

# Semi-Automatic Distributed Synthesis

Bernd Finkbeiner and Sven Schewe

Universität des Saarlandes, 66123 Saarbrücken, Germany  
{finkbeiner|schewe}@cs.uni-sb.de

**Abstract.** We propose a sound and complete compositional proof rule for distributed synthesis. Applying our proof rule only requires the manual strengthening of the specification into a conjunction of formulas that can be guaranteed by individual black-box processes. All premises of the proof rule can be checked automatically.

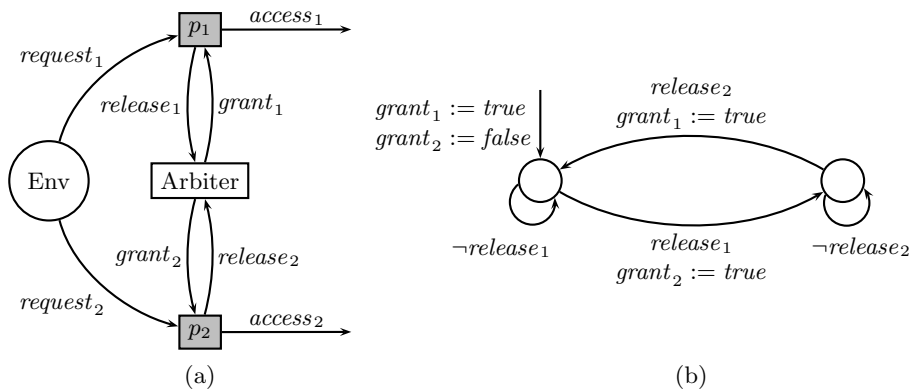
For this purpose, we give an automata-theoretic synthesis algorithm for single processes in distributed architectures. The behavior of the local environment of a process is unknown in the process of synthesis and cannot be assumed to be maximal. We therefore consider reactive environments that have the power to disable some of their own actions, and provide methods for synthesis (and realizability checking) in this setting. We establish upper bounds for CTL (2EXPTIME) and CTL\* (3EXPTIME) synthesis with incomplete information, matching the known lower bounds for these problems, and provide matching upper and lower bounds for  $\mu$ -calculus synthesis (2EXPTIME) with complete or incomplete information. Synthesis in reactive environments is harder than synthesis in maximal environments, where CTL, CTL\* and  $\mu$ -calculus synthesis are EXPTIME, 2EXPTIME and EXPTIME complete, respectively.

## 1 Introduction

In the synthesis of distributed systems, we transform a given specification into a collection of finite-state programs that satisfy the specification when composed according to a given architecture. For some restricted architectures, such as pipelines and rings in which only one designated process communicates with the environment [1], synthesis can be done automatically. However, as soon as the architecture contains an *information fork*, i.e., a pair of processes that have an incomparable degree of information about the system state, the problem becomes undecidable [2].

In this paper, we investigate a *semi-automatic* approach where we synthesize one process at a time. It turns out that the synthesis of a single process can be done automatically and it is always possible to decompose a realizable specification into a conjunction of properties that can be guaranteed by single processes. This approach therefore works for all distributed architectures, including those with information forks.

The problem of synthesizing a single process has been studied in a number of variations. *Closed synthesis* excludes any interaction with the environment



**Fig. 1.** A simple distributed shared-resource application. (a) The system architecture. An edge between two process nodes  $p$  and  $q$  labeled with variable  $v$  indicates that  $v$  is an output variable of process  $p$  and an input variable of process  $q$ . (b) The implementation of the white-box process Arbiter, represented as a finite-state automaton.

[3, 4]. *Open synthesis* finds implementations that satisfy a specification in any environment. For universal specifications (e.g., ACTL\*), it suffices to consider the *maximal* environment, which shows all possible behaviors [5, 6]. In general, it is necessary to account for *reactive environments*, which may disable some of their responses [7].

We consider the problem of synthesizing a single black-box process in a given distributed *architecture*. An architecture consists of an external environment and a set of system processes, which we partition into subsets of *white-box* and *black-box* processes: each white-box process comes with a known and fixed implementation, while the implementation of the black-box processes is yet to be found.

A single black-box process may interact with the external environment, the white-box processes, and with other black-box processes. Like in open synthesis, we assume that the behavior of the external environment is maximal. The behavior of white-box processes is known beforehand, but may be nondeterministic. The other black-box processes show reactive [7] behavior: in each state, they may disable some (but not all) of their responses. An important difference between synthesizing systems that consist of a single process, and synthesizing a single process within a general architecture is that, while the process has *complete information* about the system state in the former case, it only sees a part of the state as defined by the architecture in the latter case.

Figure 1a shows the architecture of a simple distributed *shared-resource* application. The external environment Env can request *access* to the resource by setting the *request* variable of one of the two black-box processes  $p_1$  and  $p_2$ . Mutual exclusion is accomplished using a white-box Arbiter process that alternates

a *grant* between  $p_1$  and  $p_2$ , such that each process retains the grant until the respective *release* variable is set, as shown in Figure 1b.

We can specify the expected behavior of the shared-resource system as a conjunction  $\psi = \psi_1 \wedge \psi_2 \wedge \psi_3$  of three CTL\* formulas, where the first two formulas specify that there is a way for both processes to use the resource infinitely often ( $\psi_i = \text{EGF } access_i$  for  $i \in \{0, 1\}$ ) and the third formula specifies mutual exclusion ( $\psi_3 = \text{AG } \neg(access_1 \wedge access_2)$ ).

Obviously, neither  $p_1$  nor  $p_2$  can guarantee  $\psi$  for *all* possible implementations of the other process (for example, if the other process constantly sets its *access* variable to *true*, mutual exclusion must be violated in some branch). We therefore strengthen  $\psi$  into two separate properties  $\varphi_{p_1}$  and  $\varphi_{p_2}$  that can be guaranteed by  $p_1$  and  $p_2$ , respectively. A natural assumption to be made by process  $p_{3-i}$  about process  $p_i$  is that there is path, such that process  $p_i$  infinitely often releases the grant ( $\alpha_1^{p_i} = \text{EGF } release_i$ ) and that, on every path,  $p_i$  only accesses the resource when permitted by Arbiter ( $\alpha_2^{p_i} = \text{AG } access_i \rightarrow grant_i$ ). By adding these assumptions, we obtain a strengthened specification  $\varphi = \varphi_{p_1} \wedge \varphi_{p_2}$  where

$$\varphi_{p_i} = \alpha_1^{p_i} \wedge \alpha_2^{p_i} \quad \wedge \quad (\alpha_1^{p_{3-i}} \wedge \alpha_2^{p_{3-i}} \rightarrow \psi).$$

Once the auxiliary formulas  $\varphi_{p_1}$  and  $\varphi_{p_2}$  have been defined, an implementation can be found automatically. For example, process  $p_i$  can guarantee  $\varphi_{p_i}$  against any implementation of process  $p_{3-i}$ , by setting *access<sub>i</sub>* after each *request<sub>i</sub>*, as soon as *grant<sub>i</sub>* becomes *true* and by setting *release<sub>i</sub>* in the immediately following state.

