

# Safraless Synthesis for Epistemic Temporal Specifications<sup>⊗</sup>

Rodica Bozianu<sup>1,2</sup>, Cătălin Dima<sup>1</sup>, and Emmanuel Filiot<sup>2,⊗⊗</sup>

<sup>1</sup> Université Paris Est, LACL (EA 4219), UPEC, 94010 Créteil Cedex, France

<sup>2</sup> Université Libre de Bruxelles, CP 212 - 1050 Bruxelles Belgium

We address the synthesis problem for specifications given in linear temporal single-agent epistemic logic, KLTL, over single-agent systems having imperfect information of the environment state. [19] have shown that this problem is 2Exptime complete. However, their procedure relies on complex automata constructions that are notoriously resistant to efficient implementations as they use Safra-like determinization.

We propose a "Safraless" synthesis procedure for a large fragment of KLTL. The construction transforms first the synthesis problem into the problem of checking emptiness for universal co-Büchi tree automata using an information-set construction. Then we build a safety game that can be solved using an antichain-based symbolic technique exploiting the structure of the underlying automata. The technique is implemented and applied to a couple of case studies.

## 1 Introduction

The goal of system verification is to check that a system satisfies a given property. One of the major achievements in system verification is the theory of *model checking*, that uses automata-based techniques to check properties expressed in temporal logics, for systems modelled as transitions systems. The *synthesis* problem is more ambitious: given a specification of the system, the aim is to automatically synthesise a system that fulfils the constraints defined by the specification. Therefore, the constraints do not need to be checked a posteriori, and this allows the designer to focus on defining high-level specifications, rather than designing complex computational models of the systems.

*Reactive systems* are non-terminating systems that interact with some environment, e.g., hardware or software that control transportations systems, or medical devices. One of the main challenge of synthesis of reactive systems is to cope with the uncontrollable behaviour of the environment, which usually leads to computationally harder decision problems, compared to system verification. For instance, model-checking properties expressed in *linear time temporal logic* (LTL) is PSpace-c while LTL synthesis is 2Exptime-c [16]. Synthesis of reactive systems from temporal specifications has gain a lot of interest recently as several works have shown its practical feasibility [13,4,3,11,8]. These progresses were supported by Kupferman and Vardi's breakthrough in automata-based synthesis techniques [13]. More precisely, they have shown that the complex Safra's determinization operation, used in the classical LTL synthesis algorithm [16], could be avoided by working directly with universal co-Büchi automata. Since then, several other "Safraless" procedures have been defined [13,18,9,8]. In [18,8], it is shown

<sup>⊗</sup> Work partially supported by the ANR research project "EQINOCS" no. ANR-11-BS02-0004

<sup>⊗⊗</sup> F.R.S.-FNRS research associate (chercheur qualifié)

that LTL synthesis reduces to testing the emptiness of a universal co-Büchi tree automaton, that in turn can be reduced to solving a safety game. The structure of the safety games can be exploited to define a symbolic game solving algorithm based on compact antichain representations [8].

In these works, the system is assumed to have perfect information about the state of the environment. However in many practical scenarios, this assumption is not realistic since some environment information may be hidden to the system (e.g. private variables). Towards the (more realistic) synthesis of partially informed systems, imperfect information two-player games on graphs have been studied [17,6,2,7]. However, they consider explicit state transition systems rather than synthesis from temporal specifications. Moreover, the winning objectives that they consider cannot express fine properties about imperfect information, i.e., cannot speak about knowledge.

*Epistemic Temporal Logics* [10] are logics formatted for reasoning about multi-agent situations. They are extensions of temporal logics with knowledge operators  $K_i$  for each agent. They have been successfully used for verification of various distributed systems in which the knowledge of the agents is essential for the correctness of the system specification.

*Synthesis problem with temporal epistemic objectives* Vardi and van der Meyden [19] have considered epistemic temporal logics to define specifications that can, in addition to temporal properties, also express properties that refer to the imperfect information, and they studied the synthesis problem. They define the synthesis problem in a multi-agent setting, for specifications written in LTL extended with knowledge operators  $K_i$  for each agent (KLTL). In such models, transitions between states of the environment model depend on actions of the environment and the system. The system does not see which actions are played by the environment but get some observation on the states in which the environment can be (observations are subsets of states). An execution of the environment model, from the point of view of the system, is therefore an infinite sequence alternating between its own actions and observations.

The goal of the KLTL synthesis problem is to automatically generate a strategy for the system (if it exists) that tells it which action should be played, depending on finite histories, so that whatever the environment does, all the (concrete) infinite executions resulting from this strategy satisfy the KLTL formula. In [19], this problem was shown to be undecidable even for two agents against the environment. On the other hand, for single agent against environment situations, they show that the problem is  $2\text{ExpTime-c}$ , by reduction to the emptiness of alternating Büchi automata. This theoretically elegant construction is however difficult to implement and optimize, as it relies on complex Safra-like automata operations (Muller-Schupp construction).

*Contributions* In this paper, we follow the formalisation of [19] and, as our main contribution, define and implement a Safriless synthesis procedure for the positive fragment of KLTL (KLTL<sup>+</sup>), i.e., KLTL formulas where the operator  $K$  does not occur under an odd number of negations. Our procedure relies on universal co-Büchi tree automata (UCT). More precisely, given a KLTL<sup>+</sup> formula  $\varphi$  and some environment model  $\mathcal{M}_E$ , we show how to construct a UCT  $\mathcal{T}_\varphi$  whose language is exactly the set of strategies that realize  $\varphi$  in  $\mathcal{M}_E$ .

Despite the fact that our procedure has 2-ExpTime worst-case complexity, we have implemented it and shown its practical feasibility through a set of examples. In particular, based on ideas of [8], we reduce the problem of checking the emptiness of  $\mathcal{T}_\varphi$  to solving a safety game whose state space can be ordered and compactly represented by antichains. Moreover, rather than using the reduction of [8] as a blackbox, we further optimize the antichain representations to improve their compactness. Our implementation is based on the tool Acacia [5] and, to the best of our knowledge, it is the first implementation of a synthesis procedure for epistemic temporal specifications. As an application, this implementation can be used to solve two-player games of imperfect information whose objectives are given as LTL formulas, or universal co-Büchi automata.

*Organization of the paper* In Section 2, we define the KLTL synthesis problem. In Section 3, we define universal co-Büchi automata for infinite words and trees. In Section 4, we consider the particular case of LTL synthesis in an environment model with imperfect information. The construction explained in that section will be used in the generalization to KLTL<sup>+</sup> and moreover, it can be used to solve two-player imperfect information games with LTL (and more generally  $\omega$ -regular) objectives. In Section 5, we define our Safraless procedure for KLTL<sup>+</sup>, and show in Section 6 how to implement it with antichain-based symbolic techniques. Finally, we describe our implementation in Section 7. *Full proofs can be found in Appendix in which, for self-containment, we also explain the reduction to safety games.*

## 2 KLTL Realizability and Synthesis

In this section, we define the realizability and synthesis problems for KLTL specifications, for one partially informed agent, called the *system*, against an environment.

**Environment Model** We assume to have, as input of the problem, a model of the behaviour of the environment as a transition system. This transition system is defined over two disjoint sets of actions  $\Sigma_1$  and  $\Sigma_2$ , for the system and the environment respectively. The transition relation from states to states is defined with respect to pairs of actions in  $\Sigma_1 \times \Sigma_2$ . Additionally, each state  $s$  of the environment model carries an interpretation  $\tau_e(s)$  over a (finite) set of propositions  $\mathcal{P}$ . However, the system is not perfectly informed about the value of some propositions, i.e., some propositions are visible to the system, and some are not. Therefore, we partition the set  $\mathcal{P}$  into two sets  $\mathcal{P}_v$  (the visible propositions) and  $\mathcal{P}_i$  (the invisible ones).

An *environment model* is a tuple  $\mathcal{M}_E = (\mathcal{P}, \Sigma_1, \Sigma_2, S_e, S_0, \Delta_e, \tau_e)$  where

- $\mathcal{P}$  is a finite set of propositions,  $\Sigma_1$  and  $\Sigma_2$  are finite set of actions for the system and the environment resp.,
- $S_e$  is a set of states,  $S_0 \subseteq S_e$  a set of initial states,
- $\tau_e : S_e \rightarrow 2^{\mathcal{P}}$  is a labelling function,
- $\Delta_e \subseteq S_e \times \Sigma_1 \times \Sigma_2 \times S_e$  is a transition relation.

The model is assumed to be deadlock-free, i.e. from any state, there exists at least one outgoing transition. Moreover, the model is assumed to be complete for all actions of the system, i.e. for all states and all actions of the system, there exists an outgoing transition. The set of *executions* of  $\mathcal{M}_E$ , denoted by  $\text{exec}(\mathcal{M}_E)$ , is the set of infinite sequences of states  $\rho = s_0 s_1 \dots \in S_e^\omega$  such that  $s_0 \in S_0$  and for all




Fig. 1: Environment model  $\mathcal{M}_E$  of Example 1

$i > 0$ ,  $(s_i, a_1, a_2, s_{i+1}) \in \Delta_e$  for some  $(a_1, a_2) \in \Sigma_1 \times \Sigma_2$ . Given a sequence of states  $\rho = s_0 s_1 \dots$  and a set  $P \subseteq \mathcal{P}$ , we denote by  $\text{trace}_P(\rho)$  its projection over  $P$ , i.e.  $\text{trace}_P(\rho) = (\tau_e(s_0) \cap P)(\tau_e(s_1) \cap P) \dots$ . The *visible trace* of  $\rho$  is defined by  $\text{trace}_v(\rho) = \text{trace}_{\mathcal{P}_v}(\rho)$ . The *language* of  $\mathcal{M}_E$  with respect to  $P$  is defined as  $\mathcal{L}_P(\mathcal{M}_E) = \{\text{trace}_P(\rho) \mid \rho \in \text{exec}(\mathcal{M}_E)\}$ . The language of  $\mathcal{M}_E$  is defined as  $\mathcal{L}_{\mathcal{P}}(\mathcal{M}_E)$ . The visible language of  $\mathcal{M}_E$  is defined as  $\mathcal{L}_{\mathcal{P}_v}(\mathcal{M}_E)$ . Finally, given an infinite sequence of actions  $a = a_1^0 a_2^0 \dots \in (\Sigma_1 \cdot \Sigma_2)^\omega$  and an execution  $\rho = s_0 s_1 \dots$  of  $\mathcal{M}_E$ , we say that  $a$  is compatible with  $\rho$  if for all  $i \geq 0$ ,  $(s_i, a_1^i, a_2^i, s_{i+1}) \in \Delta_e$ .

This formalization is very close to that of [19]. However in [19], partial observation is modeled as a partition of the state space. The two models are equivalent. In particular, we will see that partitioning the propositions into visible and invisible ones also induces a partition of the state space into observations.

*Example 1.* We illustrate the notion of environment model on the example of [19], that describes the behaviour of an environment against a system acting on a timed toggle switch with two positions (on,off) and a light. It is depicted in Fig. 1. The set  $\mathcal{P} = \{t, l\}$  contains two propositions  $t$  (true iff the toggle is on) and  $l$  (true iff the light is on). Actions of the system are  $\Sigma_1 = \{T, S\}$  for “toggle” and “skip” respectively. The system can change the position of the toggle only if it plays  $T$ , and  $S$  has no effect. Actions of the environment are  $\Sigma_2 = \{(t_{out}, l_{on}) \mid t_{out}, l_{on} \in \{0, 1\}\}$ . The boolean variables ( $t_{out}$  and  $l_{on}$ ) indicate that the environment times out the toggle and that it switches on the light. The transition function is depicted on the figure as well as the labelling function  $\tau_e : S_e \rightarrow 2^{\mathcal{P}}$ . The star  $*$  means “any action”. The light can be on only if the toggle is on (state  $s_1$ ), but it can be off even if the toggle is on (state  $s_2$ ), in case it is broken. This parameter is uncontrollable by the system, and therefore it is controlled by the environment (action  $l_{on}$ ). The timer is assumed to be unreliable and therefore the environment can timeout at any time (action  $t_{out}$ ). The system sees only the light, i.e.  $\mathcal{P}_v = \{l\}$  and  $\mathcal{P}_i = \{t\}$ . The goal of the system is to have a strategy such that he always knows the position of the toggle.

**Observations.** The partition of the set of propositions  $\mathcal{P}$  into a set of visible propositions  $\mathcal{P}_v$  and a set of invisible propositions  $\mathcal{P}_i$  induces an indistinguishability relation over the states  $S_e$ . Two states are *indistinguishable*, denoted  $s_1 \sim s_2$ , if they have the same visible propositions, i.e.  $\tau_e(s_1) \cap \mathcal{P}_v = \tau_e(s_2) \cap \mathcal{P}_v$ . It is easy to see that  $\sim$  is an equivalence relation over  $S_e$ . Each equivalence class of  $S_e$  induced by  $\sim$  is called

an *observation*. The equivalence class of a state  $s \in S_e$  is denoted by  $o(s)$  and the set of observations is denoted by  $\mathcal{O}$ . The relation  $\sim$  is naturally extended to (finite or infinite) executions:  $\rho_1 \sim \rho_2$  if  $\text{trace}_v(\rho_1) = \text{trace}_v(\rho_2)$ . Similarly, two executions  $\rho = s_0 s_1 \dots$  and  $\rho' = s'_0 s'_1 \dots$  are said to be indistinguishable *up to some position  $i$*  if  $\text{trace}_v(s_0 \dots s_i) = \text{trace}_v(s'_0 \dots s'_i)$ . This indistinguishability notion also is an equivalence relation over executions that we denote by  $\sim_i$ .

Coming back to Example 1, since the set of visible propositions is  $\mathcal{P}_v = \{l\}$  and the set of invisible ones is  $\mathcal{P}_i = \{t\}$ , the states  $s_2$  and  $s_3$  are indistinguishable (in both  $s_2, s_3$  the light is off) and therefore  $\mathcal{O} = \{o_0, o_1\}$  with  $o_0 = \{s_2, s_3\}$  and  $o_1 = \{s_1\}$ .

Given an infinite sequence  $u = a_1 o_1 a_2 o_2 \dots \in (\Sigma_1 \cdot \mathcal{O})^\omega$  of actions of Player 1, and observations, we associate with  $u$  the set of possible executions of  $\mathcal{M}_E$  that are compatible with  $u$ . Formally, we define  $\text{exec}(\mathcal{M}_E, u)$  the set of executions  $\rho = s_0 s_1 \dots \in \text{exec}(\mathcal{M}_E)$  such that for all  $i \geq 1$ ,  $o(s_i) = o_i$  and there exists an action  $b_i$  of the environment such that  $(s_{i-1}, a_i, b_i, s_i) \in \Delta_e$ . We also define the traces of  $u$  as the set of traces of all executions of  $\mathcal{M}_E$  compatible with  $u$ , i.e.  $\text{traces}(u) = \{\text{traces}(\rho) \mid \rho \in \text{exec}(\mathcal{M}_E, u)\}$ .

