

# What You Really Need To Know About Your Neighbor\*

Werner Damm  
CvO Universität Oldenburg,  
26111 Oldenburg  
werner.damm@offis.de

Bernd Finkbeiner  
Universität des Saarlandes,  
Campus E1 3, 66123 Saarbrücken  
finkbeiner@cs.uni-saarland.de

Astrid Rakow  
CvO Universität Oldenburg,  
26111 Oldenburg  
a.rakow@uni-oldenburg.de

A fundamental question in system design is to decide how much of the design of one component must be known in order to successfully design another component of the system. We study this question in the setting of reactive synthesis, where one constructs a system implementation from a specification given in temporal logic. In previous work, we have shown that the system can be constructed compositionally, one component at a time, if the specification admits a "dominant" (as explained in Introduction) strategy for each component. In this paper, we generalize the approach to settings where dominant strategies only exist under certain assumptions about the future behavior of the other components. We present an incremental synthesis method based on the automatic construction of such assumptions.

## 1 Introduction

We investigate the following fundamental research question in system design: Consider a system architecture  $A$  composed of two processes  $p_h, p_l$  of different priorities high and low, all with defined interfaces  $inp(A), out(A), inp(p_h), out(p_h), inp(p_l), out(p_l)$  and objectives  $\varphi_A, \varphi_{p_h}, \varphi_{p_l}$  expressed in linear time temporal logic using the interfaces of  $A, p_h, p_l$ . How much does  $p_h$  need to know about  $p_l$  so that

- (1) both jointly realize  $\varphi_A$  whenever possible,
- (2)  $p_h$  can realize its objective whenever possible,
- (3)  $p_l$  only sacrifices achieving its objectives when this is the only way to achieve (1) and (2)?

We provide a rigorous formulation of this research question, and give algorithms to compute the insight  $p_h$  must have into  $p_l$ . We then show how this can be used to generate best-in-class strategies for  $p_h$  and  $p_l$ , in the sense that if there is a strategy at all for  $p_h$  which in the context of  $p_l$  achieves both (1) and (2), then so will the best-in-class strategy for  $p_h$ , and if there is a strategy at all for  $p_l$  which achieves (3), then so will the best-in-class strategy for  $p_l$ . The algorithm works for any number of processes with priorities henceforth represented by a partial order  $<$ .

Such system design questions arise naturally in many application areas, with  $A$  being a system of systems,  $\varphi_A$  representing overarching control objectives, and constituent systems  $p_1, \dots, p_n$  with their own local control objectives, where  $<$  reflects the degree of criticality of the subsystem for the overall system. In such systems, the notion of neighborhood has its classical interpretation based on physical proximity. One example of such systems are smart grids, where subsystems represent different types of energy providers and consumers, and the highest priority subsystem is responsible for maintaining stability of the subgrid in spite of large fluctuations of inflow and outflow of energy. The running example we use comes from cooperative autonomous vehicle control, where the overall objectives relate to safety,

---

\*This work has been partly supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" (SFB/TR 14 AVACS).

avoiding congestion, fuel and  $CO_2$  reduction etc, and local objectives relate to reaching destinations in given time frames, reducing fuel consumption, and executing certain maneuvers. We focus here on simple objectives such as completing a maneuver with cooperation from neighboring cars.

As pointed out in [10, 11], for systems of systems controlling physical systems as in the two examples above, there will be no winning strategies, because rare physical events such as stemming from extreme weather conditions or physical system failures can occur, making it impossible to even achieve overarching safety objectives. No electronic stability system of a car would be able to still guarantee safety in unexpected icy road conditions in a dynamical situation with already stretched system limits. We introduced the notion of (*remorse-free*) *dominant strategies* to compare strategies in such applications where a winning strategy cannot be found, making precise what has been called "best-in-class" strategies above. Such a strategy will – for every sequence of environment actions – be at least as good in achieving the system objectives as any other strategy. We can now make precise the addressed research question: how much insight does  $p_h$  need into the plans of  $p_l$ , so that there is a dominant strategy for  $p_h$  to achieve (1) and (2) above?

In general, local strategies must cater for needs of neighboring systems, so as to enable synthesis of (*remorse-free*) dominant systems. Consider e.g. a two car system on the highway, with the ego car wanting to reach the next exit, and having an other, neighboring car on the rightmost lane. There is no remorse-free dominant strategy of the ego car to reach the exit, unless the other car to its right cooperates: if the ego car accelerates to overtake the car, the car on right might do the same; similarly, an attempt to reach the exit by slowing down might be blocked by the car to the right. Thus ego needs insight into the future moves of the other car. In fact, if the other car would signal, that its next move will be to decelerate, then the ego car can bet on accelerating so as to overtake the car and reach the exit. If, on the other hand, the other car would promise to accelerate, then ego could decelerate and thus change lane in order to reach the exit. The weakest assumption for ego to have a dominant strategy will not only consider the next move of the other car, but give the other car as much flexibility as possible, as long as a lane change will still be possible prior to reaching the exit. We formalize these as what we call *assumption trees* of ego, which branches at each point in time into the set of possible futures of settings of the speed control of the other car s.t. the ego car has a dominant strategy in all such futures, and show that we can compute the most general assumption tree under which ego has a dominant strategy. A strategy for the other car is then not only determining the speed control of other, but also picking only evolutions which are compatible with ego's assumption tree. Rather than communicating the choice of such future moves through an extended interface, we provide ego with a copy of the strategy of other, so that ego can determine the selected future by simulating the strategy of the other car.

In general, such assumption trees can have unbounded depth. Consider the slight variant of the above example, where ego only wants to change lanes (no exit targeted) and has as secondary objective to always keep its speed. The most general assumption tree of ego will then, at each depth, have choices of futures where the other car will always keep its speed, forcing ego to either accelerate or decelerate to at least meet its higher priority objective of changing lane, and those futures where eventually other either breaks or accelerates, in which case ego can achieve both its objectives.