**Contribution.** We propose a sound and complete compositional proof rule for distributed synthesis. Applying our proof rule only requires the manual strengthening of the specification into a conjunction of formulas that can be guaranteed by individual black-box processes against the other black-box processes. All premises of the proof rule can be checked automatically.

For this purpose, we give an automata-theoretic synthesis algorithm for single processes in distributed architectures. Our environment model builds on open synthesis [5], but combines the maximal external environment with reactive black-box processes. Synthesis in reactive environments was studied before, but only under the assumption of complete information [7].

Our construction turns a specification into an alternating parity automaton accepting exactly the reactive models of a specification. For a specification  $\varphi$  with length  $n = |\varphi|$  in CTL, CTL\* and  $\mu$ -calculus this automaton has  $n^{O(n)}$ ,  $2^{2^{O(n)}}$  and  $n^{O(n^3)}$  states, respectively. We establish 2EXPTIME and 3EXPTIME upper bounds for synthesis with incomplete information in case of CTL and CTL\* specifications, respectively. We defer a doubly exponential lower bound for  $\mu$ -calculus specifications from the doubly exponential lower bound for CTL and establish a matching upper bound.

**Overview.** In the following section, we formally introduce the synthesis problem studied in this paper. We explain the compositional synthesis rule in Section 3. The synthesis algorithm is presented in Section 4.

## 2 Setting

In the *distributed synthesis* problem, we decide for a pair  $(A, \varphi)$ , consisting of an architecture  $A$  and a specification  $\varphi$ , whether there exists a finite-state program (or *strategy*) for each black-box process in  $A$ , such that the joint behavior satisfies  $\varphi$ .

**Architectures.** An architecture

$$A = (B, W, \{I_p\}_{p \in B \uplus W \uplus \{env\}}, \{O_p\}_{p \in B \uplus W \uplus \{env\}}, \{s_w\}_{w \in W})$$

is given as a set of processes  $P = B \uplus W \uplus \{env\}$  that is decomposed into a set  $B$  of black-box processes that have to be developed, a set  $W$  of white-box processes that already have an implementation  $\{s_w\}_{w \in W}$ , and the external environment  $env$ . The processes communicate through a set  $V$  of shared variables, which also serve as atomic propositions in the specification. Each process  $p \in P$  has a fixed set of input and output variables  $I_p, O_p \subseteq V$ , such that the family of output variables  $\{O_p\}_{p \in P}$  decomposes  $V$ . The environment is always omniscient ( $I_{env} = V$ ).

**Implementations.** A process  $p$  is implemented by a (nondeterministic) *strategy*, i.e., a function  $s_p : (2^{I_p})^* \rightarrow 2_{\emptyset}^{O_p}$  (where  $2_{\emptyset}^X = 2^X \setminus \{\emptyset\}$  denotes the nonempty subsets of a set  $X$ ). A strategy is *finite-state* if it can be represented by a finite-state automaton. The implementations  $\{s_w\}_{w \in W}$  of the white-box processes  $W$  are fixed for the architecture. An *implementation* of an architecture is a set of strategies  $S = \{s_b\}_{b \in B}$  for the black-box processes.

We use trees as a representation for strategies and computations. As usual, an  $\mathcal{Y}$ -tree is given as a prefix-closed subset  $Y \subseteq \mathcal{Y}^*$  of all finite words over a given set of directions  $\mathcal{Y}$ . If the set of directions is not important or clear from the context, we call  $Y$  a tree. We define that every non-empty node  $x \cdot v$ ,  $x \in \mathcal{Y}^*, v \in \mathcal{Y}$ , has the direction  $dir(x \cdot v) = v$  and the empty word  $\varepsilon$  has some designated *root-direction*  $dir(\varepsilon) = v_0 \in \mathcal{Y}$ . An  $\mathcal{Y}$ -tree  $Y$  is called *total*, if it contains the empty word  $\varepsilon \in Y$  and every element  $y \in Y$  of the tree has at least one successor  $y \cdot v \in Y, v \in \mathcal{Y}$ . If  $Y = \mathcal{Y}^*$ , the tree is called *full*.

For given finite sets  $\Sigma$  and  $\mathcal{Y}$ , a  $\Sigma$ -labeled  $\mathcal{Y}$ -tree is a pair  $\langle Y, l \rangle$ , consisting of a tree  $Y \subseteq \mathcal{Y}^*$  and a labeling function  $l : Y \rightarrow \Sigma$  that maps every node of  $Y$  to a letter of  $\Sigma$ . The *successor-tree*  $\langle Y, sucset \rangle$  of a tree  $Y$  is the  $2^{\mathcal{Y}}$ -labeled  $\mathcal{Y}$ -tree, where every node is labeled with the set of its successors  $sucset : Y \rightarrow 2^{\mathcal{Y}}$ ,  $sucset : y \mapsto \{v \in \mathcal{Y} \mid y \cdot v \in Y\}$ .

For a set  $\Xi \times \mathcal{Y}$  of directions and a node  $x \in (\Xi \times \mathcal{Y})^*$ ,  $hide_{\mathcal{Y}}(x)$  denotes the node in  $\Xi^*$  obtained from  $x$  by replacing  $(\xi, v)$  by  $\xi$  in each letter of  $x$ . For a  $\Sigma$ -labeled  $\Xi$ -tree  $\langle \Xi^*, l \rangle$  we define the  $\mathcal{Y}$ -widening of  $\langle \Xi^*, l \rangle$ , denoted by  $wide_{\mathcal{Y}}(\langle \Xi^*, l \rangle)$ , as the  $\Sigma$ -labeled  $\Xi \times \mathcal{Y}$ -tree  $\langle (\Xi \times \mathcal{Y})^*, l' \rangle$  with  $l'(x) = l(hide_{\mathcal{Y}}(x))$ .

We consider specifications  $\varphi$  that are given as CTL, CTL\*, or  $\mu$ -calculus formulas. Such specifications define a set  $\mathcal{M}_{\varphi}$  of total  $2^{AP}$ -labeled  $\mathcal{Y}$ -trees, where  $AP = V$  denotes the set of atomic propositions in  $\varphi$ .

Let  $S_Q = \bigotimes_{p \in Q} 2^{2^{O_p}}$  denote the set of possible common outputs of a set of strategies for the processes in  $Q \subseteq B \cup W$ . The *composition*  $\bigoplus_{p \in Q} s_p = s_Q : (2^V)^* \rightarrow S_Q$  of a set of strategies  $\{s_p\}_{p \in Q}$  maps the global input history to the common output of the processes in  $Q$ : For  $\langle (2^V)^*, s'_p \rangle = \text{wide}_{2^{V \setminus I_p}}(\langle (2^V)^*, s'_p \rangle)$ ,  $s_Q : y \mapsto \bigoplus_{p \in Q} s'_p(y)$  naturally defines a  $S_Q$ -labeled  $2^V$ -tree.