**Epistemic Linear Time Temporal Logic (KLTL)** We now define the logic KLTL for one-agent (the system). The logic KLTL extends the logic LTL with an epistemic operator  $K\phi$ , modelling the property that the system knows that the formula  $\phi$  holds. *KLTL formulae* are defined over the set of atomic propositions  $\mathcal{P}$  by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \bigcirc\varphi \mid \varphi_1 \mathcal{U} \varphi_2 \mid K\varphi$$

in which  $p \in \mathcal{P}$  and  $\bigcirc$  and  $\mathcal{U}$  are the "next" and "until" operators from linear temporal logic. Formulas of the type  $K\varphi$  are read as "the system knows that  $\varphi$  holds". We define the macros  $\diamond$  (eventually) and  $\square$  (always) as usual. LTL is the fragment of KLTL without the  $K$  operator.

The semantics of a KLTL formula  $\varphi$  is defined for an environment model  $\mathcal{M}_E = (\mathcal{P}, \Sigma_1, \Sigma_2, S_e, S_0, \Delta_e, \tau_e)$ , a set of executions  $R \subseteq \text{exec}(\mathcal{M}_E)$ , an execution  $\rho = s_0 s_1 \dots \in R$  and a position  $i \geq 0$  in  $\rho$ . It is defined inductively:

- $R, \rho, i \models p$  if  $p \in \tau_e(s_i)$ ,
- $R, \rho, i \models \neg\varphi$  if  $R, \rho, i \not\models \varphi$ ,
- $R, \rho, i \models \varphi_1 \vee \varphi_2$  if  $R, \rho, i \models \varphi_1$  or  $R, \rho, i \models \varphi_2$ ,
- $R, \rho, i \models \bigcirc\varphi$  if  $R, \rho, i+1 \models \varphi$ ,
- $R, \rho, i \models \varphi_1 \mathcal{U} \varphi_2$  if  $\exists j \geq i$  s.t.  $R, \rho, j \models \varphi_2$  and  $\forall i \leq k < j$ ,  $R, \rho, k \models \varphi_1$ ,
- $R, \rho, i \models K\varphi$  if for all  $\rho' \in R$  s.t.  $\rho \sim_i \rho'$ , we have  $R, \rho', i \models \varphi$ .

In particular, the system knows  $\varphi$  at position  $i$  in the execution  $\rho$ , if all other executions in  $R$  whose prefix up to position  $i$  are indistinguishable from that of  $\rho$ , also satisfy  $\varphi$ . We write  $R, \rho \models \varphi$  if  $R, \rho, 0 \models \varphi$ , and  $R \models \varphi$  if  $R, \rho \models \varphi$  for all executions  $\rho \in R$ . We also write  $\mathcal{M}_E \models \varphi$  to mean  $\text{exec}(\mathcal{M}_E) \models \varphi$ . Note that  $\mathcal{M}_E \models \varphi$  iff  $\mathcal{M}_E \models K\varphi$ .

Consider Example 1 and the set  $R$  of executions that eventually loops in  $s_1$ . Pick any  $\rho$  in  $R$ . Then  $R, \rho, 0 \models \square K \diamond(l)$ . Indeed, take any position  $i$  in  $\rho$  and any other executions  $\rho' \in R$  such that  $\rho \sim_i \rho'$ . Then since  $\rho'$  will eventually loop in  $s_1$ , it will satisfy  $\diamond(l)$ . Therefore  $R, \rho, i \models K \diamond(l)$ , for all  $i \geq 0$ .

**KLTL Realizability and Synthesis** As presented in [8] for the perfect information setting, the realizability problem, given the environment model  $\mathcal{M}_E$  and the KLTL formula  $\varphi$ , is best seen as a turn-based game between the system (Player 1) and the environment (Player 2). In the first round of the play, Player 1 picks some action  $a_1^0 \in \Sigma_1$  and then Player 2 picks some action in  $a_2^0 \in \Sigma_2$  and solves the nondeterminism in  $\Delta_e$ , and a new round starts. The two players play for an infinite duration and the outcome is an infinite sequence  $w = a_1^0 a_2^0 a_1^1 a_2^1 \dots$ . The winning objective is given by some KLTL formula  $\varphi$ . Player 1 wins the play if for all executions  $\rho$  of  $\mathcal{M}_E$  that are compatible with  $w$ , we have  $\mathcal{M}_E, \rho \models \varphi$ .

Player 1 plays according to *strategies* (called *protocols* in [19]). Since Player 1 has only partial information about the state of the environment, his strategies are based on the histories of his own actions and the observations he got from the environment. Formally, a strategy for Player 1 is a mapping  $\lambda : (\Sigma_1 \mathcal{O})^* \rightarrow \Sigma_1$ , where as defined before,  $\mathcal{O}$  denotes the set of observations of Player 1 over the states of  $\mathcal{M}_E$ . Fixing a strategy  $\lambda$  of Player 1 restricts the set of executions of the environment model  $\mathcal{M}_E$ . An execution  $\rho = s_0 s_1 \dots \in \text{exec}(\mathcal{M}_E)$  is said to be *compatible* with  $\lambda$  if there exists an infinite sequence of actions  $a = a_1^0 a_2^0 \dots \in (\Sigma_1 \cdot \Sigma_2)^\omega$ , compatible with  $\rho$  such that for all  $i \geq 0$ ,  $a_1^i = \lambda(a_1^0 o(s_0) a_1^1 o(s_1) \dots a_1^{i-1} o(s_{i-1}))$ . We denote by  $\text{exec}(\mathcal{M}_E, \lambda)$  the set of executions of  $\mathcal{M}_E$  compatible with  $\lambda$ .

**Definition 1.** A KLTL formula  $\varphi$  is *realizable* in  $\mathcal{M}_E$  if there exists a strategy  $\lambda$  for the system such that  $\text{exec}(\mathcal{M}_E, \lambda) \models \varphi$ .

**Theorem 1 ([19]).** The KLTL realizability problem (for one agent) is 2ExpTime-complete.

If a formula is realizable, the *synthesis* problem asks to generate a finite-memory strategy that realizes the formula. Such a strategy always exists if the specification is realizable [19]. Finite memory strategies can be represented by Moore machines that read observations and output actions of Player 1. We refer the reader to [8] for a formal definition of finite-memory strategies.

Considering again Example 1, the formula  $\varphi = \Box(K(t) \vee K(\neg t))$  expresses the fact that the system knows at each step the position of the toggle. As argued in [19], this formula is realizable if the initial set of the environment is  $\{s_1, s_2\}$  since both states are labelled with  $t$ . Then, one strategy of the system is to play first time  $T$ , action that will lead to  $s_3$ , and then always play  $S$  in order to stay in that state. Following this strategy, in the first step the formula  $K(t)$  is satisfied and then  $K(\neg t)$  becomes true. However, the formula is not realizable if the set of initial states of the environment is  $\{s_2, s_3\}$  since from the beginning the system doesn't know the value of the toggle.

### 3 Automata for Infinite Words and Trees

**Automata on infinite Words** An *infinite word automaton* over some (finite) alphabet  $\Sigma$  is a tuple  $A = (\Sigma, Q, Q_0, \Delta, \alpha)$  where  $\Sigma$  is the finite input alphabet,  $Q$  is the finite set of states,  $Q_0 \subseteq Q$  is the set of initial states,  $\alpha \subseteq Q$  is the set of final states (accepting states) and  $\Delta \subseteq Q \times \Sigma \times Q$  is the transition relation.

For all  $q \in Q$  and all  $\sigma \in \Sigma$ , we let  $\Delta(q, \sigma) = \{q' \mid (q, \sigma, q') \in \Delta\}$ . We let  $|A| = |Q| + |\Delta|$ . We say that  $A$  is *deterministic* if  $|Q_0| = 1$  and  $\forall q \in Q, \forall \sigma \in \Sigma$

$\Sigma$ ,  $|\Delta(q, \sigma)| \leq 1$ . It is *complete* if  $\forall q \in Q, \forall \sigma \in \Sigma, \Delta(q, \sigma) \neq \emptyset$ . In this paper we assume, w.l.o.g., that the word automata are always complete.

A *run* on the automaton  $A$  over an infinite input word  $w = w_0w_1w_2\dots$ , is a sequence  $r = q_0q_1q_2\dots \in Q^\omega$  such that  $(q_i, w_i, q_{i+1}) \in \Delta$  for all  $i \geq 0$  and  $q_0 \in Q_0$ . We denote by  $Runs_A(w)$  the set of runs of  $A$  on  $w$  and by  $Visit(r, q)$  the number of times the state  $q$  is visited along the run  $r$  (or  $\infty$  if the path visit the state  $q$  infinitely often). Here, we consider two accepting conditions for infinite word automata and name the infinite word automata according to the used accepting condition. Let  $B \in \mathbb{N}$ . A word  $w \in \Sigma^\omega$  is accepted by  $A$  if (according to the accepting condition):

$$\begin{aligned} \text{Universal Co-Büchi} & : \forall r \in Runs_A(w), \forall q \in \alpha, Visit(r, q) < \infty \\ \text{Universal B-Co-Büchi} & : \forall r \in Runs_A(w), \forall q \in \alpha, Visit(r, q) \leq B \end{aligned}$$

The set of words accepted by  $A$  with the universal co-Büchi (resp.  $B$ -co-Büchi) accepting condition is denoted by  $\mathcal{L}_{uc}(A)$  (resp.  $\mathcal{L}_{uc,B}(A)$ ). We say that  $A$  is a *universal co-Büchi word automaton* (UCW) if the first acceptance condition is used and that  $(A, B)$  is an *universal  $B$ -co-Büchi word automaton* (UBCW) if the second one is used.

Given an LTL formula  $\varphi$ , we can translate it into an equivalent universal co-Büchi word automaton  $A_\varphi$ . This can be done with a single exponential blow-up by first negating  $\varphi$ , then translating  $\neg\varphi$  into an equivalent nondeterministic Büchi word automaton, and then dualize it into a universal co-Büchi word automaton [8,13].

**Automata on Infinite Trees** Given a finite set  $D$  of directions, a  $D$ -tree is a prefix-closed set  $T \subseteq D^*$ , i.e., if  $x \cdot d \in T$ , where  $d \in D$ , then  $x \in T$ . The elements of  $T$  are called *nodes* and the empty word  $\epsilon$  is the *root* of  $T$ . For every  $x \in T$ , the nodes  $x \cdot d$ , for  $d \in D$ , are the *successors* of  $x$ . A node  $x$  is a *leaf* if has no successor in  $T$ , formally,  $\forall d \in D, x \cdot d \notin T$ . The tree  $T$  is *complete* if for all nodes, there are successors in all directions, formally,  $\forall x \in T, \forall d \in D, x \cdot d \in T$ . Finite and infinite branches  $\pi$  in a tree  $T$  are naturally defined, respectively, as finite and infinite paths in  $T$  starting from the root node. Given an alphabet  $\Sigma$ , a  $\Sigma$ -labelled  $D$ -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  $\Sigma$ . We omit  $\tau$  when it is clear from the context. Then, in a tree  $T$ , an infinite (resp. finite) branch  $\pi$  induces an infinite (resp. finite) sequence of labels and directions in  $(\Sigma \cdot D)^\omega$  (resp.  $(\Sigma \cdot D)^* \Sigma$ ). We denote this sequence by  $\tau(\pi)$ . For instance, for a set of system's actions  $\Sigma_1$  and a set of observations  $\mathcal{O}$ , a strategy  $\lambda : (\Sigma_1 \mathcal{O})^* \rightarrow \Sigma_1$  of the system can be seen has a  $\Sigma_1$ -labelled  $\mathcal{O}$ -tree whose nodes are finite outcomes<sup>3</sup>.

A *universal co-Büchi tree automaton* (UCT) is a tuple  $\mathcal{T} = (\Sigma, Q, Q_0, D, \Delta, \alpha)$  where  $\Sigma$  is the finite alphabet,  $Q$  is a finite set of states,  $Q_0 \subseteq Q$  is the set of initial states,  $D$  is the set of directions,  $\Delta : Q \times \Sigma \times D \rightarrow 2^Q$  is the transition relation (assumed to be total) and  $\alpha$  is the set of final states. If the tree automaton is in some state  $q$  at some node  $x$  labelled by some  $\sigma \in \Sigma$ , it will evaluate, for all  $d \in D$ , the subtree rooted at  $x \cdot d$  in parallel from all the states of  $\Delta(q, \sigma, d)$ . Let us define the notion of run formally.

<sup>3</sup> Technically, a strategy  $\lambda$  is defined also for histories that are not accessible by  $\lambda$  itself from the initial (empty) history  $\epsilon$ . The tree represents only accessible histories but we can, in the rest of the paper, assume that strategies are only defined for their accessible histories. Formally, we assume that a strategy is a partial function whose domain  $H$  satisfies  $\epsilon \in H$  and for all  $h \in H$  and all  $o \in \mathcal{O}$ ,  $h \cdot \lambda(h) \cdot o \in H$ , and  $H$  is minimal (for inclusion) w.r.t. this property.

For all  $q \in Q$  and  $\sigma \in \Sigma$ , we denote by  $\Delta(q, \sigma) = \{(q_1, d_1), \dots, (q_n, d_n)\}$  the *disjoint* union of all sets  $\Delta(q, \sigma, d)$  for all  $d \in D$ . A *run* of  $\mathcal{T}$  on an infinite  $\Sigma$ -labelled  $D$ -tree  $\langle T, \tau \rangle$  is a  $(Q \times D^*)$ -labelled  $\mathbb{N}$ -tree  $\langle T_r, \tau_r \rangle$  such that  $\tau_r(\epsilon) \in Q_0 \times \{\epsilon\}$  and, for all  $x \in T_r$  such that  $\tau_r(x) = (q, v)$ , if  $\Delta(q, \tau(v)) = \{(q_1, d_1), \dots, (q_n, d_n)\}$ , we have  $x \cdot i \in T_r$  and  $\tau_r(x \cdot i) = (q_i, v \cdot d_i)$  for all  $0 < i \leq n$ . Note that there is at most one run per input tree (up to tree isomorphism). A run  $\langle T_r, \tau_r \rangle$  is *accepting* if for all infinite branches  $\pi$  of  $T_r$ ,  $\tau_r(\pi)$  visits a finite number of accepting states. The language of  $\mathcal{T}$ , denoted by  $\mathcal{L}_{uc}(\mathcal{T})$ , is the set of  $\Sigma$ -labelled  $D$ -trees such that there exists an accepting run on them. Similarly, we define universal  $B$ -co-Büchi tree automata by strengthening the acceptance conditions on all branches to the  $B$ -co-Büchi condition.

As noted in [8,18], testing the emptiness of a UCT automaton reduces to testing the emptiness of a universal  $B$ -co-Büchi accepting condition for a sufficiently large bound  $B$ , which in turn reduces to solving a safety game. Symbolic techniques that are also exploited in this paper have been used to solve the safety games.

## 4 LTL synthesis under imperfect information

In this section, we first explain an automata-based procedure to decide realizability under imperfect information of LTL formulas against an environment model. This procedure will be extended in the next section to handle the  $K$  operator.

Take an environment model  $\mathcal{M}_E = (\mathcal{P}, \Sigma_1, \Sigma_2, S_e, S_0, \Delta_e, \tau_e)$ . Then, a complete  $\Sigma_1$ -labelled  $\mathcal{O}$ -tree  $\langle T, \tau \rangle$  defines a strategy of the system. Any infinite branch  $\pi$  of  $\langle T, \tau \rangle$  defines an infinite sequence of actions and observations of  $\mathcal{M}_E$ , which in turn corresponds to a set of possible traces in  $\mathcal{M}_E$ . We denote by  $\text{traces}(\pi)$  this set of traces, and it is formally defined by  $\text{traces}(\pi) = \text{traces}(\tau(\pi))$  (recall that the set of traces of a sequence of actions and observations has been defined in Section 2).