In general, we assume to have  $n$  processes ordered by criticality  $<$ . Intuitively, if  $p_h$  is more critical than  $p_l$  ( $p_h > p_l$ ), then we would want  $p_l$  to adapt its behavior so as to be compliant to the assumption tree of  $p_h$ . For a totally ordered set of processes  $p_1 > p_2 > \dots > p_n$  we thus generate assumption trees inductively starting with  $p_1$ , and then propagate this along the total order. Assuming the overall environment of  $A = \{p_1, \dots, p_n\}$  meets the assumption tree of  $p_n$ , we can then synthesize for each  $p_j$  a remorse-free dominant strategy meeting the canonical generalization of (1),(2),(3) above to  $n$  processes. If processes are only partially ordered, we perform this algorithm for clusters of processes with same

priority.

**Related Work** We introduced the notion of remorse-free dominant strategies in [10], and provided methods for compositional synthesis of remorse-free dominant strategies in [11]. Compositional approaches have been studied in the setting of assume-guarantee synthesis [4], where the synthesized strategies are guaranteed to be robust with respect to changes in other components, as long as the other components do not violate their own, local, specifications. In contrast, in this paper, we derive assumptions to ensure cooperation. The automatic synthesis of strategies for reactive systems goes back to the seminal works by Church, Büchi, Landweber, and Rabin in the 1960s [6, 3, 27], and by Pnueli and Rosner in the 1980s [25]. For distributed systems, the synthesis problem is known to be undecidable in general [26], and decidable under certain assumptions such as well-connected system architectures [19]. The problem that specifications usually must be strengthened with environment assumptions to become realizable has been recognized by several authors. The construction by Chatterjee et al [5] directly operates on the game graph. A safety assumption is computed by removing a minimal set of environment edges from the graph. A liveness assumption is then computed by putting fairness assumptions on the remaining edges. Liu et al [21] mine counterexamples to realizability to find new assumptions. A key difference between the assumptions generated in these approaches and the assumptions of this paper is that we compute assumptions for the existence of dominant strategies, not for the existence of winning strategies. As a result, our assumptions are weaker. In particular, it is in general not necessary to restrict the behavioral traces of the other components, only their branching structure.

The notion of admissible strategy in [1, 15, 16, 2] is closely related to our notion of dominant strategy. In [1] Berwanger presents a method based on iterated admissibility and gives existence results for general multi-player games of infinite duration. Faella presents in [16] an efficient way to compute admissible strategies by computing the conventional winning strategy, which is applied from winning states, and a cooperatively winning strategy for the remaining states. Brenguier et al [2] introduce an approach called assume-admissible synthesis. In assume-admissible synthesis it is assumed a-priori that the other components also apply admissible strategies, while we make no such assumption.

To use formal methods and in particular formal synthesis methods for coordinated vehicle maneuvers has been proposed in among others [18, 17, 31, 12, 13, 14, 30, 24]. [18] searches for strategies controlling all vehicles, and employs heuristic methods from artificial intelligence such as tree-search to determine strategies for coordinated vehicle movements. An excellent survey for alternative methods for controlling all vehicles to perform collision free driving tasks is given in [17]. Both methods share the restriction of the analysis to a small number of vehicles, in contrast to our approach, which is based on safe abstractions guaranteeing collision freedom in achieving the driver's objectives taking into account the complete traffic situation. In [7], the authors show that LTL formulas are expressive enough to express typical traffic-flow optimization objectives, and use formal synthesis methods to automatically generate control commands for access control on highway systems. In [31], the authors describe a methodological approach for decomposing the synthesis problem, which, however, has already been described in a previous journal publication by the authors [12]. It additionally presents a tool for constructing the finite abstraction; our previous work in [13, 14] goes beyond the capabilities of this tool in also allowing for non-linear dynamics. In [30], the authors provide an improved two-level approach for control synthesis that however suffers from a flaw in constructing the abstraction relation. In [24] the authors provide robust finite abstractions with bounded estimation errors for reducing the synthesis of winning strategies for LTL objectives to finite state synthesis and demonstrate the approach for an aerospace control application; however, this approach does not cover cooperative maneuvers. We finally

mention that in the aerospace domain, high-level objectives such as maximizing throughput and energy efficiency while maintaining safety have led to new concepts for air-traffic control, such as free flight. There is a significant body of work on the verification of the safety of such control strategies as well as on design rules that ensure safety [29, 22, 23, 28, 9, 8].

**Outline** Our paper is structured as follows. We introduce preliminaries on strategy synthesis and tree automata in Section 2 and 3. Section 5 shows how to compute assumptions for cooperation. Section 6 provides the incremental distributed synthesis procedure. As a running example to illustrate our notions, we use the *side by side* example, where a car A starts side by side with car B and A’s goal is to eventually be before or after B. We conclude in Sect. 7.

## 2 Synthesis of Distributed Systems

We consider complex multi-component systems at an early design stage, where the system architecture and the design objectives are known, but the components have not been implemented yet. We are interested in synthesizing an implementation for a given system architecture  $A$  and a specification  $\varphi$ . A solution to the synthesis problem is a set of finite-state strategies  $\{s_p \mid p \in P\}$ , one for each process in the architecture, such that the joint behavior satisfies  $\varphi$ .