A *non-distributed implementation* of the processes  $B'$  is a function

$$s_{B'} : (2^{I_{B'}})^* \rightarrow S_{B'}, \text{ for } I_{B'} = \bigcup_{b \in B'} I_b.$$

A (*distributed*) *implementation* is a set of strategies  $\{s_b\}_{b \in B'}$  whose composition is the widening of a non-distributed implementation:  $\langle (2^V)^*, \bigoplus_{b \in B'} s_b \rangle = \text{wide}_{2^{V \setminus I_{B'}}}(\langle (2^{I_{B'}})^*, s_{B'} \rangle)$  for some non-distributed implementation  $s_{B'}$ .

**Realizability.** An implementation  $s_{B'}$  of a set  $B' \subseteq B$  of black-box processes *guarantees*  $\varphi$  *against the remaining black-box processes*, if for all  $S_{B \setminus B'}$ -labeled  $2^V$ -trees  $\langle (2^V)^*, s_{B \setminus B'} \rangle$ , the total  $2^V$ -labeled  $2^V$ -tree  $\langle Y, \text{dir} \rangle$ , whose branching restriction  $\text{sucset}(y) = (s_{B'} \oplus \bigoplus_{w \in W} s_w \oplus s_{B \setminus B'})(y) \times 2^{O_{\text{env}}}$  is defined by the strategies, is a model of  $\varphi$ .

We say that a specification is *realizable by the processes*  $B' \subseteq B$  for a given architecture  $(B, W, \{I_p\}_{p \in B \cup W}, \{O_p\}_{p \in B \cup W}, O_{\text{env}}, \{s_w\}_{w \in W})$ ,  $(A, B') \models \varphi$ , if there is a distributed implementation of the processes  $B'$  that guarantees  $\varphi$  against  $B \setminus B'$ . A specification is *realizable* if it is realizable by the entire set of black-box processes  $B$ .

### 3 A Compositional Synthesis Rule

The compositional synthesis rule reduces the realizability of a distributed system,  $(A, B) \models \psi$ , to the realizability of single processes,  $(A, \{b\}) \models \varphi_b$ , for each black-box process  $b \in B$ . The proof rule requires an auxiliary specification  $\varphi_b$  for each process  $b \in B$ . If each process  $b$  guarantees  $\varphi_b$  against the remaining black-box processes, the distributed system can be implemented to satisfy  $\psi$ .

For a distributed architecture  $A$  with a set of black-box processes  $B = \{b_1, \dots, b_n\}$ , and CTL\* or  $\mu$ -calculus formulas  $\psi, \varphi_{b_1}, \dots, \varphi_{b_n}$ ,

$$\begin{array}{l} \text{(R0)} \quad (A, \emptyset) \models \bigwedge_{b \in B} \varphi_b \rightarrow \psi \\ \text{(R1)} \quad (A, \{b_1\}) \models \varphi_{b_1} \\ \vdots \\ \text{(Rn)} \quad (A, \{b_n\}) \models \varphi_{b_n} \\ \hline (A, B) \models \psi \end{array}$$

Premise (R0) shows that the auxiliary formulas  $\varphi_{b_1}, \dots, \varphi_{b_n}$  strengthen the original formula  $\varphi$  and hence any implementation that satisfies  $\varphi_{b_1}, \dots, \varphi_{b_n}$  must also satisfy  $\varphi$ . Premises (R1) through (Rn) prove that there are, for all  $b_i$  in  $B$ , strategies  $s_{b_i}$  that guarantee  $\varphi_{b_i}$  against the remaining black-box processes.

**Theorem 1.** *The proof rule is sound.*

*Proof.* Premises (R1) through (Rn) guarantee that, for each  $b \in B$ , there is an implementation  $s_b$  that guarantees  $\varphi_b$  against the remaining black-box processes  $B \setminus \{b\}$ . Consequently, the strategies can be fixed independently; the distributed implementation thus obtained satisfies  $\varphi_b$  for all  $b \in B$  and hence  $\bigwedge_{b \in B} \varphi_b$ . Premise (R0) guarantees that every non-distributed implementation of  $\bigwedge_{b \in B} \varphi_b$  is also an implementation of  $\psi$ . As the distributed implementations form a subset of the non-distributed implementations, the claim holds true.  $\square$

To show the completeness of the distribution rule, we derive the auxiliary formulas from a given implementation that realizes the specification: for a given architecture, we call a specification *strict*, if it completely determines its implementation. An implementation can be described by a strict LTL specification  $\varphi$ . A distributed implementation can be described by a strict specification  $\varphi_b$  for every black-box component  $b \in B$ , such that  $\varphi = \bigwedge_{b \in B} \varphi_b$  is a strict specification for the implementation.

**Theorem 2.** *The proof rule is complete.*

*Proof.* Assume there is a distributed implementation for a specification  $\psi$  and  $\varphi = \bigwedge_{b \in B} \varphi_b$  is a strict specification for this implementation. Then  $(A, \{b\}) \models \varphi_b$  holds true for each  $b \in B$ . The implementation of  $\varphi$  is completely determined and  $(A, \emptyset) \models \varphi \rightarrow \psi$  requires that every specification of  $\varphi$  is an implementation of  $\psi$ . As the unique implementation is by definition an implementation of  $\psi$ ,  $(A, \emptyset) \models \varphi \rightarrow \psi$  also holds true.  $\square$

## 4 Single-Process Synthesis

We now develop a procedure that checks if a specification can be guaranteed by a single black-process  $b$  against the remaining black-box processes,  $(A, \{b\}) \models \varphi$ , as required for premises (R1) through (Rn), and a procedure that checks if a specification can be guaranteed by the empty set of black-processes against all black-box processes,  $(A, \emptyset) \models \varphi$ , as required for premise (R0).

Every formula of a temporal logic can be translated into an alternating tree automaton that accepts exactly its set of models. This automaton is the starting point for our construction, which consists of a series of tree automata transformations.

## 4.1 Tree automata

An *alternating parity tree automaton* is a tuple  $\mathcal{A} = (\Sigma, Q, q_0, \delta, \alpha)$ , where  $Q$  denotes a finite set of states,  $q_0 \in Q$  denotes a designated initial state,  $\delta$  denotes a transition function, and  $\alpha : Q \rightarrow C \subset \mathbb{N}$  is a coloring function. The transition function  $\delta : Q \times \Sigma \rightarrow \mathbb{B}^+(Q \times \mathcal{Y})$  maps a state and an input letter to a positive boolean combination of states and directions (for a predefined finite set  $\mathcal{Y}$  of directions).