Given an LTL formula  $\psi$ , we construct a universal co-Büchi tree automaton  $\mathcal{T} = (\Sigma_1, Q, Q_0, \mathcal{O}, \Delta, \alpha)$  that accepts all the strategies of Player 1 (the system) that realize  $\psi$  under the environment model  $\mathcal{M}_E$ . First, one converts  $\psi$  into an equivalent UCW  $\mathcal{A} = (2^{\mathcal{P}}, Q^A, Q_0^A, \Delta^A, \alpha^A)$ . Then, as a direct consequence of the definition of KLTL realizability:

**Proposition 1.** *Given a complete  $\Sigma_1$ -labelled  $\mathcal{O}$ -trees  $\langle T, \tau \rangle$ ,  $\langle T, \tau \rangle$  defines a strategy that realizes  $\psi$  under  $\mathcal{M}_E$  iff for all infinite branches  $\pi$  of  $\langle T, \tau \rangle$  and all traces  $\rho \in \text{traces}(\pi)$ ,  $\rho \in L(\mathcal{A})$ .*

We now show how to construct a universal tree automaton that checks the property mentioned in the previous proposition, for all branches of the trees. We use universal transitions to check, on every branch of the tree, that all the possible traces (possibly uncountably many) compatible with the sequence of actions in  $\Sigma_1$  and observations in  $\mathcal{O}$  defined by the branch satisfy  $\psi$ . Based on finite sequences of observations that the system has received, it can define its *knowledge*  $I$  of the possible states in which the environment can be as a subset of states of  $S_e$ . Given an action  $a_1 \in \Sigma_1$  of the system and some observation  $o \in \mathcal{O}$ , we denote by  $\text{post}_a(I, o)$  the new knowledge that the system can infer from the observation  $o$ , the action  $a$  and its previous information  $I$ . Formally,  $\text{post}_a(I, o) = \{s \in S_e \cap o \mid \exists a_2 \in \Sigma_2, \exists s' \in I \text{ s.t. } (s', a, a_2, s) \in \Delta_e\}$ .

The states of the universal tree automaton  $\mathcal{T}$  are pairs of states of  $\mathcal{A}$  and knowledges plus some extra state  $(q_w, \emptyset)$ , i.e.  $Q = Q^A \times 2^{S_e} \cup \{(q_w, \emptyset)\}$  where  $(q_w, \emptyset)$  is added



for completeness. The final states are defined as  $\alpha = \alpha^A \times 2^{S_e}$  and initial states as  $Q_0 = Q_0^A \times S_0$ . To define the transition relation, let us consider a state  $q \in Q^A$ , a knowledge set  $I \subseteq S_e$ , an action  $a \in \Sigma_1$  and some observation  $o \in \mathcal{O}$ . We now define  $\Delta((q, I), a, o)$ . It could be the case that there is no transition in  $\mathcal{M}_E$  from a state of  $I$  to a state of  $o$ , i.e.  $\text{post}_a(I, o) = \emptyset$ . In that case, all the paths from the next  $o$ -node of the tree should be accepting. This situation is modelled by going to the extra state  $(q_w, \emptyset)$ , i.e.  $\Delta((q, I), a, o) = (q_w, \emptyset)$ .

Now suppose that  $\text{post}_a(I, o)$  is non-empty. Since the automaton must check that all the traces of  $\mathcal{M}_E$  that are compatible with actions of  $\Sigma_1$  and observations are accepted by  $\mathcal{A}$ , intuitively, one would define  $\Delta((q, I), a, o)$  as the set of states of the form  $(q', \text{post}_a(I, o))$  for all states  $q'$  such that there exists  $s \in I$  such that  $(q, \tau_e(s), q') \in \Delta^A$ . However, it is not correct for several reasons. First, it could be that  $s$  has no successor in  $o$  for action  $a$ , and therefore one should not consider it because the traces up to state  $s$  die at the next step after getting observation  $o$ . Therefore, one should only consider states of  $I$  that have a successor in  $o$ . Second, it is not correct to associate the new knowledge  $\text{post}_a(I, o)$  with  $q'$ , because it could be that there exists a state  $s' \in \text{post}_a(I, o)$  such that for all its predecessors  $s''$  in  $I$ , there is not transition  $(q, \tau_e(s''), q')$  in  $\Delta^A$ , and therefore, one would also take into account sequences of interpretations of propositions that do not correspond to any trace of  $\mathcal{M}_E$ .

Taking into account these two remarks, we define, for all states  $q'$ , the set  $I_{q, q'} = \{s \in I \mid (q, \tau(s), q') \in \Delta^A\}$ . Then,  $\Delta((q, I), a, o)$  is defined as the set

$$\Delta((q, I), a, o) = \{(q', \text{post}_a(I_{q, q'}, o)) \mid \exists s \in I, (q, \tau(s), q') \in \Delta^A\}$$

Note that, since  $\bigcup_{q' \in Q^A} \text{post}_a(I_{q, q'}, o) = \text{post}_a(I, o)$  and the automaton is universal, the system does not have better knowledge by restricting the knowledge sets.

**Lemma 1.** *The LTL formula  $\psi$  is realizable in  $\mathcal{M}_E$  iff  $\mathcal{L}(\mathcal{T}) \neq \emptyset$ .*

*Proof.* If  $\psi$  is realizable in  $\mathcal{M}_E$ , there exists a strategy  $\lambda : (\Sigma_1 \times \mathcal{O})^* \rightarrow \Sigma_1$  for the system such that  $\text{exec}(\mathcal{M}_E, \lambda) \models \psi$ . Let's see this strategy as a  $\Sigma_1$ -labelled  $\mathcal{O}$ -tree  $\langle T_\lambda, \tau \rangle$  and prove that  $T_\lambda \in \mathcal{L}(\mathcal{T})$ .

The run on  $T_\lambda$  of the automaton  $\mathcal{T}$  is a  $Q \times \mathcal{O}^*$ -labelled  $\mathbb{N}$ -tree  $\langle T_r, \tau_r \rangle$ . Therefore, each branch  $\pi$  of  $T_r$  induces an infinite sequence  $\tau_r(\pi) = ((q_0, I^{(0)}), \epsilon)((q_1, I^{(1)}), o_1)((q_2, I^{(2)}), o_1 o_2) \dots$  where  $(q_0, I^{(0)}) \in Q_0$  and  $I^{(i+1)} = \text{post}_{\tau(o_1 \dots o_i)}(I_{q_i, q_{i+1}}, o_{i+1})$ , by definition of the transition relation of  $\mathcal{T}$ .

Since  $I_{q_i, q_{i+1}}^{(i)}$  is the subset of environment state in  $I^{(i)}$  whose labels can fire transitions from  $q_i$  to  $q_{i+1}$  in  $\mathcal{A}$ ,  $r = q_0 q_1 q_2 \dots$  is a run in  $\mathcal{A}$  on the traces of the executions  $\rho = s_0 s_1 s_2 \dots$  where  $s_i \in I_i$ . Hence, from the fact that  $\text{exec}(\mathcal{M}_E, \lambda) \models \psi$  and  $\rho \in \text{exec}(\mathcal{M}_E, \lambda)$ ,  $r$  is an accepting run and then  $\tau_r(\pi)$  visits a finite number of accepting states in  $\mathcal{T}$ . Therefore,  $T_\lambda \in \mathcal{L}(\mathcal{T})$  and  $\mathcal{L}(\mathcal{T}) \neq \emptyset$ .

In the other direction, if  $\mathcal{L}(\mathcal{T}) \neq \emptyset$ , there exists a  $\Sigma_1$ -labelled  $\mathcal{O}$ -tree  $\langle T_\lambda, \tau \rangle$  such that  $T_\lambda \in \mathcal{L}(\mathcal{T})$ . We have to prove that the strategy  $\lambda$  defined by  $T_\lambda$  realizes  $\psi$ , i.e., for all executions  $\rho \in \text{exec}(\mathcal{M}_E, \lambda)$ , we have  $\text{exec}(\mathcal{M}_E, \lambda), \rho, 0 \models \psi$ . Let  $\rho \in \text{exec}(\mathcal{M}_E, \lambda)$ . Since  $\psi$  is an LTL formula, the sets of executions in which  $\psi$  is evaluated does not matter, and therefore we have to prove that  $\rho, 0 \models \psi$  (where  $\models$  here denotes the classical LTL semantics). In other words, we have to show that

$\text{trace}(\rho) \in \mathcal{L}(\mathcal{A})$ , i.e. all the runs of  $\mathcal{A}$  on  $\text{trace}(\rho)$  visit finitely many accepting states. Let  $r = q_0q_1 \dots$  be a run of  $\mathcal{A}$  on  $\text{trace}(\rho)$ .

Let  $\langle T_r, \tau_r \rangle$  be the (accepting) run of  $\mathcal{T}$  on  $\langle T_\lambda, \tau \rangle$ . It is a  $Q \times \mathcal{O}^*$ -labelled N-tree. By definition of  $\text{exec}(\mathcal{M}_E, \lambda)$ , there exists a sequence  $u = a_0o_0 \dots \in \Sigma_1 \times \mathcal{O}$  such that  $\rho = s_0s_1s_2\dots$  is compatible with  $u$ . Let define  $t$  as the infinite sequence  $t = (q_0, I^{(0)}, \epsilon)((q_1, I^{(1)}), o_1)((q_2, I^{(2)}), o_1o_2)\dots$  such that  $s_i \in I^{(i)} \subseteq S$  for all  $i \geq 0$ ,  $I^{(0)} = (q_0, S_0)$  and  $I^{(i+1)} = \text{post}_{a_i}(I_{q_i, q_{i+1}}^{(i)}, o_{i+1})$  for all  $i \geq 0$ . It is easily shown that this sequence is such that there exists a branch  $\pi$  in  $\langle T_r, \tau_r \rangle$  with  $\tau_r(\pi) = t$  since the automata  $\mathcal{A}$  and  $\mathcal{T}$  are complete and  $\forall i, \bigcup_{q' \in Q} \text{post}_a(I_{q_i, q'}^{(i)}, o) = \text{post}_a(I^{(i)}, o)$ . Therefore there are finitely many accepting states in  $q_0q_1 \dots$ , since  $T_r$  is accepting. Therefore  $\text{trace}(\rho) \in \mathcal{L}(\mathcal{A})$ .  $\square$

Moreover, it is known that if a UCT has a non-empty language, then it accepts a tree that is the unfolding of a finite graph, or equivalently, that can be represented by a Moore machine. Therefore if  $\psi$  is realizable, it is realizable by a finite-memory strategy. In this paper we will also use the notation  $\mathcal{T}_{\psi, X}$  for the UCT built for the LTL formula  $\psi$  where the executions of  $\mathcal{M}_E$  start from the set  $X \subseteq S_e$ , i.e.,  $Q_0 = Q_0^A \times X$ .

## 5 Safriless procedure for positive KLTL synthesis

In this section, we extend the construction of Section 4 to the positive fragment of KLTL. Positive formulas are defined by the following grammar:

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc \varphi \mid \square \varphi \mid K\varphi \mid \varphi \mathcal{U} \varphi$$

Note that this fragment is equivalent to the fragment of KLTL in which formulas with the knowledge operator  $K$  occurring under an even number of negations. This is obtained by straightforwardly pushing the negations down towards the atoms. We denoted this fragment of KLTL by  $\text{KLTL}^+$ .

**Sketch of the construction** Given a  $\text{KLTL}^+$  formula  $\varphi$  and an environment model  $\mathcal{M}_E = (\mathcal{P}, \Sigma_1, \Sigma_2, S_e, S_0, \Delta_e, \tau_e)$ , we show how to construct a UCT  $\mathcal{T}$  such that  $\mathcal{L}(\mathcal{T}) \neq \emptyset$  iff  $\varphi$  is realizable in  $\mathcal{M}_E$ . The construction is compositional and follows, for the basic blocks, the construction of Section 4 for LTL formulas. The main idea is to replace subformulas of the form  $K\gamma$  by fresh atomic propositions  $k_\gamma$  so that we get an LTL formula for which the realizability problem can be transformed into the emptiness of a UCT. The realizability of the subformulas  $K\gamma$  that have been replaced by  $k_\gamma$  is checked by branching universally to a UCT for  $\gamma$ , constructed as in Section 4. Since transitions are universal, this will ensure that all the infinite branches of the tree from the current node where a new UCT has been triggered also satisfy  $\gamma$ . The UCTs we construct are defined over an extended alphabet that contains the new atomic propositions, but we show that we can safely project the final UCT on the alphabet  $\Sigma_1$ . The assumption on positivity of KLTL formulas implies that there is no subformulas of the form  $\neg K\gamma$ . The rewriting of subformulas by fresh atomic propositions cannot be done in any order. We now describe it formally.

We inductively define a sequence of formulas associated with  $\varphi$  as:  $\varphi^0 = \varphi$  and, for all  $i > 0$ ,  $\varphi^i$  is the formula  $\varphi^{i-1}$  in which the *innermost* subformulas  $K\gamma$  are replaced by fresh atomic propositions  $k_\gamma$ . Let  $d$  be the smallest index such that  $\varphi^d$  is

an LTL formula (in other words,  $d$  is the maximal nesting level of  $K$  operators). Let  $\mathbb{K}$  denote the set of new atomic propositions, i.e.,  $\mathbb{K} = \bigcup_{i=0}^d \{k_\gamma \mid K\gamma \in \varphi^i\}$ , and let  $\mathcal{P}' = \mathcal{P} \cup \mathbb{K}$ . Note that by definition of the formulas  $\varphi^i$ , for all atomic proposition  $k_\gamma$  occurring in  $\varphi^i$ ,  $\gamma$  is an LTL formula over  $\mathcal{P}'$ . E.g. if  $\varphi = p \rightarrow K(q \rightarrow Kr \vee Kz)$  and  $\mathcal{P} = \{p, q, r, z\}$ , then the sequence of formulas is:  $\varphi^0 = \varphi$ ,  $\varphi^1 = p \rightarrow K(q \rightarrow k_r \vee k_z)$   $\varphi^2 = p \rightarrow k_\gamma$  where  $\gamma = q \rightarrow k_r \vee k_z$ .

Then, we construct incrementally a chain of universal co-Büchi tree automata  $\mathcal{T}^d, \dots, \mathcal{T}^0$  such that  $\mathcal{L}(\mathcal{T}^d) \supseteq \mathcal{L}(\mathcal{T}^{d-1}) \supseteq \dots \supseteq \mathcal{L}(\mathcal{T}^0)$  and, the following invariant is satisfied: for all  $i \in \{0, \dots, d\}$ ,  $\mathcal{T}^i$  accepts exactly the set of strategies that realize  $\varphi^i$  in  $\mathcal{M}_E$ . Intuitively, the automaton  $\mathcal{T}^i$  is defined by adding new transitions in  $\mathcal{T}^{i+1}$ , such that for all atomic propositions  $k_\gamma$  occurring in  $\varphi^{i+1}$ ,  $\mathcal{T}^i$  will ensure that  $K\gamma$  is indeed satisfied, by branching to a UCT checking  $\gamma$  whenever the atomic proposition  $k_\gamma$  is met. Since formulas  $\varphi^i$  are defined over the extended alphabet  $\mathcal{P}' = \mathcal{P} \cup \mathbb{K}$  and  $\mathcal{M}_E$  is defined over  $\mathcal{P}$ , we now make clear what we mean by realizability of a formula  $\varphi^i$  in  $\mathcal{M}_E$ . It uses the notion of *extended model executions* and *extended strategies*.