**Architectures** An *architecture*  $A$  is a tuple  $(P, V, inp, out)$ , where  $P$  is a set of system processes,  $V$  is a set of (Boolean) variables, and  $inp, out : P \rightarrow 2^V$  are two functions that map each process to a set of input and output variables, respectively. For each process  $p$ , the inputs and outputs are disjoint,  $inp(p) \cap out(p) = \emptyset$ , and for two different processes  $p \neq q$ , the output variables are disjoint:  $out(p) \cap out(q) = \emptyset$ . We denote the set of visible variables of process  $p$  with  $V(p) = inp(p) \cup out(p)$ . Variables  $V_I = V \setminus \bigcup_{p \in P} out(p)$  that are not the output of any system process are called *external inputs*. We assume that the external inputs are available to every process, i.e.,  $V_I \subseteq inp(p)$  for every  $p \in P$ .

For two architectures  $A_1 = (P_1, V, inp_1, out_1)$  and  $A_2 = (P_2, V, inp_2, out_2)$  with the same variables, but disjoint sets of processes,  $P_1 \cap P_2 = \emptyset$ , we define the parallel composition as the architecture  $A_1 \parallel A_2 = (P_1 \cup P_2, V, p \mapsto \text{if } p \in P_1 \text{ then } inp_1(p) \text{ else } inp_2(p), p \mapsto \text{if } p \in P_1 \text{ then } out_1(p) \text{ else } out_2(p))$ .

*Example:* We use as running example the scenario from the introduction, where a car is side by side to another car and has to reach an exit. To be usable as running example, we strip this down to the bare essentials. Our formal model will consider this two car system as an architecture consisting of two “processes”, which we call ego and other. We restrict ourselves to finite state systems, and abstract the physical notion of positions of the two cars on adjacent lanes of a highway into the states *side\_by\_side*, and *was\_aside*, which are interpreted from the perspective of the ego car. Each car can control its own dynamics by accelerating, keeping the current speed, or decelerating. These actuators are captured in our formal model as output variables. E.g. for the process ego its output variables  $out(\text{ego})$  are  $a_e, k_e, d_e$  with the obvious meaning. We consider fully informed processes, hence ego has  $out(\text{other})$  as input variable.

We capture the effect of the setting of output variables and input variables of a process on the system state in what we call *world models*. Figure 1 shows the world model of the ego car. Initially, both cars are assumed to be side by side. State transitions are labeled by pairs of the chosen output of the process and the observed input of the process. E.g. if ego chooses to keep its speed, and the other car keeps its speed, too, then the system remains in state *side\_by\_side*. If, on the other hand, ego chooses to accelerate, while other keeps its velocity, the system transitions to state *was\_aside*. Multiple labels are just shorthand for

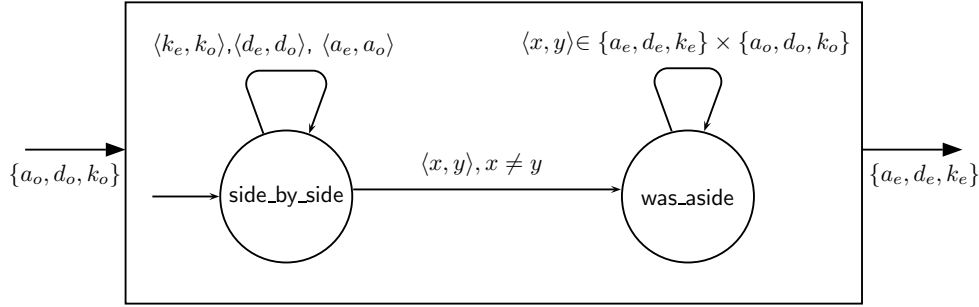


Figure 1: World model

multiple transitions with unique labels. Note that the world model of other is identical to that of ego except for exchanging inputs and outputs. From a control-theoretic perspective, the relative position of cars is a plant variable, which is influenced by the actuators of both cars.  $\square$

We now give the formal definition of the world model of a process  $p$ .

Let  $\mathbf{S}$  be a set of (system-) states. A world model  $M$  for process  $p$  with input variables  $inp(p)$  and output variables  $out(p)$  is an input deterministic transition system over  $\mathbf{S}$  with designated initial state and transitions labeled by pairs  $\langle o, i \rangle$  where  $i \in inp(p)$  and  $o \in out(p)$ . Formally  $M(p) = (\mathbf{S}, E, L, s_0)$  with  $E \subseteq \mathbf{S} \times \mathbf{S}$ ,  $L : E \rightarrow out(p) \times inp(p)$ ,  $s_0 \in \mathbf{S}$  s.t.  $(s, s_1) \in E \wedge (s, s_2) \in E \wedge L(s, s_1) = L(s, s_2) \Rightarrow s_1 = s_2$

**Implementations** An *implementation* of an architecture consists of strategies  $S = \{s_p \mid p \in P\}$  for the system processes. A system process  $p \in P$  is implemented by a *strategy*, i.e., a function  $s_p : (2^{inp(p)})^* \rightarrow 2^{out(p)}$  that maps histories of inputs to outputs. A strategy is *finite-state* if it can be represented by a finite-state *transducer*  $(Q, q_0, \delta : Q \times 2^{inp(p)} \rightarrow Q, \gamma : Q \rightarrow 2^{out(p)})$ , with a finite set of states  $Q$ , an initial state  $q_0$ , a transition function  $\delta$  and an output function  $\gamma$ .