An alternating automaton runs on full  $\Sigma$ -labeled  $\mathcal{Y}$ -trees. A *run tree*  $\langle R, r \rangle$  on a given full  $\Sigma$ -labeled  $\mathcal{Y}$ -tree  $\langle \mathcal{Y}^*, l \rangle$  is a  $Q \times \mathcal{Y}^*$ -labeled tree where the root is labeled with  $(q_0, \varepsilon)$  and where, for each node  $n$  with a label  $(q, y)$  with the set of labels of its successors  $L = \{r(n \cdot \rho) \mid \rho \in \text{succset}(n)\}$ , there is a set  $A \subseteq 2^{Q \times \mathcal{Y}}$  which satisfies  $\delta(q, l(y))$  such that  $(q', v) \in A \Leftrightarrow (q', y \cdot v) \in L$ .

An infinite path fulfills the *parity condition*, if the highest color of the states appearing infinitely often on the path is even. A run tree is *accepting* if all infinite paths fulfill the parity condition. A total  $\Sigma$ -labeled  $\mathcal{Y}$ -tree is accepted if it has an accepting run tree.

The set of trees accepted by an alternating automaton  $\mathcal{A}$  is called its *language*  $\mathcal{L}(\mathcal{A})$ .  $\overline{\mathcal{L}(\mathcal{A})}$  denotes the set of full  $\Sigma$ -labeled  $\mathcal{Y}$ -trees not accepted by  $\mathcal{A}$ . An automaton is empty, if its language is empty.

The acceptance of a tree can also be viewed as the outcome of a game, where player *accept* chooses, for a pair  $(q, \sigma) \in Q \times \Sigma$ , a set of atoms of  $\delta(q, \sigma)$ , satisfying  $\delta(q, \sigma)$ , and player *reject* chooses one of these atoms, which is executed. The input tree is accepted iff player *accept* has a strategy enforcing a path that fulfills the parity condition. One of the player has a memoryless winning strategy, i.e., a strategy where the moves only depend on the state of the automaton, the position in the tree and, for player *react*, on the choice of player *accept* in the same move.

A *nondeterministic* automaton is a special alternating automaton, where the image of  $\delta$  consists only of such formulae that, when rewritten in disjunctive normal form, contain exactly one element of  $Q \times \{v\}$  for all  $v \in \mathcal{Y}$  in every disjunct.

For nondeterministic automata, every node of a run tree corresponds to a node in the input tree. Emptiness can therefore be checked with an *emptiness game*, where player *accept* also chooses the letter of the input alphabet. A nondeterministic automaton is empty iff the emptiness game is won by *reject*.

*Symmetric alternating automata* are a variant of alternating automata that run on total  $\Sigma$ -labeled  $\mathcal{Y}$ -trees. For a symmetric alternating automaton  $\mathcal{S} = (\Sigma, Q, q_0, \delta, \alpha)$ ,  $Q$ ,  $q_0$ , and  $\alpha$  are defined as before. The transition function  $\delta : Q \times \Sigma \rightarrow \mathbb{B}^+(Q \times \{\square, \diamond\})$  now maps a state and an input letter to a positive boolean combination over atoms that refer to *some* ( $\diamond$ ) or *all* ( $\square$ ) successor states.

A *run tree* on a given  $\Sigma$ -labeled  $\mathcal{Y}$ -tree  $\langle R, r \rangle$  is a  $Q \times \mathcal{Y}^*$ -labeled tree where the root is labeled with  $(q_0, \varepsilon)$  and where, for a node  $n$  with a label  $(q, y)$  and a set of labels of its successors  $L = \{r(n \cdot \rho) \mid \rho \in \text{succset}(n)\}$ , the following property holds: there is a set of atoms  $A \subseteq 2^{Q \times \{\square, \diamond\}}$  satisfying  $\delta(q, l(y))$  such

that  $\forall q' \in Q. ((q', \square) \in A \Rightarrow \forall v \in \text{succset}(x). (q', y \cdot v) \in L) \wedge ((q', \diamond) \in A \Rightarrow \exists v \in \text{succset}(x). (q', y \cdot v) \in L)$ .

We introduce a function  $\text{succ} : (Q \times \Sigma \rightarrow \mathbb{B}^+(Q \times \{\square, \diamond\})) \rightarrow (Q \times \Sigma \times 2_\emptyset^{\mathcal{X}} \rightarrow \mathbb{B}^+(Q \times \mathcal{Y}))$  that translates the transition function of a symmetric alternating automaton running on total  $\Sigma$ -labeled  $\mathcal{Y}$ -trees into the corresponding transition function of an alternating automaton running on full  $\Sigma \times 2_\emptyset^{\mathcal{X}}$ -labeled  $\mathcal{Y}$ -trees. For the set  $2_\emptyset^{\mathcal{X}} = 2^{\mathcal{X}} \setminus \{\emptyset\}$  of possible sets of successors,  $\text{succ}(\delta) : Q \times \Sigma \times 2_\emptyset^{\mathcal{X}} \rightarrow \mathbb{B}^+(Q \times \mathcal{Y})$  maps a state, an input letter and a set of successors to a positive boolean combination of states and directions.

## 4.2 Overview

We represent the joint behavior of a system as a total  $2^V$ -labeled  $2^V$ -tree  $\langle Y, \text{dir} \rangle$ , where the label is completely determined by the direction. The process strategies determine the tree: By the proper widening of a strategy  $s'_p : (2^V)^* \rightarrow 2_\emptyset^{2^{O_p}}$ , each input history (or initial sequence of a path) is mapped to a nonempty subset of  $2^{O_p}$ , restricting the set of successors. The nodes of  $Y$  consist of the root and all nodes  $y \cdot v$  whose predecessor  $y$  is in  $Y$ , and whose direction agrees with the decisions of the processes:  $y \cdot v \in Y \Leftrightarrow y \in Y \wedge \forall p \in B \uplus W. v \cap O_p \in s'_p(y)$ .

We start our construction with a symmetric automaton  $\mathcal{S}_\varphi$  that accepts the models of the specification  $\varphi$ . Automata transformations are simpler for automata running on full trees; we therefore represent total trees as full trees by decorating each node with its own set of successors. Considering a full  $2^V \times 2_\emptyset^{2^V}$ -labeled  $2^V$ -tree  $\langle (2^V)^*, l' \rangle$ , where the nodes are additionally decorated with the sets of relevant successors, one can easily determine the original total  $2^V$ -labeled  $2^V$ -tree  $\langle Y, l \rangle$ , which we call its *characteristic tree*.