**Extended actions, model executions and strategies** We extend the actions of the system to  $\Sigma'_1 = \Sigma_1 \times 2^{\mathbb{K}}$  (call  $e$ -actions). Informally, the system plays an  $e$ -action  $(a, K)$  if it considers formulas  $K\gamma$  for all  $k_\gamma \in K$  to be true. An *extended execution* ( $e$ -execution) of  $\mathcal{M}_E$  is a infinite sequence  $\rho = (s_0, K_0) \dots \in (S_e \times 2^{\mathbb{K}})^\omega$  such that  $s_0 s_1 \dots \in \text{exec}(\mathcal{M}_E)$ . We denote  $s_0 s_1 \dots$  by  $\text{proj}_1(\rho)$  and  $K_0 K_1 \dots$  by  $\text{proj}_2(\rho)$ . The extended labelling function  $\tau'_e$  is a function from  $S_e \times 2^{\mathbb{K}}$  to  $\mathcal{P}'$  defined by  $\tau'_e(s, K) = \tau_e(s) \cup K$ . The indistinguishability relation between extended executions is defined, for any two extended executions  $\rho_1, \rho_2$ , by  $\rho_1 \sim \rho_2$  iff  $\text{proj}_1(\rho_1) \sim \text{proj}_1(\rho_2)$  and  $\text{proj}_2(\rho_1) = \text{proj}_2(\rho_2)$ , i.e., the propositions in  $\mathbb{K}$  are visible to the system. We define  $\sim_i$  over extended executions similarly. Given the extended labelling functions and indistinguishability relation, the KLTL satisfiability notion  $R, \rho, i \models \psi$  can be naturally defined for a set of  $e$ -execution  $R$ ,  $\rho \in R$  and  $\psi$  a KLTL formula over  $\mathcal{P}' = \mathcal{P} \cup \mathbb{K}$ .

An extended strategy is a strategy defined over  $e$ -actions, i.e. a function from  $(\Sigma'_1 \mathcal{O})^*$  to  $\Sigma'_1$ . For an infinite sequence  $u = (a_0, K_0) o_0 (a_1, K_1) o_1 \dots \in (\Sigma'_1 \mathcal{O})^\omega$ , we define  $\text{proj}_1(u)$  as  $a_0 o_0 \dots$ . The sequence  $u$  defines a set of compatible  $e$ -executions  $\text{exec}(\mathcal{M}_E, u)$  as follows: it is the set of  $e$ -executions  $\rho = (s_0, K_0)(s_1, K_1) \dots \in (S_e \times 2^{\mathbb{K}})^\omega$  such that  $\text{proj}_1(\rho) \in \text{exec}(\mathcal{M}_E, \text{proj}_1(u))$ . Similarly, we define for  $e$ -strategies  $\lambda'$  the set  $\text{exec}(\mathcal{M}_E, \lambda')$  of  $e$ -executions compatible with  $\lambda'$ . A KLTL formula  $\psi$  over  $\mathcal{P}'$  is realizable in  $\mathcal{M}_E$  if there exists an  $e$ -strategy  $\lambda'$  such that for all runs  $\rho \in \text{exec}(\mathcal{M}_E, \lambda')$ , we have  $\text{exec}(\mathcal{M}_E, \lambda'), \rho, 0 \models \psi$ .

**Proposition 2.** *There exists an  $e$ -strategy  $\lambda' : (\Sigma'_1 \mathcal{O})^* \rightarrow \Sigma'_1$  realizing  $\varphi^0$  in  $\mathcal{M}_E$  iff there exists a strategy  $\lambda : (\Sigma_1 \mathcal{O})^* \rightarrow \Sigma_1$  realizing  $\varphi^0$  in  $\mathcal{M}_E$ .*

*Proof.* Let see  $e$ -strategies and strategies as  $\Sigma'_1$ -labelled (resp.  $\Sigma_1$ -labelled)  $\mathcal{O}$ -trees. Given a tree representing  $\lambda'$ , we project its labels on  $\Sigma_1$  to get a tree representing  $\lambda$ . The strategy  $\lambda$  defined in this way realises  $\varphi^0$ , as  $\varphi^0$  does not contain any occurrence of propositions in  $\mathbb{K}$ . Conversely, given a tree representing  $\lambda$ , we extend its labels with  $\emptyset$  to get a tree representing  $\lambda'$ . It can be shown for the same reasons that  $\lambda'$  realizes  $\varphi^0$ .  $\square$

**Incremental tree automata construction** The invariant mentioned before can now be stated more precisely: for all  $i$ ,  $\mathcal{T}^i$  accepts the  $e$ -strategies  $\lambda' : (\Sigma'_1 \mathcal{O}) \rightarrow \Sigma'_1$  that realise  $\varphi^i$  in  $\mathcal{M}_E$ . Therefore, the UCT  $\mathcal{T}^i$  are labelled with  $e$ -actions  $\Sigma'_1$ . We now explain how they are constructed.

Since  $\varphi^d$  is an LTL formula, we follow the construction of Section 4 to build the UCT  $\mathcal{T}^d$ . Then, we construct  $\mathcal{T}^i$  from  $\mathcal{T}^{i+1}$ , for  $0 \leq i < d$ . The invariant tells us that  $\mathcal{T}^{i+1}$  defines all the  $e$ -strategies that realize  $\varphi^{i+1}$  in  $\mathcal{M}_E$ . It is only an over-approximation of the set of  $e$ -strategies that realize  $\varphi^i$  in  $\mathcal{M}_E$  (and *a fortiori*  $\varphi^0$ ), since the subformulas of  $\varphi^i$  of the form  $K\gamma$  correspond to atomic propositions  $k_\gamma$  in  $\varphi^{i+1}$ , and therefore  $\mathcal{T}^{i+1}$  does not check that they are satisfied. Therefore to maintain the invariant,  $\mathcal{T}^i$  is obtained from  $\mathcal{T}^{i+1}$  such that whenever an action that contains some formula  $k_\gamma \in \text{sub}(\varphi^{i+1})$  occurs on a transition of  $\mathcal{T}^{i+1}$ , we trigger (universally) a new transition to a UCT  $\mathcal{T}_{\gamma,I}$ , for the current information set  $I$  in  $\mathcal{T}^{i+1}$ , that will check that  $K\gamma$  indeed holds. The assumption on positivity of KLTL formulas is necessary here as we do not have to check for formulas of the form  $\neg K\gamma$ , which could not be done without an involved “non Safrasless” complementation step. Since  $\gamma$  is necessarily an LTL formula over  $\mathcal{P}'$  by definition of the formula  $\varphi^{i+1}$ , we can apply the construction of Section 4 to build  $\mathcal{T}_{\gamma,I}$ .

Formally, from the incremental way of constructing the automata  $\mathcal{T}^j$  for  $j \geq i$ , we know that  $\mathcal{T}^{i+1}$  has a set of states  $Q_{i+1}$  where all states are of the form  $(q, I)$  where  $I \subseteq S_e$  is some knowledge. In particular, it can be verified to be true for the state space of  $\mathcal{T}^d$  by definition of the construction of Section 4. Let also  $\Delta_{i+1}$  be the transition relation of  $\mathcal{T}^{i+1}$ . For all formulas  $\gamma$  such that  $k_\gamma$  occurs in  $\varphi^{i+1}$ , we let  $Q_\gamma$  be the set of states of  $\mathcal{T}_{\gamma,I}$  and  $\Delta_\gamma$  its set of transitions. Again from the construction of Section 4, we know that  $Q_\gamma = Q \times 2^{S_e}$  where  $Q$  is the set of states of a UCW associated with  $\gamma$  (assumed to be disjoint from that of  $\mathcal{T}^{i+1}$ ) and  $Q_\gamma^0 = Q \times I$ .

We define the set of states  $Q_i$  of  $\mathcal{T}^i$  by  $Q_{i+1} \cup \bigcup_{k_\gamma \in \text{sub}(\varphi^{i+1})} Q_\gamma$ . Its set of transitions  $\Delta_i$  is defined as follows. Assume w.l.o.g. that there is a unique initial state  $q_0 \in Q$  in the UCW  $\mathcal{A}_\gamma$ . If  $(q', I') \in \Delta_{i+1}((q, I), (a, K), o)$  where  $I, I' \subseteq S_e$ ,  $a \in \Sigma_1$ ,  $K \subseteq \mathbb{K}$ ,  $o \in \mathcal{O}$  and  $k_\gamma \in K$  is such that  $k_\gamma$  occurs in  $\varphi^{i+1}$ , then we let  $(q', I') \in \Delta_i((q, I), (a, K), o)$  and  $\Delta_\gamma((q_0, I), (a, K), o) \subseteq \Delta_i((q, I), (a, K), o)$ .

In the following, we prove that the automaton  $\mathcal{T}^i$  accepts exactly the strategies that realize  $\varphi^i$  in the extended runs of the environment  $\mathcal{M}_E$ . A  $e$ -strategy of the system in this case can be seen as a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree  $T$  and then we extend the notation for the set of executions induced by the strategy in the environment model by  $\text{exec}(\mathcal{M}_E, T)$ . A branch of  $T$  is a sequence  $\pi = \tilde{a}_1 o_1 \tilde{a}_2 o_2 \dots$  where  $\forall i, \tilde{a}_i \in \Sigma_1 \times 2^{\mathbb{K}}$  and  $o_i \in \mathcal{O}$ .

As we mentioned before, we start our construction with the LTL formula  $\varphi^d$  and construct the universal co-Büchi tree automaton  $\mathcal{T}^d = \mathcal{T}_{\varphi^d, S_0}$  for the LTL formula  $\varphi^d$ . Then, we construct the sequence of automata  $\mathcal{T}^d, \mathcal{T}^{d-1}, \dots, \mathcal{T}^0$  such that  $\mathcal{T}^i$  checks for the satisfaction of  $\varphi^i$  on the branches of accepted trees. In other words,  $\forall T \in \mathcal{L}(\mathcal{T}^i)$ ,  $\forall \rho \in \text{exec}(\mathcal{M}_E, T)$ , we have that

$$\text{exec}(\mathcal{M}_E, T), \rho, 0 \models \varphi^i$$

In order to prove that the invariant holds, let define

$$\text{exec}(\mathcal{M}_E, T, \pi, j) = \bigcup_{\pi' \in T: \pi'[0\dots j] = \pi[0\dots j]} \text{exec}(\mathcal{M}_E, \pi')$$

be the set of all extended executions that are compatible with the branch  $\pi$  of the tree  $T$  up to position  $j$ . This runs are indistinguishable from the runs in  $\text{exec}(\mathcal{M}_E, \pi)$  up to position  $j$ . Observe that by definition we have  $\text{exec}(\mathcal{M}_E, T, \pi, 0) = \text{exec}(\mathcal{M}_E, T)$ .

**Definition 2.** A  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree  $T$  is called fair with respect to a  $K$ -positive one-agent KLTL formula  $\phi$  if:

$$\begin{aligned} \forall \pi &= (a_0, K_0)o_0(a_1, K_1)o_1\dots, \forall j, \forall k_\gamma \in K_j \cap \text{sub}(\phi), \\ \forall \pi' \text{ s.t. } \pi'[0\dots j] &= \pi[0\dots j], \\ \forall \rho \in \text{exec}(\mathcal{M}_E, \pi'), & \text{ we have } \text{exec}(\mathcal{M}_E, \pi'), \rho, j \models \gamma \end{aligned}$$

Intuitively, whenever we have  $k_\gamma$  in a node of a branch  $\pi$  of the tree, the formula  $\gamma$  holds at that position on all the other branches that pass by that node. And, because  $T$  is a tree, the branches  $\pi'$  that pass by this node, have the same prefix up to that position as  $\pi$  and then the extended runs compatible with  $\pi'$  have the same observations up to there.

**Lemma 2.** If a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree  $T$  is fair with respect to  $\phi$ , then

$$\begin{aligned} \forall \psi \in \text{sub}(\phi), \forall \pi \text{ branch of } T, \forall j, \forall \rho \in \text{exec}(\mathcal{M}_E, T, \pi, j), \\ \text{if } \text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi, \\ \text{then } \text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi^{-1} \end{aligned}$$

where  $\psi^{-1}$  is obtained from  $\psi$  by replacing the atomic propositions  $k_\gamma$  with the formulas  $K\gamma$ .

*Proof.* Let fix a branch  $\pi = (a_0, K_0)o_0(a_1, K_1)o_1(a_2, K_2)o_2\dots$  of  $T$ . Then, the proof is by induction on the structure of  $\psi$ .

- if  $\psi = p \in \mathcal{P}$ ,  $\psi^{-1} = \psi$  and then,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models p$ .
- if  $\psi = k_\gamma$ , since  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models k_\gamma$ ,  $k_\gamma \in \text{proj}_2(\rho[j]) = K_j$ . Then, because  $k_\gamma \in \text{sub}(\psi)$ , we have  $k_\gamma \in K_j \cap \text{sub}(\psi)$ .  
Because  $T$  is a fair tree with respect to  $\phi$  and  $\psi \in \text{sub}(\phi)$ ,  $\forall \pi' \in T$  s.t.  $\pi'[0\dots j] = \pi'[0\dots j]$ , for all runs  $\rho' \in \text{exec}(\mathcal{M}_E, \pi')$ , we have that  $\text{exec}(\mathcal{M}_E, T), \rho', j \models \gamma$ . Then, since  $\text{exec}(\mathcal{M}_E, T, \pi, j) = \bigcup_{\pi' \text{ s.t. } \pi'[0\dots j] = \pi'[0\dots j]} \text{exec}(\mathcal{M}_E, \pi')$ , we have  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models K\gamma$ .
- if  $\psi = \psi_1 \wedge \psi_2$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1 \wedge \psi_2$ . That is,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_2$ . From the induction, we have  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1^{-1}$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_2^{-1}$  which means that  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1^{-1} \wedge \psi_2^{-1}$ . The proof is the same for  $\psi = \psi_1 \vee \psi_2$ .