The parallel composition  $s_p || s_q$  of the strategies of two processes  $p, q \in P$  is a function  $s_p || s_q : (2^I)^* \rightarrow 2^O$  that maps histories of the remaining inputs  $I = (inp(p) \cup inp(q)) \setminus (out(p) \cup out(q))$  to the union  $O = out(p) \cup out(q)$  of the outputs:  $s_p || s_q(\sigma) = s_p(\alpha_p(\sigma)) \cup s_q(\alpha_q(\sigma))$ , where  $\alpha_p : (2^I)^* \rightarrow (2^{inp(p)})^*$  fills in for process  $p$  what  $q$  contributes to the  $p$ 's input, that is  $\alpha_p(\varepsilon) = \varepsilon$  and  $\alpha_p(v_0 v_1 \dots v_k) = ((v_0 \cup s_q(\varepsilon)) \cap inp(p))((v_1 \cup s_q(\alpha_q(v_0))) \cap inp(p)) \dots ((v_k \cup s_q(\alpha_q(v_1 v_2 \dots v_{k-1}))) \cap inp(p))$ , and, analogously,  $\alpha_q(\varepsilon) = \varepsilon$  and  $\alpha_q(v_0 v_1 \dots v_k) = ((v_0 \cup s_p(\varepsilon)) \cap inp(q))((v_1 \cup s_p(\alpha_p(v_0))) \cap inp(q)) \dots ((v_k \cup s_p(\alpha_p(v_1 v_2 \dots v_{k-1}))) \cap inp(q))$ .

A *computation* is an infinite sequence of variable valuations. For a sequence  $\gamma = v_1 v_2 \dots \in (2^{V \setminus out(p)})^\omega$  of valuations of the variables outside the control of a process  $p$ , the computation resulting from  $s$  is denoted by  $comp(s, \gamma) = (s(\varepsilon) \cup v_1)(s(v_1 \cap inp(p)) \cup v_2)(s(v_1 \cap inp(p) v_2 \cap inp(p)) \cup v_3) \dots$

**Specification** We use linear-time temporal logic (LTL) to specify properties. In the following, we will denote the next-time operator with  $\bigcirc$ , globally with  $\square$  and eventually with  $\diamond$ . For a computation  $\sigma$  and an  $\omega$ -regular language  $\varphi$ , we write  $\sigma \models \varphi$  if  $\sigma$  satisfies  $\varphi$ . Design objectives are often given as a *prioritized* conjunction of LTL formulas  $\varphi = \varphi_1 \wedge \varphi_2 \wedge \dots \wedge \varphi_n$ , where  $\varphi_1$  is the most important objective and  $\varphi_n$  is the least important objective. For a priority  $k$ , with  $1 \leq k \leq n$ , we consider the partial conjunction  $\varphi^k = \bigwedge_{i=1 \dots k} \varphi_i$  and say that  $\sigma$  satisfies  $\varphi$  *up to priority*  $k$  if  $\sigma \models \varphi^k$ .

*Example:* In the side by side example, the control objective of the ego car is to reach a state where it can change lane, i.e., to eventually drive the plant into a state different from `side_by_side`. From the perspective of ego, the actuator settings of the other car are uncontrollable disturbances. Suppose that the

ego car has the secondary objective to reduce fuel consumption, which very abstractedly can be captured by avoiding accelerating and decelerating altogether, in other words to always output  $k_e$ . Suppose finally that other is owned by a driver who may avoid getting tired by changing speed every now and then. We capture these overall control objectives as list of prioritized LTL formulas:

1.  $\diamond \neg \text{side\_by\_side}$
2.  $\square k_e$
3.  $\square k_o \vee \square \diamond a_o \wedge \square \diamond d_o$ .

The synthesis problem, we then are solving, is to find strategies for ego and other which ultimately allow ego to change lane and avoid violating the lower priority objectives.  $\square$

A strategy  $s_p : (2^I)^* \rightarrow 2^O$  is *winning* for a property  $\varphi$ , denoted by  $s_p \models \varphi$ , iff, for every sequence  $\gamma = v_1 v_2 \dots \in (2^{V \setminus O})^\omega$  of valuations of the variables outside the control of  $p$ , the computation  $\text{comp}(s_p, \gamma)$  resulting from  $s_p$  satisfies  $\varphi$ . We generalize the notion of winning from strategies to implementations (and, analogously, the notions of dominance and bounded dominance later in the paper), by defining that an implementation  $S$  is winning for  $\varphi$  iff the parallel composition of the strategies in  $S$  is winning (for their combined sets of inputs and outputs).

**Synthesis** A property  $\varphi$  is *realizable* in an architecture  $A$  iff there exists an implementation that is winning for  $\varphi$ .

### 3 Preliminaries: Automata over Infinite Words and Trees

We assume familiarity with automata over infinite words and trees. In the following, we only give a quick summary of the standard terminology, the reader is referred to [20] for a full exposition.

A *full tree* is given as the set  $\Upsilon^*$  of all finite words over a given set of directions  $\Upsilon$ . For given finite sets  $\Sigma$  and  $\Upsilon$ , a  $\Sigma$ -*labeled  $\Upsilon$ -tree* is a pair  $\langle \Upsilon^*, l \rangle$  with a labeling function  $l : \Upsilon^* \rightarrow \Sigma$  that maps every node of  $\Upsilon^*$  to a letter of  $\Sigma$ .

An *alternating tree automaton*  $\mathcal{A} = (\Sigma, \Upsilon, Q, q_0, \delta, \alpha)$  runs on full  $\Sigma$ -labeled  $\Upsilon$ -trees.  $Q$  is a finite set of states,  $q_0 \in Q$  a designated initial state,  $\delta$  a transition function  $\delta : Q \times \Sigma \rightarrow \mathbb{B}^+(Q \times \Upsilon)$ , where  $\mathbb{B}^+(Q \times \Upsilon)$  denotes the positive Boolean combinations of  $Q \times \Upsilon$ , and  $\alpha$  is an acceptance condition. Intuitively, disjunctions in the transition function represent nondeterministic choice; conjunctions start an additional branch in the run tree of the automaton, corresponding to an additional check that must be passed by the input tree. A *run tree* on a given  $\Sigma$ -labeled  $\Upsilon$ -tree  $\langle \Upsilon^*, l \rangle$  is a  $Q \times \Upsilon^*$ -labeled tree where the root is labeled with  $(q_0, l(\varepsilon))$  and where for a node  $n$  with a label  $(q, x)$  and a set of children  $\text{child}(n)$ , the labels of these children have the following properties:

- for all  $m \in \text{child}(n)$ : the label of  $m$  is  $(q_m, x \cdot v_m)$ ,  $q_m \in Q$ ,  $v_m \in \Upsilon$  such that  $(q_m, v_m)$  is an atom of  $\delta(q, l(x))$ , and
- the set of atoms defined by the children of  $n$  satisfies  $\delta(q, l(x))$ .