We continue with an automaton that accepts those full  $2^V \times 2_\emptyset^{2^V}$ -labeled  $2^V$ -trees whose characteristic tree is a model of  $\varphi$ . The labeling of the nodes  $(2^V)^* \setminus Y$  of  $\langle (2^V)^*, l' \rangle$  that are not on the characteristic tree has no influence on the acceptance of the tree. We restrict the language under consideration to  $2^V \times S_b \times S_W \times S_{B'} \times S$ -labeled  $2^V$ -trees, where  $S_b, S_W, S_{B'}$  and  $S$  describe the possible restrictions on the successor sets induced by the black-box process  $b$ , the set of white-box processes  $W$ , the remaining black-box processes  $B' = B \setminus \{b\}$ , and the environment, respectively. By that, we obtain an automaton  $\mathcal{A}_\varphi$  that accepts  $2^V \times S_b \times S_W \times S_{B'} \times S$ -labeled  $2^V$ -trees. Since  $S_p = 2_\emptyset^{2^{O_p}}$  for all processes<sup>1</sup>  $p \in B \uplus W$ , the sets of possible restrictions can be identified with  $S_W = \bigotimes_{w \in W} S_w$ ,  $S_{B'} = \bigotimes_{b \neq b' \in B} S_{b'}$  and  $S = \{2^{O_{env}}\}$  (as we assume the environment to be maximal).

It remains to find a strategy  $s_b$  such that for its proper widening  $s'_b$ ,  $\langle (2^V)^*, \text{dir} \times s'_b \times (\bigoplus_{w \in W} s_w) \times s_{B'} \times \{O_{env}\} \rangle$  is accepted for all strategies  $s_{B'} : (2^V)^* \rightarrow \bigotimes_{b \neq b' \in B} 2_\emptyset^{2^{O_{b'}}}$ .

<sup>1</sup> For generality, we allow all processes to be nondeterministic. If a subset  $D \subseteq B \uplus W$  of the processes is to be deterministic, one can simply choose, for all  $p \in D$ , the set of singleton subsets of  $2^{O_p}$  instead of the set of non-empty subsets.



We first build an automaton  $\mathcal{R}_\varphi$  that accepts a  $2^V \times S_b \times S_W$ -labeled  $2^V$ -tree, if its complete cylinder is accepted by  $\mathcal{A}_\varphi$ , establishing independence from the decision of the black-box processes. Subsequently, we use the determination of the  $2^V$  and  $S_W$  fraction of the label to defer an automaton  $\mathcal{D}_\varphi$  that accepts all strategy trees  $\langle (2^V)^*, s_b \rangle$  that would guarantee  $\varphi$  against the remaining black-box processes if process  $b$  were omniscient. This automaton is then transformed into an automaton  $\mathcal{B}_\varphi$  accepting the strategies of  $b$  ( $2_\emptyset^{2^{O_b}}$ -labeled  $2^{I_b}$ -trees).

Checking this automaton for emptiness answers the question of realizability. In case of realizability, the emptiness test can be extended to synthesize a finite-state strategy for  $b$ .

In summary, our construction consists of seven steps:

1. **From formulas to automata:** We construct a symmetric alternating automaton  $\mathcal{S}_\varphi$  that accepts the models of  $\varphi$ .
2. **Characteristic trees:** The alternating automaton  $\mathcal{A}_\varphi$  accepts a  $2^V \times S_b \times S_W \times S_{B'} \times S$ -labeled  $2^V$ -tree if its characteristic tree is accepted by  $\mathcal{S}_\varphi$ .
3. **Quantification:** The alternating automaton  $\mathcal{R}_\varphi$  accepts a  $2^V \times S_b \times S_W$ -labeled  $2^V$ -tree if all  $S_{B'} \times S$  extensions are accepted by  $\mathcal{A}_\varphi$ .
4. **Adjusting for white-box processes:** The alternating automaton  $\mathcal{W}_\varphi$  accepts a  $2^V \times S_b$ -labeled  $2^V$ -tree if the  $2^V \times S_b \times S_W$ -labeled  $2^V$ -tree obtained by adding the decisions of the white-box processes is accepted by  $\mathcal{R}_\varphi$ .
5. **Pruning directions from the labeling:** The alternating automaton  $\mathcal{D}_\varphi$  accepts a  $S_b$ -labeled  $2^V$ -tree if the  $2^V \times S_b$ -labeled  $2^V$ -tree obtained by adding the direction of a node to the label is accepted by  $\mathcal{W}_\varphi$ .
6. **Narrowing:** The alternating automaton  $\mathcal{B}_\varphi$  accepts a  $S_b$ -labeled  $2^{I_b}$ -tree if its proper widening is accepted by  $\mathcal{D}_\varphi$ .
7. **Emptiness check:** The realizability claim  $(A, \{b\}) \models \varphi$  holds true iff  $\mathcal{B}_\varphi$  is not empty. To perform an emptiness test,  $\mathcal{B}_\varphi$  can be transformed into an equivalent nondeterministic automaton  $\mathcal{C}_\varphi$ , which can be checked for emptiness by solving the emptiness game. A winning strategy in the emptiness game implies an implementation for the process  $b$ .

In the following, we discuss the automata transformations in detail.

### 4.3 Automata Transformations

**From formulas to automata.** We use standard constructions to translate a temporal specification  $\varphi$  into a symmetric alternating automaton  $\mathcal{S}_\varphi$  that accepts the models of the formula:  $\mathcal{L}(\mathcal{S}_\varphi) = \mathcal{M}_\varphi$ .

**Theorem 3.** *Given a CTL specification  $\varphi$ , we can construct a symmetric alternating automaton  $\mathcal{S}_\varphi$  with  $O(|\varphi|)$  states and two colors such that  $\mathcal{L}(\mathcal{S}_\varphi) = \mathcal{M}_\varphi$  [8]. Given a CTL\* specification  $\varphi$ , we can construct a symmetric alternating automaton  $\mathcal{S}_\varphi$  with  $2^{O(|\varphi|)}$  states and five colors such that  $\mathcal{L}(\mathcal{S}_\varphi) = \mathcal{M}_\varphi$  [8]. Given a  $\mu$ -calculus specification  $\varphi$ , we can construct a symmetric alternating automaton  $\mathcal{S}_\varphi$  with  $O(|\varphi|^2)$  states and  $O(|\varphi|)$  colors such that  $\mathcal{L}(\mathcal{S}_\varphi) = \mathcal{M}_\varphi$  [6].  $\square$*

**Characteristic trees.** For a  $\Sigma \times \Xi$ -labeled  $\Upsilon$ -tree  $\langle Y, l \rangle$ , we denote the  $\Sigma$ -projection  $proj_\Sigma : \langle Y, l \rangle \mapsto \langle (Y, l_\Sigma) \rangle$  with  $l(y) = (\sigma, \xi) \Rightarrow l_\Sigma : y \mapsto \sigma$  that maps  $\Sigma \times \Xi$ -labeled  $\Upsilon$ -trees to  $\Sigma$ -labeled  $\Upsilon$ -trees.

For a full  $\Sigma \times 2_0^\Upsilon$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$ , we define the *characteristic tree* as the total  $\Sigma$ -labeled  $\Upsilon$ -tree  $\langle Y, l_c \rangle = char(\langle \Upsilon^*, l \rangle)$  to be the sub-tree of  $proj_\Sigma(\langle \Upsilon^*, l \rangle)$  with  $y \in Y \Rightarrow \forall v \in \Upsilon.y \cdot v \in \Upsilon \Leftrightarrow v \in proj_{2_0^\Upsilon}(\langle \Upsilon^*, l \rangle)$ . Intuitively, the second argument in the label defines the set of successors of a node.