- if  $\psi = \psi_1 \mathcal{U} \psi_2$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1 \mathcal{U} \psi_2$ . This means that  $\exists j' \geq j$  s.t.  $\forall j \leq k < j'$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, k \models \psi_1$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j'), \rho, j' \models \psi_2$ . From the fact that  $\text{exec}(\mathcal{M}_E, T, \pi, k) \subseteq RT, \pi, j$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j') \subseteq \text{exec}(\mathcal{M}_E, T, \pi, j)$ , we have that  $\forall j \leq k < j'$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, k), \rho, k \models \psi_1$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j'), \rho, j' \models \psi_2$ . Then, from the inductive hypothesis,  $\forall j \leq k < j'$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, k), \rho, k \models \psi_1^{-1}$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j'), \rho, j' \models \psi_2^{-1}$  which means that  $\forall j \leq k < j'$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, k \models \psi_1^{-1}$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j' \models \psi_2^{-1}$ . This is because  $\text{exec}(\mathcal{M}_E, T, \pi, j)$  is the set of runs that are consistent with  $\pi$  from position 0 to position  $j$  and  $\text{exec}(\mathcal{M}_E, T, \pi, j')$  and  $\text{exec}(\mathcal{M}_E, T, \pi, k)$  are subsets of  $\text{exec}(\mathcal{M}_E, T, \pi, j)$  (since  $j \leq k < j'$ ). Furthermore, all the runs in  $\text{exec}(\mathcal{M}_E, T, \pi, j')$  (and  $\text{exec}(\mathcal{M}_E, T, \pi, k)$  respectively) are distinguishable from the rest of words in  $\text{exec}(\mathcal{M}_E, T, \pi, j)$ . Then,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \psi_1^{-1} \mathcal{U} \psi_2^{-1}$ .
- if  $\psi = \bigcirc \psi_1$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \bigcirc \psi_1$   $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j+1 \models \psi_1$ . Because  $\text{exec}(\mathcal{M}_E, T, \pi, j+1) \subseteq \text{exec}(\mathcal{M}_E, T, \pi, j)$  and from the induction hypothesis, we have  $\text{exec}(\mathcal{M}_E, T, \pi, j+1), \rho, j+1 \models \psi_1^{-1}$ . Using the same argument as before,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j+1 \models \psi_1^{-1}$  and then  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \bigcirc \psi_1^{-1}$ .
- if  $\psi = \square \psi_1$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \square \psi_1$ . Then,  $\forall k \geq j$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, k \models \psi_1$ . Because  $\forall k \geq j$   $\text{exec}(\mathcal{M}_E, T, \pi, k) \subseteq \text{exec}(\mathcal{M}_E, T, \pi, j)$  and from the inductive hypothesis,  $\forall k \geq j$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, k), \rho, k \models \psi_1^{-1}$ . Using the same argument as before, we get  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, k \models \psi_1^{-1}$ ,  $\forall k \geq j$ . That is,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \square \psi_1^{-1}$ .
- if  $\psi = K\psi_1$ , the fact that  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models K\psi_1$  means that  $\forall \rho' \in \text{exec}(\mathcal{M}_E, T, \pi, j)$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho', j \models \psi_1$ . From the induction hypothesis applied for  $\rho'$ , we have that  $\forall \rho' \in \text{exec}(\mathcal{M}_E, T, \pi, j)$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho', j \models \psi_1^{-1}$ . That is,  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models K\psi_1^{-1}$ .
- if  $\psi = \neg \psi_1$ , then  $K\gamma$  is not a subformula of  $\psi_1$  because  $K\gamma$  does not occur under negations in  $\psi$ . This means that  $\psi_1 = \psi_1^{-1}$  and then  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \neg \psi_1$  implies that  $\text{exec}(\mathcal{M}_E, T, \pi, j), \rho, j \models \neg \psi_1^{-1}$ .

Then, denoting by  $\text{exec}(\mathcal{M}_E[X], \pi)$  the set of  $e$ -executions in  $\mathcal{M}_E$  that are compatible with  $\pi$  and start when the current possible states of  $\mathcal{M}_E$  are  $X$ , we can prove that:

**Lemma 3.** For all  $i \in \{0, \dots, d-1\}$ ,  $\forall T \in \mathcal{L}(\mathcal{T}^i)$ ,  $T$  is fair with respect to  $\varphi^{i+1}$ .

*Proof.* The proof of this theorem comes directly from the construction of  $\mathcal{T}^i$  from  $\mathcal{T}^{i+1}$ .

Let  $T \in \mathcal{L}(\mathcal{T}^i)$ . By the construction,  $\mathcal{T}^i$  is  $\mathcal{T}^{i+1}$  to which is "plugged" the automata  $\mathcal{T}_{\gamma, X}$  for the LTL formulas  $\gamma$  where  $k_\gamma \in \text{sub}(\varphi^{i+1})$  appears on a transition of  $\mathcal{T}^{i+1}$ .

Therefore, if at position  $j$  on a branch  $\pi = (a_0, K_0)o_0(a_1, K_1)o_1\dots$  of  $T$  we have  $k_\gamma \in K_j \cap \text{sub}(\varphi^{i+1})$ , there also starts the execution of  $\mathcal{T}_{\gamma, X}$  that accepts the "maximal" subtree of  $T$  with the root at position  $j$  on branch  $\pi$ . Let denote this subtree  $T_{\pi, j}$  and observe that it is formed from the suffixes  $\pi'[j\dots]$  of all the branches  $\pi'$  of  $T$  such that  $\pi'[0\dots j] = \pi[0\dots j]$ . Therefore, since  $T_{\pi, j} \in \mathcal{L}(\mathcal{T}_{\gamma, X})$  for some  $X$  and  $\gamma$  is an LTL formula, using the results in Section 4 we have that  $\text{exec}(\mathcal{M}_E[X], \pi'[j\dots]) \models \gamma$ .

That is,  $\forall \pi'$  branch of  $T$  such that  $\pi'[0\dots j] = \pi[0\dots j]$ ,  $\forall \rho \in \text{exec}(\mathcal{M}_E, \pi')$ , we have  $\text{exec}(\mathcal{M}_E, \pi'), \rho, j \models \gamma$ . This means that  $T$  is fair with respect to  $\varphi^{i+1}$ .

**Proposition 3.** For all  $i \in \{0, \dots, d\}$ ,  $\forall T \in \mathcal{L}(\mathcal{T}^i)$ ,  $\forall \pi$  branch of  $T$ ,  $\forall \rho \in \text{exec}(\mathcal{M}_E, \pi)$ ,

$$\text{exec}(\mathcal{M}_E, \pi), \rho, 0 \models \varphi^i$$

*Proof.* We do the proof of the theorem by induction on  $i$ . For the base case, if  $i = d$ ,  $\mathcal{T}^d = \mathcal{T}_{\varphi^d, S_0}$  and a sample adaptation of the proof of Lemma 1 yields that  $\forall \pi$  a branch of  $T \in \mathcal{L}(\mathcal{T}^d)$ , we have directly that  $\text{exec}(\mathcal{M}_E, \pi) \models \varphi^d$ .

For the inductive step, we suppose that the theorem holds for  $\mathcal{T}^{i+1}$  and prove it for  $\mathcal{T}^i$ . From the inductive hypothesis, we have that  $\forall T \in \mathcal{L}(\mathcal{T}^{i+1})$ ,  $\forall \pi$  branch of  $T$ ,  $\forall \rho \in \text{exec}(\mathcal{M}_E, \pi)$ ,  $\text{exec}(\mathcal{M}_E, \pi), \rho, 0 \models \varphi^{i+1}$ .

But, since  $\mathcal{L}(\mathcal{T}^{i+1}) \supseteq \mathcal{L}(\mathcal{T}^i)$ , it is true that  $\forall T \in \mathcal{L}(\mathcal{T}^i)$ ,  $\forall \pi$  branch of  $T$ ,  $\forall \rho \in \text{exec}(\mathcal{M}_E, \pi)$ , we have  $\text{exec}(\mathcal{M}_E, \pi), \rho, 0 \models \varphi^{i+1}$ .

Let fix  $T \in \mathcal{L}(\mathcal{T}^i)$  and a branch  $\pi$  of  $T$ . By Theorem 3,  $T$  is fair with respect to  $\varphi^{i+1}$ . Then, since  $\text{exec}(\mathcal{M}_E, \pi) \subseteq \text{exec}(\mathcal{M}_E, T, \pi, 0)$  and since all the runs  $\rho \in \text{exec}(\mathcal{M}_E, \pi)$  are distinguishable from the other runs from  $\text{exec}(\mathcal{M}_E, T, \pi, 0) \setminus \text{exec}(\mathcal{M}_E, \pi)$  (because they are compatible with different sequences  $\pi'$  of actions  $\tilde{a}_i$  and observations  $o_i$ ), we have that  $\forall \rho \in \text{exec}(\mathcal{M}_E, \pi)$ ,  $\text{exec}(\mathcal{M}_E, T, \pi, 0), \rho, 0 \models \varphi^{i+1}$ .

Then, by applying Lemma 2, results that  $\text{exec}(\mathcal{M}_E, T, \pi, 0), \rho, 0 \models \varphi^i$  and because of the inclusion  $\text{exec}(\mathcal{M}_E, \pi) \subseteq \text{exec}(\mathcal{M}_E, T, \pi, 0)$  and because the runs  $\rho \in \text{exec}(\mathcal{M}_E, \pi)$  are distinguishable from the other runs from  $\text{exec}(\mathcal{M}_E, T, \pi, 0) \setminus \text{exec}(\mathcal{M}_E, \pi)$ , we have that  $\text{exec}(\mathcal{M}_E, \pi), \rho, 0 \models \varphi^i$ .

**Lemma 4.** For all  $i \geq 0$ ,  $\mathcal{L}(\mathcal{T}^i)$  accepts the set of  $e$ -strategies that realize  $\varphi^i$  in  $\mathcal{M}_E$ .

From Lemma 4, we know that  $\mathcal{L}(\mathcal{T}^0)$  accepts the set of  $e$ -strategies that realize  $\varphi^0 = \varphi$  in  $\mathcal{M}_E$ . Then by Proposition 2 we get:

**Corollary 1.** The  $\text{KLTL}^+$  formula  $\varphi$  is realizable in  $\mathcal{M}_E$  iff  $\mathcal{L}(\mathcal{T}^0) \neq \emptyset$ .

*Proof.* From left to right, if  $\mathcal{L}(\mathcal{T}^0) \neq \emptyset$ , there exists a tree  $T \in \mathcal{L}(\mathcal{T}^0)$  that, by Theorem 3, satisfies the property that  $\forall \pi = (a_0, K_0)o_0(a_1, K_1)o_1\dots$  branch of  $T$ ,  $\forall r \in \text{exec}(\mathcal{M}_E, \pi)$ ,  $\text{exec}(\mathcal{M}_E, \pi), r, 0 \models \varphi^0$ . This means that all paths in the model  $\mathcal{M}_E$  compatible with the branches of the tree  $T$  satisfy the formula  $\varphi^0 = \varphi$ . That is, since  $\varphi$  is a  $K$ -positive  $\text{KLTL}$  formula over  $2^{\mathcal{P}}$ ,  $\forall r \in \text{exec}(\mathcal{M}_E, T)$ , we have that  $\text{proj}_1(\text{exec}(\mathcal{M}_E, T)), \text{proj}_1(r), 0 \models \varphi$ .

Then, there exists a strategy  $\lambda_1 : (\Sigma_1 \times \mathcal{O})^* \rightarrow \Sigma_1$  of the system, represented by the tree  $T$ , that, since  $\text{exec}(\mathcal{M}_E, \lambda_1) = \bigcup_{\pi \in T} \text{proj}_1(\text{exec}(\mathcal{M}_E, \pi))$ , satisfies the property that  $\forall \rho \in \text{exec}(\mathcal{M}_E, \lambda_1)$ , we have  $\text{exec}(\mathcal{M}_E, \lambda_1), \rho, 0 \models \varphi$ . That is,  $\varphi$  is realizable in  $\mathcal{M}_E$ .

From right to left, if  $\varphi$  is realizable in  $\mathcal{M}_E$ , there exists a strategy  $\lambda_1 : (\Sigma_1 \times \mathcal{O})^* \rightarrow \Sigma_1$  of the system such that  $\forall \rho \in \text{exec}(\mathcal{M}_E, \lambda_1)$ , we have  $\text{exec}(\mathcal{M}_E, \lambda_1), \rho, 0 \models \varphi$ . This strategy can be seen as a  $\Sigma_1$ -labelled  $\mathcal{O}$ -tree  $T$ . Note that any branch  $a_0 \xrightarrow{o_0} a_1 \xrightarrow{o_1} a_2\dots$  of  $T$  can be seen as a sequence  $\hat{\pi} = a_0o_0a_1o_1a_2o_2\dots$  where  $a_i \in \Sigma_1 \times \emptyset$  and  $o_i \in \mathcal{O}$ . Therefore,  $\forall r \in \text{exec}(\mathcal{M}_E, \hat{\pi})$ , we have  $\text{exec}(\mathcal{M}_E, \hat{\pi}), r, 0 \models \varphi$ .

Now, we annotate all the nodes  $a_j$  of the tree  $T$  with fresh atomic propositions  $k_\gamma \in \mathbb{K}$  whenever on the branch  $\hat{\pi}$ , the formula  $K\gamma$  has to be true at position  $j$ , i.e., whenever

for all the branches  $\hat{\pi}'$  of  $T$  such that  $\hat{\pi}[0\dots j] = \hat{\pi}'[0\dots j]$ ,  $\forall r \in \text{exec}(\mathcal{M}_E, \hat{\pi}')$ , it is true that  $\text{exec}(\mathcal{M}_E, \hat{\pi}'), r, j \models \gamma$ . Then, all branches of the obtained tree  $\tilde{T}$  will be of the form  $\pi = (a_0, K_0)o_0(a_1, K_1)o_1(a_2, K_2)o_2\dots$  where  $K_j = \{k_\gamma \mid K\gamma \text{ is true at position } j\}$ . Also,  $\forall r \in \text{exec}(\mathcal{M}_E, \pi)$ ,  $\text{exec}(\mathcal{M}_E, \pi), r, 0 \models \varphi^d$  thanks to the annotations. Therefore, the annotated tree  $\tilde{T}$  is accepted by  $\mathcal{T}^d$  and then, by the construction of the chain of over-approximations and by the fact that  $k_\gamma$  appears only where  $K\gamma$  is true, we have that  $\tilde{T}$  is also accepted by  $\mathcal{T}^0$ . This means that  $\mathcal{L}(\mathcal{T}^0) \neq \emptyset$ .

We now let  $\mathcal{T}_\varphi$  be the UCT obtained by projecting  $\mathcal{T}^0$  on  $\Sigma_1$ . We have:

**Theorem 2.** *For any KLTL<sup>+</sup> formula  $\varphi$ , one can construct a UCT  $\mathcal{T}_\varphi$  such that  $\mathcal{L}(\mathcal{T}_\varphi)$  is the set of strategies that realize  $\varphi$  in  $\mathcal{M}_E$ .*

The number of states of  $\mathcal{T}_\varphi$  is (in the worst-case)  $2^{|S_e|} \cdot (2^{|\varphi^d|} + \sum_{k_\gamma \in \mathbb{K}} 2^{|\gamma|})$ , and since  $|\varphi^d| + \sum_{\gamma \in \mathbb{K}} |\gamma|$  is bounded by  $|\varphi|$ , the number of states of  $\mathcal{T}_\varphi$  is  $O(2^{|S_e| + |\varphi|})$ .

## 6 Antichain Algorithm

In the previous sections, we have shown how to reduce the problem of checking the realizability of some KLTL<sup>+</sup> formula  $\varphi$  to the emptiness of a UCT  $\mathcal{T}_\varphi$  (Theorem 2). In this section, we describe an antichain symbolic algorithm to test the emptiness of  $\mathcal{T}_\varphi$ .

It is already known from [8] that checking emptiness of the language defined by a UCW  $\mathcal{A}$  can be reduced to checking the emptiness of  $\mathcal{L}_{uc,B}(\mathcal{A})$  for a sufficiently large bound  $B$ , which in turn can be reduced to solving a safety game. The same thing result holds in the case of UCT automata. This results from [15,14] that proved that if a UCT automaton with  $n$  states is not empty, then it accepts a finite state machine of width bounded by  $2n(n!)^2$ . We prove here the result.