A run tree is *accepting* if all its paths fulfill the acceptance condition. A *parity condition* is a function  $\alpha$  from  $Q$  to a finite set of colors  $C \subset \mathbb{N}$ . A path is accepted if the highest color appearing infinitely often is even. The *safety condition* is the special case of the parity condition where all states are colored with 0. The *Büchi condition* is the special case of the parity condition where all states are colored with either 1 or 2, the *co-Büchi condition* is the special case of the parity condition where all states are colored with either 0 or 1. For Büchi and co-Büchi automata we usually state the coloring function in terms of a set  $F$  of states. For the Büchi condition,  $F$  contains all states with color 2 and is called the set of *accepting*

states. For the co-Büchi condition,  $F$  contains all states with color 1 and is called the set of *rejecting* states. The Büchi condition is satisfied if some accepting state occurs infinitely often, the co-Büchi condition is satisfied if all rejecting states only occur finitely often. A  $\Sigma$ -labeled  $\Upsilon$ -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*  $\mathcal{L}(\mathcal{A})$ . An automaton is empty iff its language is empty. In addition to full trees, we also consider *partial trees*, which are given as prefix-closed subsets of  $\Upsilon^*$ . As partial trees can easily be embedded in full trees (for example, using a labeling  $\Upsilon^* \mapsto \mathbb{B}$  that indicates whether a node is present in the partial tree), tree automata can also be used to represent sets of partial trees.

A *nondeterministic* automaton is an alternating automaton where the image of  $\delta$  consists only of such formulas that, when rewritten in disjunctive normal form, contain at most one element of  $Q \times \{v\}$  for every direction  $v$  in every disjunct. A *universal* automaton is an alternating automaton where the image of  $\delta$  contains no disjunctions. A *deterministic* automaton is an alternating automaton that is both universal and nondeterministic, i.e., the image of  $\delta$  has no disjunctions and contains at most one element of  $Q \times \{v\}$  for every direction  $v$ .

A *word automaton* is the special case of a tree automaton where the set  $\Upsilon$  of directions is singleton. For word automata, we omit the direction in the transition function.

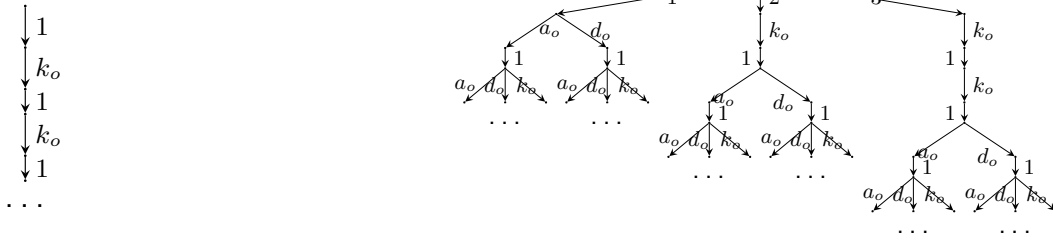
## 4 Dominant Strategies

In previous work, we introduced the notion of *remorse-free dominance* [10] in order to deal with situations where it is impossible to achieve the specified objective. Dominance is a weaker version of winning. A strategy  $t : (2^I)^* \rightarrow 2^O$  is *dominated* by a strategy  $s : (2^I)^* \rightarrow 2^O$ , denoted by  $t \preceq s$ , iff, for every sequence  $\gamma \in (2^{V \setminus O})^\omega$  for which the computation  $\text{comp}(t, \gamma)$  resulting from  $t$  satisfies the objective specification up to priority  $m$ , the computation  $\text{comp}(s, \gamma)$  resulting from  $s$  satisfies the objective specification up to priority  $n$  and  $m \leq n$ . A strategy  $s$  is *dominant* iff, for all strategies  $t$ ,  $t \preceq s$ . Analogously to the definition of winning implementations, we say that an implementation  $S$  is dominant iff the parallel composition of the strategies in  $S$  is dominant.

Finally, we say that a property  $\varphi$  is *admissible* in an architecture  $A$  iff there is a dominant implementation. Informally, a specification is admissible if the question whether it can be satisfied does not depend on variables that are not visible to the process or on *future inputs*. For example, the specification  $\varphi = (\bigcirc a) \leftrightarrow b$ , where  $a$  is an input variable and  $b$  is an output variable is not admissible, because in order to know whether it is best to set  $b$  in the first step, one needs to know the value of  $a$  in the second step. No matter whether the strategy sets  $b$  or not, there is an input sequence that causes *remorse*, because  $\varphi$  is violated for the chosen strategy while it would have been satisfied for the same sequence of inputs if the other strategy had been chosen.

