. Hence, as detailed in the proof of Theorem 2, we can construct from  $\mathcal{U}$  an NBT  $\mathcal{N}$  of size polynomial in  $|G|$  and doubly-exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$  such that  $L(\mathcal{U}) \neq \emptyset$  iff  $L(\mathcal{N}) \neq \emptyset$ . Since the nonemptiness problem for NBT can be solved in quadratic time [27], the desired complexity follows.

Finally, as turn-based games are a special case of concurrent ones, the lower bound from Theorem 5 applies here.  $\square$

### 5.3 General rational synthesis

Consider a  $k$ -player game  $\mathcal{G} = \langle G, \{\psi_i\}_{i \in [k]} \rangle$ . Recall that the problem of rational synthesis is to return a  $(\{1\} \cup \text{CU})$ -profile  $\pi'$  such that there is a 1-fixed NE that extends  $\pi'$ , and for every 1-fixed NE  $\pi$  that extends  $\pi'$ , we have that  $1 \in \text{Win}(\pi)$ .

A  $(\{1\} \cup \text{CU})$ -profile  $\pi'$  is an RS-solution iff for every path  $\rho$  such that  $\rho \not\equiv \psi_1$ , there is no 1-fixed NE  $\pi$  that extends  $\pi'$  and  $\text{Outcome}(\pi) = \rho$ . It is easy to see that we can define certified  $(\{1\} \cup \text{CU})$ -profile trees in a similar way we defined certified strategy trees for Player 1, inducing an algorithm of the same complexity for checking the existence of a certified  $(\{1\} \cup \text{CU})$ -profile tree. Also, as NRS is a special case of RS, the NRS lower bound provide a lower bound for RS. Hence, we can conclude with the following.

**Theorem 9.** *Solving RS can be done in time polynomial in  $|G|$  and doubly-exponential in  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ . The problem is 2EXPTIME-hard in each of  $|\psi_1|, |\psi_2|, \dots, |\psi_k|$ .*

## References

1. S. Almagor, O. Kupferman, and G. Perelli. Synthesis of controllable Nash equilibria in quantitative objective game. In *Proc. 27th Int. Joint Conf. on Artificial Intelligence*, pages 35–41, 2018.
2. R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, pages 295–304. IEEE Computer Society, 2004.
4. P. Bouyer-Decitre, O. Kupferman, N. Markey, B. Maubert, A. Murano, and G. Perelli. Reasoning about quality and fuzziness of strategic behaviours. In *Proc. 28th Int. Joint Conf. on Artificial Intelligence*, pages 1588–1594, 2019.
5. A.K. Chandra, D.C. Kozen, and L.J. Stockmeyer. Alternation. *Journal of the Association for Computing Machinery*, 28(1):114–133, 1981.
6. K. Chatterjee, T. Henzinger, and B. Jobstmann. Environment assumptions for synthesis. In *Proc. 19th Int. Conf. on Concurrency Theory*, volume 5201 of *Lecture Notes in Computer Science*, pages 147–161. Springer, 2008.
7. K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In *Proc. 18th Int. Conf. on Concurrency Theory*, pages 59–73, 2007.
8. K. Chatterjee, R. Majumdar, and M. Jurdzinski. On Nash equilibria in stochastic games. In *Proc. 13th Annual Conf. of the European Association for Computer Science Logic*, volume 3210 of *Lecture Notes in Computer Science*, pages 26–40. Springer, 2004.
9. A. Church. Logic, arithmetics, and automata. In *Proc. Int. Congress of Mathematicians, 1962*, pages 23–35. Institut Mittag-Leffler, 1963.
10. R. Condurache, E. Filiot, R. Gentilini, and J.-F. Raskin. The complexity of rational synthesis. In *Proc. 43th Int. Colloq. on Automata, Languages, and Programming*, volume 55 of *LIPICs*, pages 121:1–121:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
11. E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *Proc. 29th IEEE Symp. on Foundations of Computer Science*, pages 328–337, 1988.
12. D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *Proc. 16th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 190–204. Springer, 2010.
13. J. Gutierrez, M. Najib, G. Perelli, and M. Wooldridge. On computational tractability for rational verification. In *Proc. 28th Int. Joint Conf. on Artificial Intelligence*, pages 329–335. International Joint Conferences on Artificial Intelligence Organization, 2019.
14. J. Gutierrez, G. Perelli, and M. J. Wooldridge. Imperfect information in reactive modules games. *Inf. Comput.*, 261:650–675, 2018.
15. O. Kupferman, Y. Lustig, M.Y. Vardi, and M. Yannakakis. Temporal synthesis for bounded systems and environments. In *Proc. 28th Symp. on Theoretical Aspects of Computer Science*, pages 615–626, 2011.
16. O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with rational environments. *Annals of Mathematics and Artificial Intelligence*, 78(1):3–20, 2016.
17. O. Kupferman and N. Piterman. Lower bounds on witnesses for nonemptiness of universal co-Büchi automata. In *Proc. 12th Int. Conf. on Foundations of Software*

- Science and Computation Structures*, volume 5504 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2009.
18. O. Kupferman and N. Shenwald. Perspective games with notifications. In *Proc. 40th Conf. on Foundations of Software Technology and Theoretical Computer Science*, volume 182 of *LIPICs*, pages 51:1–51:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
  19. O. Kupferman and N. Shenwald. Perspective multi-player games. In *Proc. 36th IEEE Symp. on Logic in Computer Science*, 2021.
  20. O. Kupferman and G. Vardi. Perspective games. In *Proc. 34th IEEE Symp. on Logic in Computer Science*, pages 1 – 13, 2019.
  21. O. Kupferman and M.Y. Vardi. Safraless decision procedures. In *Proc. 46th IEEE Symp. on Foundations of Computer Science*, pages 531–540, 2005.
  22. J.F. Nash. Equilibrium points in  $n$ -person games. In *Proceedings of the National Academy of Sciences of the United States of America*, 1950.
  23. N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
  24. C. H. Papadimitriou. Algorithms, games, and the internet. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 749–753, 2001.
  25. A. Pnueli. The temporal semantics of concurrent programs. *Theoretical Computer Science*, 13:45–60, 1981.
  26. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proc. 16th ACM Symp. on Principles of Programming Languages*, pages 179–190, 1989.
  27. M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logics of programs. *Journal of Computer and Systems Science*, 32(2):182–221, 1986.
  28. M. Wooldridge, J. Gutierrez, P. Harrenstein, E. Marchioni, G. Perelli, and A. Toumi. Rational verification: From model checking to equilibrium checking. In *Proc. of 30th National Conf. on Artificial Intelligence*, pages 4184–4190, 2016.