**Theorem 3 ([15], [14]).** *For every NBW  $\mathcal{N}$  with  $n$  states, there is a DPW  $\mathcal{D}$  with  $2n(n!)^2$  states and index  $2n$  such that  $\mathcal{L}(\mathcal{D}) = \mathcal{L}(\mathcal{N})$ .*

**Theorem 4.** *A UCT  $\mathcal{T}$  with  $n$  states is not empty iff there is a non-empty Moore machine with at most  $2n^2(n!)^2$  states that generates a tree accepted by  $\mathcal{T}$ .*

*Proof.* The proof of the theorem follows the steps of one in [13], but passes through parity automata. The interesting part of the proof is from left to right. Let consider that the language of  $\mathcal{T}$  is not empty.

We construct the NBW  $\mathcal{A}'$  that runs over paths of a run tree of  $\mathcal{T}$  and checks that they are rejecting.  $\mathcal{A}'$  has the same state space as  $\mathcal{T}$ . By 3, we have an DPW  $\mathcal{A}''$  with  $2n(n!)^2$  states equivalent to  $\mathcal{A}'$ . Then, we complement  $\mathcal{A}''$  by incrementing the priorities and get the DPW  $\tilde{\mathcal{A}}$  over all paths of the input tree to check that all paths of the run tree of  $\mathcal{T}$  are accepting. And finally, we get the NPT  $A^n$  that is equivalent to  $\mathcal{T}$ .

We then view  $A^n$  as a two-player parity game  $G(A^n)$  with perfect information and  $2n(n!)^2$  states such that  $\mathcal{L}(A^n) \neq \emptyset$  iff there is a winning strategy for the protagonist in  $G(A^n)$ .

Then, it is known that parity games admit memoryless strategies. Therefore, if  $\mathcal{L}(A^n) \neq \emptyset$ , there is a strategy for the protagonist in  $G(A^n)$  that can be obtained by keeping the outgoing transitions labelled with a specific action of the protagonist. And we can then transform it into a Moore machine with  $2n(n!)^2$  states.



**Lemma 5.** *Given an UCT  $\mathcal{T}$  that accepts  $\Sigma$ -labelled  $D$ -trees with  $n$  states and a Moore machine  $\mathcal{M}_\lambda$  with  $m$  states representing a strategy  $\lambda : (\Sigma D)^* \rightarrow \Sigma$ . Let  $T_\lambda$  be the tree representation of  $\lambda$ . Then,*

$$T_\lambda \in \mathcal{L}_{UC}(\mathcal{T}) \text{ iff } T_\lambda \in \mathcal{L}_{UC,mn}(\mathcal{T})$$

*Proof.* From right to left, the proof is obvious since  $\mathcal{L}_{UC,mn}(\mathcal{T}) \subseteq \mathcal{L}_{UC}(\mathcal{T})$ . In the other sense, let  $\mathcal{M}_\lambda = (\Sigma, D, Q_M, p_0, \delta_M, g_M)$  be the Moore machine representing the strategy  $\sigma$ .

First, we transform  $\mathcal{M}_\lambda$  into a nondeterministic Büchi word automata  $\mathcal{A}_M = (Q_M, p_0, \delta'_M, \Sigma' = \Sigma \times D, \alpha' = Q_M)$  for which  $\mathcal{L}(\mathcal{A}_M)$  is the set of traces of  $\mathcal{M}_\lambda$  (Accepts all the branches of  $T_\lambda$  since it is a "tree unfolding" of  $\mathcal{M}_\lambda$ ). Note that  $\mathcal{A}_M$  has  $m$  states.

Let  $\mathcal{T} = (Q, Q_0, \Delta, \alpha, \Sigma, D)$ . We define the product  $\mathcal{T} \times \mathcal{A}_M = (\bar{Q}, \bar{Q}_0, \bar{\Delta}, \bar{\alpha}, \Sigma, D)$  where  $\bar{Q} = Q \times Q_M, \bar{Q}_0 = Q_0 \times p_0, \bar{\alpha} = \alpha \times Q_M$  and

$$\bar{\Delta}((q, p), \sigma, d) = \{(q', p') \mid q' \in \Delta(q, \sigma, d) \text{ and } p' = \delta'_M(p, (\sigma, d))\}$$

Now, since  $T_\lambda \in \mathcal{L}_{UC}(\mathcal{T})$ , for all accepting run  $r_\lambda$  on  $T_\lambda$  in  $\mathcal{T}$ , for all  $\rho$  a path of  $r_\lambda$  with  $\rho = q_0 q_1 q_2 \dots, \text{inf}(\rho) \cap \alpha = \emptyset$ . That is,  $\exists 0 \leq i \leq k < j$  s.t.  $q_i = q_j$  and  $q_k \in \alpha$  (there are no cycles in  $r_\lambda$ ). Then, from the construction of  $\mathcal{T} \times \mathcal{A}_M$ , there is no cycle reachable from  $\bar{Q}_0$  that contains a final state. Therefore, all the paths  $\rho$  of runs  $r_\lambda$  on  $T_\lambda$  visit at most  $m \times n$  final states. That is,  $T_\lambda \in \mathcal{L}_{UC,mn}(\mathcal{T})$ .

Then, using this results we can turn the emptiness problem for the UCT automaton  $\mathcal{T}$  into the emptiness problem of the UCB automaton  $(\mathcal{T}, B)$  as follows:

**Theorem 5.** *Given the UCT  $\mathcal{T}$  over  $\Sigma_1 \times 2^{\mathbb{K}}$  with  $n$  states and  $B = 2n^2(n!)^2, \mathcal{L}_{uc}(\mathcal{T}) \neq \emptyset$  iff  $\mathcal{L}_{uc,B}(\mathcal{T}) \neq \emptyset$ .*

*Proof.* If  $\mathcal{L}_{uc}(\mathcal{T}) \neq \emptyset$ , then, by Theorem 4, there exists a regular tree  $T_\lambda \in \mathcal{L}_{uc}(\mathcal{T})$  generated by a finite state machine with at most  $m$  states ( $m < 2n(n!)^2$ ). Then, by the Lemma 5,  $T_\lambda \in \mathcal{L}_{uc,mn}(\mathcal{T})$  and  $\mathcal{L}_{uc,B}(\mathcal{T}) \neq \emptyset$ . In the other sense, the proof is obvious since  $\mathcal{L}_{uc,B}(\mathcal{T}) \subseteq \mathcal{L}_{uc}(\mathcal{T})$ .

Clearly, for all  $B \geq 0$ , if  $\mathcal{L}_{uc,B}(\mathcal{T}) \neq \emptyset$ , then  $\mathcal{L}_{uc}(\mathcal{T}) \neq \emptyset$ . This has led to an incremental algorithm by starting with some small bound  $b$  and the experiments have shown that in general, a small bound  $b$  is necessary to conclude for realizability of an LTL formula (transformed into the emptiness of a UCT). We also exploit this idea in our implementation and show that for KLTL<sup>+</sup> specifications that we considered, this observation still holds: small bounds are enough.

In [8], it is shown that the safety games can be solve on-the-fly without constructing them explicitly, and that the fixpoint algorithm used to solve these safety games could be optimized by using some antichain representation of the sets constructed during the fixpoint computation. Rather than using the algorithm of [8] as a black box, we study the state space of the safety games constructed from the UCT  $\mathcal{T}_\varphi$  and show that they are also equipped with a partial order that allows one to get more compact antichain representations.

## 6.1 Reduction to Safety Game

Given a bound  $b \geq 0$  and a UCT  $\mathcal{T}$ , the idea is to construct a safety game  $G(\mathcal{T}, b)$  such that Player 1 has a winning strategy in  $G(\mathcal{T}, b)$  iff  $L_{uc,b}(\mathcal{T})$  is non-empty. The game  $G(\mathcal{T}, b)$  is obtained by extending the classical automata subset construction with counters which count, up to  $b$ , the maximal number of times all the runs, up to the current point, have visited accepting states. We construct a complete deterministic 0-Co-Büchi tree automaton by extending the subset construction with counters. If  $Q$  is the set of states of  $\mathcal{T}$ , the set of states of the safety game  $G(\mathcal{T}, b)$  is all the functions  $F : Q \rightarrow \{-1, 0, \dots, b+1\}$ . The value  $F(q) = -1$  means that no run have reached  $q$  and  $F(q) \in \{0, \dots, b\}$  means that the maximal number of accepting states that has been visited by the runs reaching  $q$  is  $F(q)$ . The safe states are all the functions  $F$  such that  $F(q) \leq b$  for all  $q \in Q$ . The set of states can be partially ordered by the pairwise comparison between functions and it is shown that the sets of states manipulated by the fixpoint algorithm are downward closed for this order.

Consider the UCT  $\mathcal{T}_\varphi$  constructed from the KLTL<sup>+</sup> formula  $\varphi$ . The deterministic 0-Co-Büchi automaton for  $(\mathcal{T}, B)$  is the universal tree automaton  $Det(\mathcal{T}, B) = \{\mathcal{F}, F_0, \alpha', \delta\}$  where  $\mathcal{F}$  is the set of states,  $F_0$  is the initial state,  $\alpha'$  is the set of final states and  $\delta$  is the transition relation with:

$$\begin{aligned} & - \mathcal{F} = \{F | F : Q \rightarrow \{-1, 0, \dots, B+1\}\} \\ & - \forall (q, I) \in Q, F_0(q, I) = -1 \text{ if } (q, I) \notin Q_0 \text{ and } F_0(q, I) = ((q, I) \in \alpha) \text{ otherwise} \\ & - \alpha' = \{F \in \mathcal{F} | \exists (q, I) \in Q \text{ s.t. } F(q, I) > B\} \\ & - \forall o \in \mathcal{O}, \delta(F, (a, K), o) = (F', o) \text{ if} \\ & \quad F'(q', I') = \max\{\min\{B+1, F(q, I) + ((q', I') \in \alpha) | \\ & \quad \quad ((q', I'), o) \in \Delta((q, I), (a, K), o) \text{ and } F(q, I) \neq -1\}\} \end{aligned}$$

where  $\max(\emptyset) = -1$  and, for all states  $(q, I) \in Q$ ,  $((q, I) \in \alpha) = 1$  if  $(q, I)$  is in  $\alpha$  and 0 otherwise. We say that a state  $F$  is unsafe if there exists  $(q, I) \in Q$  such that  $F(q, I) = B+1$ .

**Theorem 6.** *Let  $\mathcal{T}$  be the UCT for for the KLTL<sup>+</sup> formula  $\varphi$  and  $B \in \mathbb{N}$ . Then,  $Det(\mathcal{T}, B)$  is complete, deterministic and  $\mathcal{L}_{uc,0}(Det(\mathcal{T}, B)) = \mathcal{L}_{uc,B}(\mathcal{T})$ .*

*Proof.* It is obvious that the constructed automaton is complete by construction and from the completeness of  $(\mathcal{T}, B)$ . Also, it is deterministic because of the construction. From a state  $F$ , for one action of the agent and one observation we get only one successor state  $F'$ .

Then, we prove the language equality by double inclusion. From right to left, if  $\langle T, \tau \rangle$  is a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree accepted by  $(\mathcal{T}, B)$ , there exists an accepting run on  $\langle T, \tau \rangle$  in  $(\mathcal{T}, B)$ . Let it be  $\langle T_r, \tau_r \rangle$  such that  $\tau_r(\epsilon) \in Q_0 \times \{\epsilon\}$  and  $\forall x \in T_r$  s.t.  $\tau_r(x) = ((q, I), v)$  and  $\tau(v) = (a, K)$ , if  $\Delta((q, I), (a, K)) = \{((q_1, I_1), o_1), \dots, ((q_n, I_n), o_n)\}$ , then  $\forall 0 < j \leq n$ ,  $x_j \in T_r$  and  $\tau_r(x_j) = ((q_j, I_j), v \cdot o_j)$ .

Since the run on  $T$  is accepting, each branch  $\pi$  of  $t_\lambda$  induces a infinite sequence  $\tau_r(\pi) = ((q_0, I_0), \epsilon)((q_1, I_1), v_1)((q_2, I_2), v_2) \dots$  which visits at most  $B$  times the final states. Then, by the construction of  $Det(\mathcal{T}, B)$  using subset construction on the states with the same observation, there exists a run on  $\langle T, \tau \rangle$  which is accepting since the set of final state  $\alpha' = \{F \in \mathcal{F} | \exists (q, I) \in Q' \text{ s.t. } F(q, I) > B\}$  it is not reached.

From left to right, let  $\langle T, \tau \rangle$  be a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree accepted by  $Det(\mathcal{T}, B)$ . Then, there exists a run on  $\langle T, \tau \rangle$  in  $Det(\mathcal{T}, B)$ . Let it be  $\langle T_r, \tau_r \rangle$  such that  $\tau_r(\epsilon) \in F_0 \times \{\epsilon\}$  and  $\forall x \in T_r$  s.t.  $\tau_r(x) = (F, v)$  and  $\tau(v) = (a, K)$ , if  $\delta(F, (a, K)) = \{(F_1, o_1), \dots, (F_n, o_n)\}$ , then  $\forall 0 < j \leq n$ ,  $xj \in T_r$  and  $\tau_r(xj) = (F_j, v \cdot o_j)$ .

Since the run is accepting in the 0-Co-Büchi tree automaton, it means that it doesn't visit any  $F$  such that  $\exists(q, I)$  for which  $F(q, I) \geq B + 1$ . This means that, using the construction of  $Det(\mathcal{F}, B)$ , the branches of the run in  $(\mathcal{T}, B)$  on  $\langle T, \tau \rangle$  visit at most  $B$  times the final states in  $\alpha$ . Therefore,  $T \in \mathcal{L}_{uc, B}(\mathcal{T})$ .

A two-player safety game is played on a game arena. To reflect the game point of view, the automaton  $Det(\mathcal{T}, B)$  can be seen as a *game arena*  $G(\mathcal{T}, B) = (\mathcal{F}_E, \mathcal{F}, q_0, T = T_1 \cup T_E, Safe)$  with  $|\mathcal{F}| + |\delta|$  states where  $\mathcal{F}_E \cup \mathcal{F}$  is the set of states of the game,  $\mathcal{F}$  being the set of states controlled by the system and  $\mathcal{F}_E$  the set of states controlled by the environment,  $q_0 = F_0 \in \mathcal{F}$  is the initial state,  $T \subseteq \mathcal{F} \times \mathcal{F}_E \cup \mathcal{F}_E \times \mathcal{F}$  is the transition relation and  $Safe \subseteq \mathcal{F}$  is the safety winning condition defined by

- $\mathcal{F}_E = \{F_{a,K} \mid F \in \mathcal{F} \text{ and } (a, K) \in \Sigma_1 \times 2^{\mathbb{K}}\}$
- $T_1 = \{(F, F_{a,K}) \mid F \in \mathcal{F} \text{ and } (a, K) \in \Sigma_1 \times 2^{\mathbb{K}}\}$
- $T_E = \{(F_{a,K}, F') \mid \forall F' \in \delta(F, (a, K))\}$
- $Safe = \mathcal{F} \setminus \alpha'$

Then,  $\mathcal{F}_E \cup \mathcal{F}$  is the total set of states of the game and  $T = T_1 \cup T_E$  is the transition relation.