**Theorem 1.** [11] *For a process  $p$ , one can construct a parity tree automaton such that the trees accepted by the automaton define exactly the dominant strategies of  $p$ .*

*Example:* In the side by side example, the specification is not admissible: Consider  $\text{KEEP}^\omega$ , a strategy where ego always keeps its speed, i.e., applies  $k_e^\omega$ . Following  $\text{KEEP}^\omega$ , ego will achieve objectives 1 and 2 on p. 26, if other accelerates or decelerates at least once. But in case the other car will always keep its speed, following  $\text{KEEP}^\omega$  ego will not even achieve its most important objective,  $\diamond \neg \text{side\_by\_side}$ . In this case though, following strategy ACC, which is to accelerate at the first step and then keep speed, ego will achieve objective 1. So, neither ACC nor  $\text{KEEP}^\omega$  is dominant. Ego will feel remorse, if it follows strategy ACC but other eventually chooses  $a_o$  or  $d_o$ , and ego will also feel remorse, if it follows  $\text{KEEP}^\omega$



a) ego assumes, other does  $\text{KEEP}^\omega$       b) ego assumes, other changes speed with the first three steps

Figure 2: Two assumption trees of ego

while other also always keeps its speed. In order to know whether to choose  $\text{KEEP}^\omega$  or  $\text{ACC}$ , ego needs to know whether other will eventually change its speed or always keep speed.  $\square$

## 5 Computing cooperation assumptions

A component does not necessarily have a dominant strategy, even if the composite system has a dominant strategy. Suppose that, in the example specification  $\varphi = (\bigcirc a) \leftrightarrow b$  discussed above, the architecture consists of two processes  $p$  and  $q$ , with outputs  $b$  and  $a$ , respectively. Clearly, the system has a dominant, and even winning strategy, e.g., always set  $a$  and  $b$  to *false*, while  $p$ , as discussed above, does not have a dominant strategy. An interesting observation is that  $p$  does have a dominant strategy if it is guaranteed that  $q$  will set  $a$  to *false* in the next step, and that it also has a dominant strategy if it is guaranteed that  $q$  will set  $a$  to *true* (in which case,  $p$  would set  $b$  to *true* as well).

We formalize such assumptions on the environment behavior that guarantee the existence of a dominant strategy as *assumption trees*. An *assumption tree* for process  $p$  is a subtree of the full  $\mathbb{N} \cup 2^{V \setminus \text{out}(p)}$ -tree where, on every path, the direction alternates between  $\mathbb{N}$  and  $2^{V \setminus \text{out}(p)}$ . Intuitively, each number corresponds to a promise of the environment to a particular restriction on its future behavior. The values of  $2^{V \setminus \text{out}(p)}$  are then the execution of this promise. The future behavior following different numbers can differ in the inputs provided by the environment at specific points in time in the future and also in the number of branchings with natural numbers that occur in the future.

*Example:* In our side by side example, a very restrictive assumption is "other always chooses  $k_o$ ". This assumption can be envisioned as a tree that degenerates to a simple linear graph as in Fig. 2 a). At the root the single edge announces that only one future is assumed initially. From there the single edge labeled  $k_o$  encodes that the only allowed first step of this future is  $k_o$ . Likewise, the single edge at node  $1 k_o$  announces that only one future development is distinguished at that point. The single edge at  $1 k_o 1$  encodes that the next step has to be  $k_o$  again, and so on. With respect to this assumption, ego cannot achieve  $\diamond \neg \text{side\_by\_side} \wedge \square k_e$ , but ego can achieve objective 1 without remorse following any strategy that chooses at least once  $a_e$  or  $d_e$ . Any such strategy is dominant wrt. the assumption tree in Fig. 2 a).

Ego may also assume that "other eventually chooses  $\neg k_o$ ", where  $\neg k_o$  abbreviates  $a_o \vee d_o$ . The strategy  $\text{KEEP}^\omega$  is the only dominant strategy with respect to this assumption and any stronger assumption – including the one depicted in Fig. 2 b). The tree of Fig. 2 b) roughly describes ego's assumption that other changes its speed within the first three steps. At the tree's root, three assumptions on other's behavior are distinguished. The first, at node 1, describes that other chooses initially  $\neg k_o$ . The subtrees at  $1 a_o$  and  $1 d_o$ , respectively, both describe the assumption that all future development is possible: The



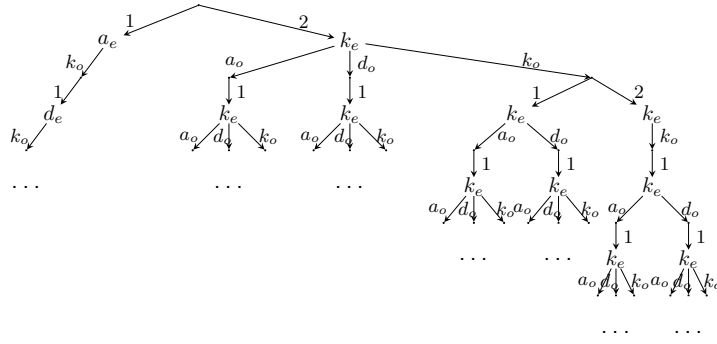


Figure 3: Side\_by\_side example: A strategy annotated assumption tree of ego.

outgoing single edge at node 1  $a_o$  announces that then only a single future is distinguished. The three outgoing edges at 1  $a_o$  1 specify that this future assumption allows to apply any acceleration, i.e.,  $a_o$ ,  $d_o$  or  $k_o$ , as second step. From there again, but not depicted, a single edge leads to a node with three outgoing edges labeled  $a_o$ ,  $d_o$  and  $k_o$  denoting that any choice for other’s acceleration is accepted for this step as well; and so on. The subtree at node 2 in Fig. 2 b) describes that other initially chooses  $k_o$ , then  $\neg k_o$  and thereafter the future is unconstrained. The subtree at node 3 in Fig. 2 b) describes that other is assumed to choose  $k_o k_o$  followed by  $\neg k_o$  and then the future is no further constrained.