**Lemma 1.** *Given a symmetric alternating automaton  $\mathcal{S} = (\Sigma, Q, q_0, \delta, \alpha)$ , running on total  $\Sigma$ -labeled  $\Upsilon$ -trees, we can construct an alternating automaton  $\mathcal{A} = (\Sigma \times 2_0^\Upsilon, Q, q_0, suc(\delta), \alpha)$  that accepts a full  $\Sigma \times 2_0^\Upsilon$  labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$ , iff  $proj_\Sigma(char(\langle \Upsilon^*, l \rangle))$  is accepted by  $\mathcal{S}$ .*

*Proof.* Let  $\langle T, l_T \rangle = char(\langle \Upsilon^*, l \rangle)$ . Then the successor set of a node  $x \in T$  is defined by the label:  $sucset(x) = proj_{2_0^\Upsilon}(l_{T(x)}) = proj_{2_0^\Upsilon}(l(x))$ .  $\square$

**Quantification.** To construct an alternating automaton  $\mathcal{R}_\varphi$  that accepts a  $2^V \times S_b \times S_W$ -labeled  $2^V$ -tree if all  $S_{B'} \times S$  extensions are accepted by  $\mathcal{A}_\varphi$ , we

1. complement  $\mathcal{A}_\varphi$ , i.e., we compute an alternating automaton  $\mathcal{I}_\varphi$  with  $\mathcal{L}(\mathcal{I}_\varphi) = \overline{\mathcal{L}(\mathcal{A}_\varphi)}$ ,
2. build a nondeterministic automaton  $\mathcal{N}_\varphi$  with the same language  $\mathcal{L}(\mathcal{N}_\varphi) = \mathcal{L}(\mathcal{A}_\varphi)$ ,
3. compute a nondeterministic automaton  $\mathcal{P}_\varphi$  that accepts a  $2^V \times S_b \times S_W$ -labeled  $2^V$ -tree if it is the the  $S_{B'} \times S$ -projection of a tree accepted by  $\mathcal{N}_\varphi$ ,
4. complement  $\mathcal{P}_\varphi$ , i.e., we compute an alternating automaton  $\mathcal{R}_\varphi$  with  $\mathcal{L}(\mathcal{R}_\varphi) = \overline{\mathcal{L}(\mathcal{P}_\varphi)}$ .

**Lemma 2.** [9] *Given an alternating automaton  $\mathcal{A} = (\Sigma, Q, q_0, \delta, \alpha)$  that runs on  $\Sigma$ -labeled  $\Upsilon$ -trees, the dual automaton  $\mathcal{I} = (\Sigma, Q, q_0, \bar{\delta}, \alpha + 1)$ , where  $\bar{\delta}$  is the function dual to  $\delta$ , accepts a tree  $\langle \Upsilon^*, l \rangle$  iff  $\langle \Upsilon^*, l \rangle$  is not accepted by  $\mathcal{A}$ .  $\square$*

**Lemma 3.** [2, 10] *Given an alternating automaton  $\mathcal{A}$  with  $n$  states and  $c$  colors, we can construct an equivalent nondeterministic automaton  $\mathcal{N}$  with  $n^{O(c \cdot n)}$  states and  $O(c \cdot n)$  colors.  $\square$*

**Lemma 4.** *Given a nondeterministic automaton  $\mathcal{N} = (\Sigma \times \Xi, Q, q_0, \delta, \alpha)$  that runs on  $\Sigma \times \Xi$ -labeled  $\Upsilon$ -trees, we can construct a nondeterministic automaton  $\mathcal{P} = (\Sigma, Q, q_0, \delta', \alpha)$  that accepts a  $\Sigma$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$  iff there is a  $\Sigma \times \Xi$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l_\Xi \rangle$  accepted by  $\mathcal{N}$  with  $\langle \Upsilon^*, l \rangle = proj_\Sigma(\langle \Upsilon^*, l_\Xi \rangle)$ .*

*Proof.*  $\mathcal{P}$  can be constructed by using  $\delta'$  to guess the correct tree: we set  $\delta' : (q, \sigma) \mapsto \bigvee_{\xi \in \Xi} \delta(q, (\sigma, \xi))$ .  $\square$

In the following two transformations, the decisions of the white-box processes and the labeling imposed by the directions are deleted from the label.

**Adjusting for white-box processes.** The  $S_W$  fraction of the label represents the decisions made by the white box processes. Consequently, we are only interested in those trees, where the label of every node is in accordance with these decisions. This information is then redundant and can be pruned. We assume that the composed strategy  $\bigoplus_{w \in W} s_w$  of the white-box processes is represented as a finite-state automaton  $\mathcal{O} = (2^V, O, o_0, d_W, o_W)$ , where  $O$  is a set of states,  $o_0$  the initial state, the transition function  $d_W : 2^V \times O \rightarrow O$  is a mapping from the input alphabet and the set of states to the set of states, and the output function  $o_W : O \rightarrow 2_{\emptyset}^{O_W}$  maps each state to a nonempty set of output letters. The following operation performs the pruning; the state-space of the resulting automaton is linear in the state-space of the original automaton and the number of states of  $\mathcal{O}$ , while the set of colors remains unchanged.

**Lemma 5.** *Given an alternating automaton  $\mathcal{R} = (\Sigma \times \Xi, Q, q_0, \delta, \alpha)$  over  $\Sigma \times \Xi$ -labeled  $\Upsilon$ -trees and a finite automaton  $\mathcal{O} = (\Sigma, O, o_0, d_W, o_W)$  that produces a  $\Xi$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$ , we can construct an alternating automaton  $\mathcal{W} = (\Sigma, Q \times O, (q_0, o_0), \delta', \alpha')$  over  $\Sigma$ -labeled  $\Upsilon$ -trees, such that  $\mathcal{W}$  accepts  $\langle \Upsilon^*, l' \rangle$  iff  $\mathcal{R}$  accepts  $\langle \Upsilon^*, l'' \rangle$  with  $l'' : y \mapsto (l'(y), l(y))$ .*

*Proof.* If  $\delta : (q, \sigma, \xi) \mapsto b_{(q, \sigma, \xi)}(\{q_i, v_i\}_{i \in I})$ , we can set  $\delta' : (q, o, \sigma) \mapsto b_{(q, \sigma, o_W(o))}(\{q_i, d_W(\sigma, o), v_i\}_{i \in I})$ . The coloring function can simply be set to  $\alpha' : (q, o) \mapsto \alpha(q)$ .  $\square$

**Pruning directions from the labeling.** We are only interested in those trees where the label of every node is in accordance with its direction. This information then becomes redundant and can be pruned. The following operation performs this pruning; the state-space of the resulting automaton is linear in the state-space of the original automaton, while the set of colors remains unchanged.