A infinite *play* on  $G(\mathcal{T}, B)$  is a path  $\rho = \rho_0 \rho_1 \rho_2 \dots \in (\mathcal{F} \mathcal{F}_E)^\omega$  such that  $\forall i \geq 0$ ,  $(\rho_i, \rho_{i+1}) \in T$ . Finite plays are similarly defined and they belong to  $(\mathcal{F} \mathcal{F}_E)^*$ . A strategy for the system is a mapping  $\lambda_1 : (\mathcal{F} \mathcal{F}_E)^* \mathcal{F} \rightarrow \mathcal{F}_E$  that maps every finite play  $\rho$  whose last state is  $\rho_n \in \mathcal{F}$  to a state  $\rho_{n+1}$  such that  $(\rho_n, \rho_{n+1}) \in T_1$ . A strategy for the environment is a mapping  $\lambda_2 : (\mathcal{F} \mathcal{F}_E)^* \rightarrow \mathcal{F}$  that maps an finite play ending in  $\rho_n \in \mathcal{F}_E$  to a state  $\rho_{n+1} \in \mathcal{F}$  such that  $(\rho_n, \rho_{n+1}) \in T_E$ . The *outcome* of a strategy  $\lambda_1$  of the system is a set  $Outcome_{G(\mathcal{T}, B)}(\lambda_1)$  of infinite plays  $\rho = \rho_0 \rho_1 \rho_2 \dots \in (\mathcal{F} \mathcal{F}_E)^\omega$  such that if  $\rho_i \in \mathcal{F}$ , then  $\rho_{i+1} = \lambda_1(\rho_0 \dots \rho_i)$ . A strategy for the system is winning if  $Outcome_{G(\mathcal{T}, B)}(\lambda_1) \subseteq (Safe \mathcal{F}_E)^\omega$ .

**Theorem 7.**  $\mathcal{L}_{uc, 0}(Det(\mathcal{T}, B)) \neq \emptyset$  iff the system has a winning strategy  $\lambda_1$  in the game  $G(\mathcal{T}, B)$ .

*Proof.* If  $\mathcal{L}_{uc, 0}(Det(\mathcal{T}, B)) \neq \emptyset$ , then there exists a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree  $T$  which is accepted by  $Det(\mathcal{T}, B)$ . Then, there exists an accepting run  $\langle T_r, \tau_r \rangle$  on  $T$  whose paths don't visit the final states in  $\alpha'$ . Then, by the construction of the game  $G(\mathcal{T}, B)$ , there exists a strategy (represented by the tree  $T$ ) which is winning in  $G(\mathcal{T}, B)$  because all the paths in the game arena that follow it visit exactly the same states in  $\mathcal{F}$  as  $\langle T_r, \tau_r \rangle$  does and the safe set is defined as  $Safe = \mathcal{F} \setminus \alpha'$ .

Now, if there exists a winning strategy  $\lambda_1$  in  $G(\mathcal{T}, B)$  (which can be seen as a  $\Sigma_1 \times 2^{\mathbb{K}}$ -labelled  $\mathcal{O}$ -tree  $T$ ), then all the outcomes of  $\lambda_1$  will stay in the set  $Safe$ . Then, by construction of the game arena, there exists a run on  $T$  in the automaton  $Det(\mathcal{T}, B)$  which is also accepting because  $Safe = \mathcal{F} \setminus \alpha'$  which means that the run doesn't visit the set  $\alpha'$ .

Therefore, by Theorems 2, 5, 6 and 7, we have that the KLTTL<sup>+</sup> formula  $\varphi$  is realizable in the environment model  $\mathcal{M}_E$  if and only if the system has a winning strategy in the game  $G(\mathcal{T}, B)$ .

## 6.2 Complexity

Consider now the UCT  $\mathcal{T}_\varphi$  constructed from the KLTL<sup>+</sup> formula  $\varphi$ . Its state space is of the form  $Q \times 2^{S_e}$  where  $S_e$  is the set of states of the environment, because the construction also take into account the knowledge the system has from the environment. Given a bound  $b$ , the state space of the safety game  $G(\mathcal{T}_\varphi, b)$  is therefore functions from  $Q \times 2^{S_e}$  to  $\{-1, \dots, b+1\}$ . However, we can reduce this state space thanks to the following result:

**Proposition 4.** *For all runs  $\langle T_r, \tau \rangle$  of  $\mathcal{T}_\varphi$  on some tree  $T$ , for all branches  $\pi, \pi'$  in  $T_r$  of the same length such that they follow the same sequence of observations, if  $\tau(\pi) = (q, I)$  and  $\tau(\pi') = (q, I')$ , then  $I = I'$ .*

*Proof.* It is clear from the definition of  $\mathcal{F}$  that for any sequence of observations there is only one corresponding function  $F$ . Then, since  $\tau(\pi) = (q, I)$  and  $\tau(\pi') = (q, I')$ , we have that  $F(q, I) \neq \emptyset$  and  $F(q, I') \neq \emptyset$ .

We prove the lemma by contradiction. Suppose that  $F$  is reachable and  $I \neq I'$ . The, by induction on the length of the path leading to a state  $F$ , we have:

If  $F = F_0$ , then since  $F(q, I) \neq -1$ , we have  $(q, I) \in Q_0 = \{(q_0, S_0) \mid q_0 \in Q_\psi^0\}$ . Then,  $\pi$  and  $\pi'$  have length 1 and therefore  $I = I'$ .

Now, suppose that there is reached a state  $F_i \in \mathcal{F}$  such that for all  $(q, I), (q, I') \in Q$  with  $F_i(q, I) \neq -1$  and  $F_i(q, I') \neq -1$ , we have  $I = I'$ . Let this set be denoted by  $I_i$ . We then prove that, after one round in the game, the state  $F_{i+1}$  that is reached satisfy the property that for all  $(q, I), (q, I') \in Q$  such that  $F_{i+1}(q, I) \neq -1$  and  $F_{i+1}(q, I') \neq -1$ , we have  $I = I'$ .

From the construction of the game arena,  $F_{i+1}$  is computed from  $F_i$  using the action  $(a, K)$  proposed by the agent and observation  $o$  given by the environment. That is,  $F_{i+1}(q, I') = \max\{\min\{B+1, F_i(q'', I_i) + ((q, I') \in \alpha) \mid ((q, I'), o) \in \Delta((q'', I_i), (a, K))\} \text{ and } F_i(q'', I_i) \neq -1\}$ . From the definition, the states  $(q, I')$  for which  $F_{i+1}(q, I') \neq -1$  are states for which there exists  $(q'', I_i)$  with  $F_i(q'', I_i) \neq -1$  and  $I' = \text{Post}_a(I_{q'', q}^i, o)$ . Therefore, since the state  $F_{i+1}$  is uniquely determined by a sequence of actions and observations, and the states  $q$  and sets  $I$  are synchronized along a path, we have that all the states  $(q, I')$  with  $F_{i+1}(q, I') \neq -1$  contain the same set  $I_{i+1} = I'$  of states of the environment.

In other words, given the same sequence of observations, the tree automaton  $\mathcal{T}_\varphi$  computes, for a given state  $q$ , the same knowledge.

Based on this proposition, it is clear that reachable states  $F$  of  $G(\mathcal{T}_\varphi, b)$  satisfy, for all states  $q \in Q$  and knowledges  $I, I'$ , if  $F(q, I) \neq -1$  and  $F(q, I') \neq -1$  then  $I = I'$ . We can therefore focus on the reachable states and define the state space of  $G(\mathcal{T}_\varphi, b)$  as the set of pairs  $(F, \bar{K})$  such that  $F : Q \rightarrow \{-1, \dots, b+1\}$  and  $\bar{K} : Q \rightarrow 2^{S_e}$  associates with each state  $q$  a knowledge (we let  $G(q) = \emptyset$  if  $F(q) = -1$ ).

**Lemma 6.** *The number of states of the safety game  $G(\mathcal{T}, B)$  built for the KLTL<sup>+</sup> formula  $\varphi$  in the environment  $\mathcal{M}_E$  has at most  $\frac{(B+3)2^{|\Sigma|+|\varphi|+1}-1}{B+2} \times (1+2^{|\Sigma|}) \times (1+|\Sigma| \times 2^{\mathbb{K}})$  states.*

*Proof.*  $\mathcal{F} = \{F \mid F : Q \rightarrow \{-1, 0, \dots, B+1\}\}$ . But, according to the previous lemma,  $\forall (q, I), (q', I') \in Q$  such that  $F(q, I) \neq -1$  and  $F(q', I') \neq -1$ , we have  $I = I'$ . It means that, for a function  $F$ , there are at most  $n = |Q| + \sum_{\gamma: \exists i \text{ s.t. } K\gamma \in \text{Sub}(\varphi^i)} |Q_\gamma|$  states  $(q, I)$  for which  $F(q, I) \neq -1$ .

Then, because of the way the set  $I$  is computed using  $\text{Post}_a(Y, o)$ , excepting the initial state that may fall into several observations, there exists  $o \in \mathcal{O}$  such that  $I \subseteq o$ . This means that there exist  $1 + \sum_{o \in \mathcal{O}} 2^{|o|}$  sets of states of the environment that are reached, where 1 is for the initial state.

In conclusion, the number of functions  $F$  is  $[(B+2)^0 + (B+2)^1 + \dots + (B+2)^n] \times (1 + \sum_{o \in \mathcal{O}} 2^{|o|})$ . That is,  $\frac{(B+3)^{n+1} - 1}{B+2} \times (1 + \sum_{o \in \mathcal{O}} 2^{|o|})$  functions.

Then, for each state  $F$ , we have  $|\Sigma_1 \times 2^{\mathbb{K}}|$  states  $F_{a,K}$  controlled by the environment. This gives a total number of  $\frac{(B+3)^{n+1} - 1}{B+2} \times (1 + \sum_{o \in \mathcal{O}} 2^{|o|}) \times (1 + |\Sigma_1 \times 2^{\mathbb{K}}|)$  of states of the game  $G(\mathcal{T}, B)$ .

Now, since  $n$  can be bounded by  $2^{|S|+|\varphi|}$  and  $\sum_{o \in \mathcal{O}} 2^{|o|} \leq 2^{|S|}$ , we get that the number of states of  $G(\mathcal{T}, B)$  is bounded by  $\frac{(B+3)^{2^{|S|+|\varphi|+1}} - 1}{B+2} \times (1 + 2^{|S|}) \times (1 + |\Sigma_1 \times 2^{\mathbb{K}}|)$ .

This state space is naturally ordered by  $(F_1, \overline{K}_1) \preceq (F_2, \overline{K}_2)$  if for all  $q \in Q$ ,  $F_1(q) \leq F_2(q)$  and  $\overline{K}_1(q) \subseteq \overline{K}_2(q)$ . We show that all the sets manipulated during the fixpoint computation used to solve the safety games are downward closed for this order and therefore can be represented by the antichain of their maximal elements. A detailed analysis of the size of the safety game shows that  $G(\mathcal{T}_\phi, B)$  is doubly exponential in the size of  $\varphi$ , and therefore, since safety games can be solved in linear time, one gets a 2EXPTIME upper bound for  $\text{KLTL}^+$  realizability.

**Proposition 5.** *The realizability of a  $K$ -positive  $\text{KLTL}$  formula  $\varphi$  is decidable in 2EXPTIME.*

*Proof.* As also showed in Section 5, the number of states if the co-Büchi tree automaton  $\mathcal{T}$  is bounded by  $2^{|S|+|\varphi|}$  and is built in EXPTIME since there is a polynomial number of formulas  $K\gamma$  in  $\phi$ .

Then, when transforming  $\mathcal{T}$  into a  $B$ -co-Büchi tree automaton, the bound  $B$  equals to  $4n^2(n!)^2$ . By applying Lemma 6, we have that the safety game  $G(\mathcal{T}, B)$  is built in 2EXPTIME.

## 7 Implementation and Case Studies

In this section we briefly present our prototype implementation Acacia-K for  $\text{KLTL}^+$  synthesis [1], and provide some interesting examples on which we tested the tool, on a laptop equipped with an Intel Core i7 2.10Ghz CPU. Acacia-K extends the LTL synthesis tool Acacia+[5]. As Acacia+, the implementation is made in Python together with C for the low level operations that need efficiency.

As Acacia+, the tool is available in one version working on both Linux and MacOSX and can be executed using the command-line interface. As parameters, in addition to the files containing the  $\text{KLTL}^+$  formula and the partition of the signals and actions, Acacia-K requires a file with the environment model. The output of the tool is a winning

strategy, if the formula is realizable, given as a Moore machine described in Verilog and if this strategy is small, Acacia-K also outputs it as a picture.

In order to have a more efficient implementation, the construction of the automata for the LTL formulas  $\gamma$  is made on demand. That is, we construct the UCT  $\mathcal{T}_\gamma$  incrementally by updating it as soon as it needs to be triggered from some state  $(q, I)$  which has not been constructed yet.

As said before, the synthesis problem is reduced to the problem of solving a safety game for some bound  $b$  on the number of visits to accepting states. The tool is incremental: it tests realizability for small values of  $b$  first and increments it as long as it cannot conclude for realizability. In practice, we have observed, as for classical LTL synthesis, that small bounds  $b$  are sufficient to conclude for realizability. However if the formula is not realizable, we have to iterate up to a large upper bound, which in practice is too large to give an efficient procedure for testing unrealizability. We leave as future work the implementation of an efficient procedure for testing unrealizability.

Taking now Example 1, the strategy provided by the tool is depicted in Figure 2. It asks to play first "toggle" and then keep play "skip" and, depending on the observation he gets, the system goes in a different state. The state 0 is for the start, the state 1 is the "error" state in which the system goes if he receives a wrong observation. That is, the environment gives an observation even if he cannot go in a state having that observation. Then, if the observation is correct, after playing the action "toggle" from the initial states  $\{s_1, s_2\}$ , the environment is forced to go in  $s_3$  and by playing the action "skip", the system forces the environment to stay in  $s_3$  and he will know that  $t$  is false. In the strategy, this situation corresponds to the state 2. For this example, Acacia-K constructed a UCT with 31 states and the total running time is 0.2s.




Fig. 2: Winning strategy synthesized by Acacia-K for Example 1

*Example 2 (The 3-Coin Game).* Another example that we tried is a game played using three coins which are arranged on a table with either *head* or *tail* up. The system doesn't see the coins, but knows at each time the number of tails and heads. Then, the game is infinitely played as follows. At the beginning the environment chooses an initial configuration and then at each round, the system chooses a coin and the environment has to flip that coin and inform the system about the new number of heads and tails. The objective of the system is to reach, at least once, the state in which all the coins have the *heads* up and to avoid all the time the state in which all the coins are *tails*. Depending on the initial number of tails up, the system may or may not have a winning strategy.

In order to model this, we considered an environment model whose states are labelled with atomic propositions  $c_1, c_2, c_3$  for the three coins, which are not visible for

the system, and two other variables  $b_1, b_0$  which are visible and represent the bits encoding the number of *heads* in the configuration. The actions of the system are  $C_1, C_2, C_3$  with which he chooses a coin and the environment has to flip the coin chosen by the system by playing only the action *done*. A picture of the environment is in Figure 3.