For an example of an assumption tree with infinite branching at the root, consider ego’s assumption that other announces in subtree  $i$  to do  $\neg k_o$  at step  $i \in \mathbb{N}$ .  $\square$

To ensure that a given assumption tree guarantees the existence of a dominant strategy, we annotate each node that is reached by a natural number with an output value from  $2^{out(p)}$ . One such strategy annotation  $t$  is *dominated* by another strategy annotation  $s$  of the same assumption tree, denoted by  $t \preceq s$ , iff for every alternating sequence of natural numbers and environment outputs from  $2^{out(p)}$  that is present in the tree, if the computation resulting from  $t$  satisfies the objective specification up to priority  $m$ , then the computation resulting from  $s$  satisfies the objective specification up to priority  $n$  and  $m \leq n$ . A strategy annotation  $s$  is *dominant* iff, for all strategies  $t$ ,  $t \preceq s$ . We say that an assumption tree is *dominant* (and it *guarantees the existence of a dominant strategy*) iff it has a dominant strategy annotation.

*Example:* In Fig. 2 b) where the assumption tree basically expresses that other guarantees to choose  $\neg k_o$  initially or within the next two steps, we can annotate the assumption tree with strategy  $KEEP^\omega$  to get a dominant strategy annotation.

An example of a strategy annotated assumption tree is shown in Fig. 3. There we basically combined the assumptions “other chooses always  $k_o$ ” and “other guarantees to choose  $\neg k_o$  initially or within its next two steps” as subtrees 1 and 2, respectively. The strategy annotation results from applying the strategy that alternates  $a_e$  and  $d_e$ ,  $(a_e d_e)^\omega$  on subtree 1, and the  $KEEP^\omega$  strategy on subtree 2.  $\square$

We now describe the computation of a tree automaton that recognizes the set of all assumption trees for a given process with a dominant annotation. A small technical difficulty is that assumption trees have an infinite branching degree. In order to represent the set of assumption trees as a tree automaton, we use the standard unary representation of  $\mathbb{N}^*$  in  $\mathbb{B}^*$ , where a sequence  $abc \dots \in \mathbb{N}^*$  is encoded as  $0^a 10^b 10^c 1 \dots$ . The construction then proceeds in the following steps. (1) We build a word automaton  $\mathcal{A}$  that checks that no alternative would do better on the present path. (2) We build a tree automaton  $\mathcal{B}$  that reads in a strategy annotated assumption tree and uses  $\mathcal{A}$  to verify that this annotation is dominant.

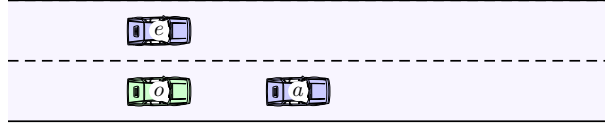


Figure 4: Initial scenario for the incremental synthesis: Three cars on a freeway.

**Theorem 2.** *For a single process  $p$ , one can construct a parity tree automaton such that the trees accepted by the automaton are the dominant strategy annotated assumption trees of  $p$ .*

## 6 Incremental synthesis of cooperation strategies

We now use the construction of the environment assumptions from Section 5 to incrementally synthesize a distributed system by propagating the assumptions. We assume that the processes  $P = \{p_1, p_2, \dots, p_n\}$  are ordered  $p_1 > p_2 > \dots > p_n$  according to criticality. We begin by computing the set  $A_1$  of strategy annotated assumption trees for process  $p_1$ . Next, we compute the set  $A_2$  of strategy annotated assumption trees for process  $p_2$  that additionally satisfy  $A_1$ . We continue up to process  $p_n$ .

**Theorem 3.** *For an architecture with a set  $P = \{p_1, p_2, \dots, p_n\}$  of processes, ordered  $p_1 > p_2 > \dots > p_n$  according to criticality, one can construct, for each process  $p_i, i = 1..n$ , a parity tree automaton, such that the trees accepted by the automaton are assumption trees that guarantee the existence of dominant strategies for all  $p_j$  with  $j \leq i$ .*

The set of assumption trees constructed in this way represent the remaining assumptions on the external inputs. Since we cannot constrain the external environment, we find the final distributed implementation by selecting some assumption tree that allows for all possible future environment behaviors.

*Example:* Let us consider three cars as in Fig. 4. Ego and other are initially side by side and the car ahead starts in front of other. The synthesis problem is to find a strategy for ego, other and ahead that satisfies the prioritized specification given by the following objectives listed in order of decreasing priority.

1.  $\diamond \neg \text{side\_by\_side}_e \wedge \square (\text{curve} \Rightarrow \bigcirc \bigcirc \neg a_a)$
2.  $\square k_e$
3.  $\square (\bigcirc d_a \Rightarrow d_o) \wedge \square (\bigcirc k_a \Rightarrow \neg a_o)$

The objectives for ego remain the same. ahead is required not to accelerate in a curve, i.e., when ahead notices a curve two steps ahead, it may not accelerate two steps later, while in the curve. The car other has to ensure not to get faster than ahead. This is here expressed via a sufficient constraint on its acceleration.

We order processes  $\text{ego} > \text{other} > \text{ahead}$  and start the incremental synthesis with ego. We have already seen elements of the set of ego's assumption trees for the simpler setting of two cars in Sect. 5. An assumption tree for the three car example will additionally be labeled with acceleration values for ahead and the environmental input announcing bends  $\text{curve\_ahead} \in \{\text{curve}, \text{no\_curve}\}$ . But since ego depends on neither of them, they are of no relevance for the assumption tree, i.e., for each assumption tree in the two car setting there will be now an assumption tree where ahead and  $\text{curve\_ahead}$  are unconstrained.