For a  $\Sigma$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$ , we define the function  $xray : \langle \Upsilon^*, l \rangle \mapsto \langle \Upsilon^*, l' \rangle$  with  $l'(x) = (dir(x), l(x))$  that maps  $\Sigma$ -labeled  $\Upsilon$ -trees to  $\Upsilon \times \Sigma$ -labeled  $\Upsilon$ -trees.

**Lemma 6.** [8] *Given an alternating automaton  $\mathcal{W} = (\Upsilon \times \Sigma, Q, q_0, \delta, \alpha)$  over  $\Upsilon \times \Sigma$ -labeled  $\Upsilon$ -trees, we can construct an alternating automaton  $\mathcal{D} = (\Sigma, Q \times \Upsilon, (q_0, v_0), \delta', \alpha')$  over  $\Sigma$ -labeled  $\Upsilon$ -trees, such that  $\mathcal{D}$  accepts  $\langle \Upsilon^*, l \rangle$  iff  $\mathcal{R}$  accepts  $xray(\langle \Upsilon^*, l \rangle)$ .  $\square$*

The transition function  $\delta' : Q \times \Upsilon \times \Sigma \rightarrow \mathcal{B}^+(Q \times \Upsilon \times \Upsilon)$  can be constructed from  $\delta : Q \times \Upsilon \times \Sigma \rightarrow \mathcal{B}^+(Q \times \Upsilon)$  by replacing all occurrences of  $(q, v)$  in each  $\delta(q', v', \sigma')$  by  $(q, v, v)$ , storing the direction as quasi-input.  $\alpha' : (q, c) \mapsto \alpha(q)$  simply evaluates the first component of the new state-space.

**Narrowing.** The process  $b$  is in general not omniscient, and its output may only depend on the history of the input visible to  $b$ . The following transformation therefore accepts a  $2^{O_p}$ -labeled  $2^{I_p}$ -tree if its proper widening is accepted by  $\mathcal{D}_\varphi$ . The state-space and the set of colors remain unchanged.

**Lemma 7.** [8] *Given an alternating automaton  $\mathcal{D} = (\Sigma, Q, q_0, \delta, \alpha)$  over  $\Sigma$ -labeled  $\Xi \times \mathcal{T}$ -trees, we can construct an alternating automaton  $\mathcal{B} = (\Sigma, Q, q_0, \delta', \alpha)$  over  $\Sigma$ -labeled  $\Xi$ -trees, such that  $\mathcal{B}$  accepts  $\langle \Xi^*, l \rangle$  iff  $\mathcal{W}$  accepts  $\text{wide}_{\mathcal{T}}(\langle \Xi^*, l \rangle)$ .  $\square$*

$\delta'$  can be constructed from  $\delta$  by replacing all occurrences of  $(q, (\xi, v))$  by  $(q, \xi)$  in  $\delta(q', \sigma)$  for all  $q, q' \in Q, \sigma \in \Sigma, \xi \in \Xi$  and  $v \in \mathcal{T}$ .

**Emptiness check.** To perform an emptiness test,  $\mathcal{B}_\varphi$  can be transformed into an equivalent nondeterministic automaton  $\mathcal{C}_\varphi$ .

**Theorem 4.** *Given a symmetric alternating automaton  $\mathcal{S}_\varphi$  that accepts the models of  $\varphi$ , an architecture  $(B, W, \{I_p\}_{p \in B \uplus W \uplus \{env\}}, \{O_p\}_{p \in B \uplus W \uplus \{env\}}, \{s_w\}_{w \in W})$  and a designated black-box process  $b \in B$ , we can construct a nondeterministic automaton  $\mathcal{C}_\varphi$  that accepts a full  $2^{O_p}$ -labeled  $2^{I_p}$ -tree  $\langle (2^{I_p})^*, s_b \rangle$  iff  $s_b$  guarantees  $\varphi$  against  $B \setminus \{b\}$ . If  $\mathcal{S}$  has  $n$  states and  $c$  colors,  $\mathcal{C}$  has  $2^{n^{O(n \cdot c)}}$  states and  $n^{O(n \cdot c)}$  colors.*

*Proof.* By applying the transformation steps in the order described in the overview of the algorithm, we obtain an alternating automaton  $\mathcal{B}_\varphi$  with  $n^{O(n \cdot c)}$  states and  $O(n \cdot c)$  colors that accepts an implementation  $\langle (2^{I_b})^*, s_b \rangle$  of a process  $b$  if it guarantees  $\varphi$  against  $B \setminus \{b\}$ . A nondeterminisation of  $\mathcal{B}_\varphi$  by the construction of Lemma 3 provides the required automaton.  $\square$

**Theorem 5.** *For a given architecture  $A$  and a black-box process  $b$ , we can check  $(A, \{b\}) \models \varphi$  and, if the claim is true, provide an implementation for  $b$  guaranteeing  $\varphi$ , in 2EXPTIME in the length  $|\varphi|$  if  $\varphi$  is a CTL or  $\mu$ -calculus specification, and in 3EXPTIME in  $|\varphi|$  if  $\varphi$  is a CTL\* specification, respectively.*

*Proof.* By Theorem 3, we can turn a specifications  $\varphi$  in CTL,  $\mu$ -calculus or CTL\* with length  $n = |\varphi|$  into a symmetric alternating automaton  $\mathcal{S}$  with  $O(n)$  states and two colors,  $O(n^2)$  states and  $O(n)$  colors or  $2^{O(n)}$  states and five colors, respectively.

By Theorem 4, we can transform the symmetrical alternating automaton  $\mathcal{S}$  into a nondeterministic automaton  $\mathcal{C}$ , accepting the strategies of  $b$  that guarantee  $\varphi$  against the remaining black-box processes.  $\mathcal{C}$  has  $2^{n^{O(n)}}$  states and  $n^{O(n)}$  colors,  $2^{n^{O(n^3)}}$  states and  $n^{O(n^3)}$  colors or  $2^{2^{O(n)}}$  states and  $2^{2^{O(n)}}$  colors, respectively.

The actual emptiness test or the synthesis of a strategy for process  $n$  can be done in time polynomial in the state-space and exponential in the number of colors. More precisely, if  $\mathcal{C}$  has  $m$  states and  $c$  colors, a strategy (or the proof of emptiness) can be found in  $m^{O(c)}$  time [11]. The overall time complexity is hence  $2^{n^{O(n)}}$ ,  $2^{n^{O(n^3)}}$  and  $2^{2^{2^{O(n)}}}$ , respectively.  $\square$

**Lower Bounds.** To demonstrate that the upper bounds are sharp, we give a reduction from the synthesis problem in reactive environments with complete information, which is known to be 2EXPTIME and 3EXPTIME hard for CTL

and CTL\*, respectively [7]. In synthesis with reactive environments and complete information, we have only one process  $b$ , for which a (deterministic) strategy  $s_b : (2^{O_{env}})^* \rightarrow S_b$  is sought (where  $S_b$  is the set of singleton subsets of  $2^{O_{env}}$ ). The environment can react on the input by restricting its actions to a non-empty subset of its output variables  $O_e$ , which can be viewed as a non-deterministic strategy  $s_e : (2^{O_{env} \cup O_b})^* \rightarrow 2_{\emptyset}^{2^{O_e}}$ . In our terms, a strategy  $s_b : (2^{O_{env}})^* \rightarrow S_b$  implements a specification  $\varphi$  if, for all strategies  $s_e : (2^{O_{env} \cup O_b})^* \rightarrow 2_{\emptyset}^{2^{O_e}}$  of the environment,  $s_b \times s_e$  is a model of  $\varphi$ .