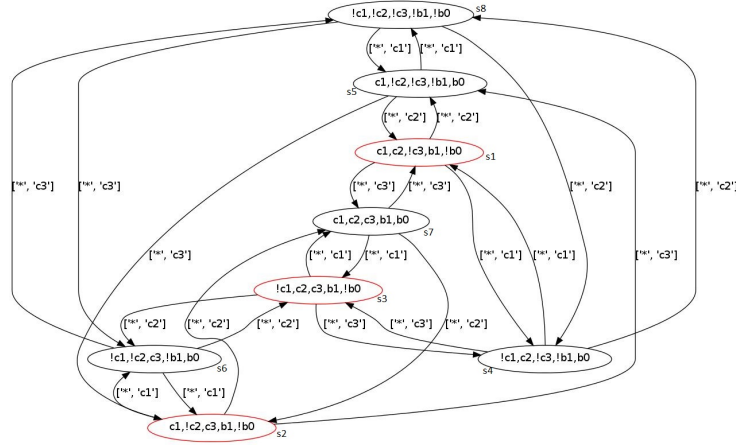


Fig. 3: The environment for 3-Coins Game

Then, the specification is translated into the KLTL<sup>+</sup> formula  $K\Diamond(c_1 \wedge c_2 \wedge c_3) \wedge \Box K(c_1 \vee c_2 \vee c_3)$ . Then, assuming that the initial state of the environment has two *heads*, the synthesized strategy proposes to ”check” the position of every coin by double flipping. If after one flip, the winning state is not reached, the system flips back the coin and at the third round he chooses another coin to check. A picture of the strategy can be found in Figure 4.

For this example, Acacia-K constructs a UCT with 79 states, synthesises a strategy with 10 states, and the total running time is 3.9s. In the Moore machine representing the strategy of Fig.4, the state 2 corresponds to the configuration in which there are only *heads* and the states 3 and 7 are the states in which he goes when the environment cheated. Then, the strategy has two parts. One that leads to the state 2 by checking the coins and the second part in which the system plays only  $c_3$  and moves between the states  $s_1$  and  $s_7$ .

*Example 3 (n-Prisoners Enigma).* Finally, the last example is about  $n$  prisoners in a prison, each one in his own cell and they cannot communicate. Also, there is a room with a light bulb and a switch and a policeman that, at each moment of time, sends only one prisoner in that room and gives him the possibility to turn on or off the light. The prisoners can observe only the light when they are in the room. The guardians ensure that each prisoner is sent into the room an infinite number of times (fairness assumption). Before the game starts, the prisoners are allowed to communicate, and they know the initial state of the light. The goal of the prisoners is to learn whether all of them have visited the room at least once – more specifically, whenever all prisoners have visited the room, one specially designated prisoner must know that fact.

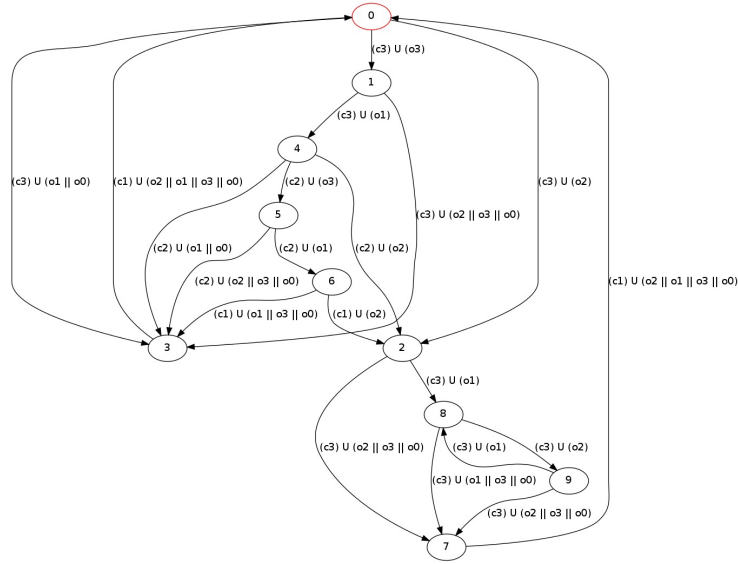


Fig. 4: The winning strategy for 3-Coins Game

Assume that the light is initially off. Then the winning strategy is that the special prisoner, say prisoner  $n$ , will count up to  $n - 1$ . For all  $1 \leq j \leq n - 1$ , the fairness assumption ensures that prisoner  $j$  will visit the room again and again until the game stops. The first time  $j$  visits the room and the light is off, he turns it on, otherwise he does nothing. Prisoner  $n$  will turn the light off next time he enters the room, and increment his counter by 1. When the counter reaches  $n - 1$ , prisoner  $n$  will be sure that all prisoners have visited the room at least once.

To model this problem, it is natural to represent the guardians by the environment and the prisoners by multi-agents. However, our framework only allows for one agent. Therefore we fix the strategy of prisoners 1 to  $n - 1$  and encode them in the environment model. Prisoner  $n$  (the system) must figure out a winning strategy (ideally the counting strategy described above). We have modelled this example in Acacia-K and indeed, our tool find the strategy described above. Let us now give more details about the formalization.

For three prisoners,  $\mathcal{P} = \{on, x_1, x_2, p_1, p_2, p_3\}$  where the atomic proposition  $on$  corresponds to the light, values of  $x_i$  for  $i \in \{1, 2\}$  is *true* if the prisoner  $i$  already turned the light on, and the proposition  $p_i$  for  $i \in \{1, 2, 3\}$  indicates the prisoner that is inside the room. Then,  $\mathcal{P}_v = \{on, p_3\}$  and  $\mathcal{P}_i = \{x_1, x_2, p_1, p_2\}$  indicating that the last prisoner sees all the time the light and can observe when he is inside the special room but cannot see what the other prisoners do.

We assume that at the beginning there is no one in the room. Then, the environment can propose an action in  $\Sigma_2 = \{P_1, P_2, P_3\}$  deciding which prisoner is going in the room and prisoner 3 will decide if he wants the light on or off by choosing an action in  $\Sigma_1 = \{t_{on}, t_{off}\}$ . Note that the action of  $P_3$  is ignored if he is not chosen by the environment. The transition relation asks that if the prisoner  $p_i, i \in \{1, 2\}$  finds for the



first time the light off ( $x_1 = false$  and  $on = false$ ), he turns on the light and the value of  $x_i$  changes and remains *true* for all the reachable states from there.

Then, assuming that the environment is restricted to send all the prisoners in the special room infinitely many times, the KLTL<sup>+</sup> formula that translates the goal is  $\Box \wedge_{i=1}^n (\Diamond p_i) \rightarrow \Diamond K(\wedge_{i=1}^{n-1} x_i)$ . A winning strategy for the prisoner  $n$  would be to turn off the light whenever he is sent to the special room and to let it off if it already is. Then, after he finds the light on  $n - 1$  times when he is sent in that room, thanks to the strategy of the other prisoners, he will know that all of them passed by that room, and even more, all of them switched an the light. Assuming that the observations set  $\mathcal{O} = \{o_0, o_1, o_2, o_3\}$  where  $o_0 = \{s \in S \mid on \in \tau(s) \text{ and } p_3 \notin \tau(s)\}$ ,  $o_1 = \{s \in S \mid on \notin \tau(s) \text{ and } p_3 \in \tau(s)\}$ ,  $o_2 = \{s \in S \mid on \in \tau(s) \text{ and } p_3 \in \tau(s)\}$  and  $o_3 = \{s \in S \mid on \notin \tau(s) \text{ and } p_3 \notin \tau(s)\}$ , the strategy synthesized by Acacia-K for three prisoners and corresponds to the intuitive strategy is illustrated in Figure 5 where the state 10 in the generated Moore machine corresponds to the moment when the prisoner  $p_3$  knows that all the other prisoners passed through the special room and turned on the light.

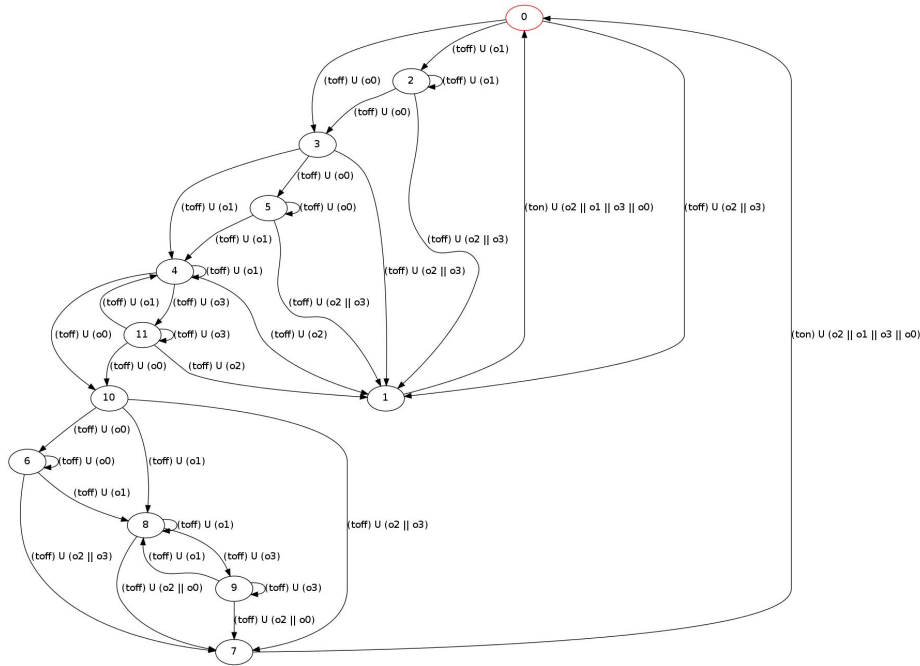


Fig. 5: The winning strategy for 3-Prisoners Game

For this example, Acacia-K constructed a UCT with 144 states, synthesised a strategy with 12 states, and the total running time is 1.87s.

We have tried 3/4/5/6 prisoners versions (including the protagonist) of this problem, obtaining a one hour timeout for 6 agents. The statistics we obtained are the following:

Pris #	$ \mathcal{M}_E $	$ \text{UCT} $	$ \text{tb} - \text{UCT} $	Aut constr (s)	$ \mathcal{M}_\lambda $	Total time(s)
3	21	144	692	1.79s	12	1.87s
4	53	447	2203	1.98s	16	13.20s
5	129	1310	6514	199.06s	20	553.45s ( $\simeq$ 9 min)
6	305	3633	18125	6081.69s	N/A	N/A

Again, Acacia-K generates strategies that are natural, the same that one would synthesize intuitively. This fact is remarkable itself since, in synthesis, it is often a difficult task to generate small and natural strategies.

## 8 Conclusion

In this paper, we have defined a Safraless procedure for the synthesis of  $\text{KLTL}^+$  specifications in environment with imperfect information. This problem is  $2\text{ExpTime-c}$  but we have shown that our procedure, based on universal co-Büchi tree automata, can be implemented efficiently thanks to an antichain symbolic approach. We have implemented a prototype and run some preliminary experiments that prove the feasibility of our method. While the UCT constructed by the tool are not small (around 1300 states), our tool can handle them, although in theory, the safety games could be exponentially larger than the UCT. Moreover, our tool synthesises small strategies that correspond to the intuitive strategies we would expect, although it goes through a non-trivial automata construction. As a future work, we want to see if Acacia-K scales well on larger examples. We also want to extend the tool to handle the full  $\text{KLTL}$  logic in an efficient way. This paper is an encouraging (and necessary) step towards this objective. In a first attempt to generalize the specifications, we plan to consider assume-guarantees specifications  $K\phi \rightarrow \psi$ , where  $\phi$  is an LTL formula and  $\psi$  a  $\text{KLTL}^+$  formula.

## References

1. Acacia-k. Available at <http://lacl.fr/~rbozianu/Acacia-K/>.
2. D. Berwanger and L. Doyen. On the power of imperfect information. In R. Hariharan, M. Mukund, and V. Vinay, editors, *FSTTCS*, volume 2 of *LIPICs*, pages 73–82. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2008.
3. R. Bloem, A. Cimatti, K. Greimel, G. Hofferek, R. Knighofer, M. Roveri, V. Schuppan, and R. Seeber. Ratsy a new requirements analysis tool with synthesis. In T. Touili, B. Cook, and P. Jackson, editors, *Computer Aided Verification*, volume 6174 of *Lecture Notes in Computer Science*, pages 425–429. Springer Berlin Heidelberg, 2010.
4. R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa’ar. Synthesis of reactive(1) designs. *J. Comput. Syst. Sci.*, 78(3):911–938, 2012.
5. A. Bohy, V. Bruyère, E. Filiot, N. Jin, and J.-F. Raskin. Acacia+, a tool for LTL synthesis. In P. Madhusudan and S. A. Seshia, editors, *CAV*, volume 7358 of *Lecture Notes in Computer Science*, pages 652–657. Springer, 2012.
6. K. Chatterjee, L. Doyen, E. Filiot, and J.-F. Raskin. Doomsday equilibria for omega-regular games. In K. L. McMillan and X. Rival, editors, *VMCAI*, volume 8318 of *Lecture Notes in Computer Science*, pages 78–97. Springer, 2014.
7. K. Chatterjee, L. Doyen, and T. A. Henzinger. A survey of partial-observation stochastic parity games. *Formal Methods in System Design*, 43(2):268–284, 2013.
8. E. Filiot, N. Jin, and J.-F. Raskin. Antichains and compositional algorithms for LTL synthesis. *Formal Methods in System Design*, 39(3):261–296, 2011.

9. B. D. Giampaolo, G. Geeraerts, J.-F. Raskin, and N. Sznajder. Safrless procedures for timed specifications. In *FORMATS*, volume 6246 of *Lecture Notes in Computer Science*, pages 2–22. Springer, 2010.
10. J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. In T. Kameda, J. Misra, J. G. Peters, and N. Santoro, editors, *PODC*, pages 50–61. ACM, 1984.
11. B. Jobstmann and R. Bloem. Optimizations for LTL synthesis. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 117–124. IEEE Computer Society, 2006.
12. O. Kupferman and N. Piterman. Lower bounds on witnesses for nonemptiness of universal co-Büchi automata. In L. de Alfaro, editor, *FOSFACS*, volume 5504 of *Lecture Notes in Computer Science*, pages 182–196. Springer, 2009.
13. O. Kupferman and M. Y. Vardi. Safrless decision procedures. In *FOCS*, pages 531–542. IEEE Computer Society, 2005.
14. W. Liu and J. Wang. A tighter analysis of piterman’s büchi determinization. *Inf. Process. Lett.*, 109(16):941–945, 2009.
15. N. Piterman. From nondeterministic büchi and streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3(3), 2007.
16. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190. ACM Press, 1989.
17. J.-F. Raskin, K. Chatterjee, L. Doyen, and T. A. Henzinger. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3(3), 2007.
18. S. Schewe and B. Finkbeiner. Bounded synthesis. In *International Symposium on Automated Technology for Verification and Analysis (ATVA)*, volume 4762 of *LNCS*, pages 474–488. Springer, 2007.
19. R. van der Meyden and M. Vardi. Synthesis from knowledge-based specifications. In D. Sangiorgi and R. Simone, editors, *CONCUR’98 Concurrency Theory*, volume 1466 of *Lecture Notes in Computer Science*, pages 34–49. Springer Berlin Heidelberg, 1998.