We encode this synthesis problem as the realizability of  $\varphi$  by  $b$  against a black-box process  $e$  with output  $O_e$  and an environment without output. The second black-box process  $e$  plays the rôle of the reactive environment. Formally, we define the architecture  $A = (\{b, e\}, \emptyset, \{I_b = O_e, I_e = I_{env} = V\}, \{O_b, O_e, O_{env} = \emptyset\}, \emptyset)$ . The determinacy of  $s_b$  can be guaranteed by the construction (by setting  $S_b$  to the set of singleton subsets of  $2^{O_b}$ ). Alternatively, we can ensure the determinacy of  $s_b$  by strengthening the specification  $\varphi$  such that only deterministic strategies are allowed: For  $\psi = \bigwedge_{o \in O_b} AG (EXo \rightarrow AXo)$ , we can solve the realizability problem for  $\varphi' = \varphi \wedge \psi$  (which is linear in  $\varphi$ ).

**Theorem 6.** *The realizability problem  $(A, \{b\}) \models \varphi$  is 3EXPTIME complete for CTL\* and 2EXPTIME complete for CTL and  $\mu$ -calculus specifications in the size  $|\varphi|$  of the specification.*

*Proof.* The lower bounds for CTL and CTL\* follow from the equal lower bounds for the synthesis problem with reactive environments. The lower bound for the  $\mu$ -calculus is established by the lower bound for CTL. The upper bound is demonstrated by Theorem 5.  $\square$

**Premise R0.** The correctness of premise R0 can be checked along the same lines: we check whether the empty strategy guarantees  $\bigwedge_{b \in B} \varphi_b \rightarrow \psi$  against all black-box processes. Since  $S_b = \{\emptyset\}$  and  $I_b = \emptyset$ , the automaton  $\mathcal{B}_\varphi$  (with  $n$  states and  $c$  colors) is an alternating word automaton over the single-letter alphabet, whose emptiness can be checked in  $n^{O(c)}$  time. Checking (R0) is therefore in EXPTIME for CTL and  $\mu$ -calculus specifications and in 2EXPTIME for CTL\* specifications, respectively, in  $|\bigwedge_{b \in B} \varphi_b \rightarrow \psi|$ .

## 5 Conclusions

In open synthesis, where we synthesize a system that consists of a single process, it is safe to assume that the environment behavior is *maximal*. For the synthesis of a black-box process in the architecture of a general distributed system, the environment model needs two extensions: (1) The other black-box processes add a *reactive* component to the environment and (2) the process only has *incomplete information* about the environment behavior.

Extension (1) turns out to be expensive. Adding the reactive component increases the complexity for CTL specifications from EXPTIME [8] to 2EXPTIME [7], and for CTL\* specifications from 2EXPTIME [8] to 3EXPTIME [7]. As shown in Section 4, extension (2) has no extra cost. This settles an open question of [7]: The complexity of synthesizing a single process in a distributed architecture is still 2EXPTIME and 3EXPTIME, respectively.

The complexity of single-process synthesis is especially convincing in comparison to the cost of distributed synthesis: in the rare cases where distributed synthesis is decidable, the cost of synthesizing a distributed system with  $n$  processes (with distinguishable degree of information about the system state) is  $n$ -exponential in the size of the specification [1, 2].

Dividing the synthesis problem into several synthesis problems for single processes therefore appears as a promising approach to cope with the complexity and general undecidability of distributed synthesis. The situation is similar to the *verification* of distributed systems, where the compositional approach is well-established [12]. Our proof rule in Section 3 is a first example of a compositional synthesis technique. The rule is complete and therefore sufficient to decompose any realizable specification. The rule may, however, be less convenient to use than some compositional verification rules that, for example, apply circular assume-guarantee reasoning [13]. Defining such rules for the synthesis problem is an interesting topic of future research.

## References

1. Kupferman, O., Vardi, M.Y.: Synthesizing distributed systems. In: IEEE Symposium on Logic in Computer Science. (2001)
2. Finkbeiner, B., Schewe, S.: Uniform distributed synthesis. In: IEEE Symposium on Logic in Computer Science. (2005)
3. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Proc. IBM Workshop on Logics of Programs. Volume 131 of LNCS., Springer-Verlag (1981) 52–71
4. Wolper, P.: Synthesis of Communicating Processes from Temporal-Logic Specifications. PhD thesis, Stanford University (1982)
5. Kupferman, O., Vardi, M.Y.: Synthesis with incomplete informatio. In: Proc. 2nd International Conference on Temporal Logic (ICTL'97). (1997)
6. Kupferman, O., Vardi, M.Y.:  $\mu$ -calculus synthesis. In: Proc. 25th International Symposium on Mathematical Foundations of Computer Science. Volume 1893 of LNCS., Springer-Verlag (2000) 497–507
7. Kupferman, O., Madhusudan, P., Thiagarajan, P., Vardi, M.Y.: Open systems in reactive environments: Control and synthesis. In: Proc. 11th Int. Conf. on Concurrency Theory. Volume 1877 of LNCS., Springer-Verlag (2000) 92–107
8. Kupferman, O., Vardi, M.Y.: Church's problem revisited. *The bulletin of Symbolic Logic* **5** (1999) 245–263
9. Muller, D.E., Schupp, P.E.: Alternating automata on infinite trees. *Theor. Comput. Sci.* **54** (1987) 267–276
10. Muller, D.E., Schupp, P.E.: Simulating alternating tree automata by nondeterministic automata: new results and new proofs of the theorems of rabin, mcnaughton and safra. *Theor. Comput. Sci.* **141** (1995) 69–107

11. Jurdziński, M.: Small progress measures for solving parity games. In: 17th Annual Symposium on Theoretical Aspects of Computer Science. Volume 1770 of LNCS., Springer-Verlag (2000) 290–301
12. de Roever, W.P., Langmaack, H., Pnueli, A., eds.: Compositionality: The Significant Difference. COMPOS'97. Volume 1536 of LNCS., Springer Verlag (1998)
13. Maier, P.: A Lattice-Theoretic Framework For Circular Assume-Guarantee Reasoning. PhD thesis, Universität des Saarlandes, Saarbrücken (2003)