In the following we will illustrate the assumption propagation, starting with the exemplary assumption tree of ego depicted in Fig. 3. Next, we consider a strategy annotated assumption tree of other that matches with ego's assumption tree of Fig. 3.. For the sake of a simple illustration, we have chosen a strong assumption of other about ahead, that is annotated with a simple strategy for other: The strategy

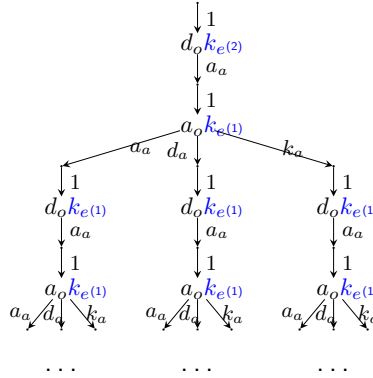


Figure 5: A strategy annotated assumption tree of other: ahead is assumed to apply  $a_a$  at least every other time. Other’s dominant strategy is to alternate between  $d_o$  and  $a_o$ . The tree satisfies  $\text{saat}(\text{ego})$  of Fig. 3.

annotated assumption tree,  $\text{saat}(\text{other})$ , given in Fig. 5, describes the assumption that ahead chooses  $a_a$  at least every other time. The strategy  $(d_o a_o)^\omega$  for other is dominant strategy wrt. the assumption tree in Fig. 5.

Now, the  $\text{saat}(\text{other})$  satisfies ego’s assumption “other guarantees to choose  $\neg k_o$  within its next two steps”. A compatible dominant strategy for ego is determined based on the propagated strategy annotated assumption tree. The match with ego’s strategy is highlighted in blue in Fig. 5. The nodes are also annotated with ego’s acceleration choices. The numbers in brackets refer to the respective edge labels in the assumption tree in Fig. 3.

Intuitively, the weakest assumption for other is “ahead announces when it will do  $d_a$  or  $k_a$  next” and a dominant strategy for other is, to then choose the acceleration appropriately.

At this point the inputs curve ahead and the acceleration of ahead are open. Since the exemplary strategies do not refer to these, we again omitted these variables at the edge annotation in Fig. 5.

Next and final element according to our process order is ahead. So we choose a dominant strategy for ahead, for instance a strategy that chooses  $d_a$  in two steps time after the environment announces a curve two steps ahead, otherwise it chooses to accelerate. The strategies for ego and ahead are now determined based on the propagated (ego,other)-strategy annotated assumption trees, as illustrated in the previous step. Intuitively, when ahead gets informed that a bend is two steps ahead, other knows from the propagated strategy assumption tree that ahead will decelerate in two steps time and chooses to decelerate at its the next step.  $\square$

## 7 Conclusion

Relying on cyber-physical systems increases our vulnerability towards technical failures caused by un-predicted emergent behavior. Understanding what pieces of information must be exchanged with whom, understanding what level of cooperation is necessary for the overall objectives –topics addressed in this paper– constitute initial steps towards a rigorous design discipline for such systems. They also expose the vulnerability of the system towards intruders: the analysis clearly can be used to highlight the chaos, which can be inserted into the system by pretending to pass information and then not acting accordingly.

The paper thus both provides an initial step towards cooperatively establishing systems objectives, as well as exposing its vulnerabilities. We will explore both avenues in our future research in the application domain of cooperative driver assistance systems.

## References

- [1] Dietmar Berwanger (2007): *Admissibility in Infinite Games*. In Wolfgang Thomas & Pascal Weil, editors: *STACS 2007: 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007. Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 188–199, doi:10.1007/978-3-540-70918-3\_17.
- [2] Romain Brenguier, Jean-François Raskin & Ocan Sankur (2015): *Assume-Admissible Synthesis*. In: *26th International Conference on Concurrency Theory (CONCUR 2015)*, LIPIcs 42, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp. 100–113, doi:10.4230/LIPIcs.CONCUR.2015.100. Available at <http://arxiv.org/abs/1507.00623>.
- [3] J. Richard Büchi & Lawrence H. Landweber (1990): *Solving Sequential Conditions by Finite-State Strategies*. In Saunders Mac Lane & Dirk Siefkes, editors: *The Collected Works of J. Richard Büchi*, Springer New York, pp. 525–541, doi:10.1007/978-1-4613-8928-6\_29.
- [4] Krishnendu Chatterjee & Thomas A. Henzinger (2007): *Assume-Guarantee Synthesis*. In Tiziana Margaria & Wang Yi, editors: *Tools and Algorithms for the Construction and Analysis of Systems, LNCS 4424*, Springer Berlin Heidelberg, pp. 261–275, doi:10.1007/978-3-540-71209-1\_21.
- [5] Krishnendu Chatterjee, Thomas A. Henzinger & Barbara Jobstmann (2008): *Environment Assumptions for Synthesis*. In: *CONCUR 2008 - Concurrency Theory: 19th International Conference, CONCUR 2008, Toronto, Canada, August 19-22, 2008. Proceedings, Lecture Notes in Computer Science 5201*, Springer Berlin Heidelberg, pp. 147–161, doi:10.1007/978-3-540-85361-9\_14.
- [6] Alonzo Church (1963): *Logic, Arithmetic and Automata*. In: *Proc. 1962 Intl. Congr. Math.*, pp. 23–25, doi:10.2307/2270398.
- [7] Samuel Coogan & M. Arcak (2014): *Freeway traffic control from linear temporal logic specifications*. In: *Cyber-Physical Systems (ICCPs), 2014 ACM/IEEE International Conference on*, pp. 36–47, doi:10.1109/ICCPs.2014.6843709.
- [8] Werner Damm, Stefan Disch, Hardi Hungar, Swen Jacobs, Jun Pang, Florian Pigorsch, Christoph Scholl, Uwe Waldmann & Boris Wirtz (2007): *Exact State Set Representations in the Verification of Linear Hybrid Systems with Large Discrete State Space*. In Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino & Yoshio Okamura, editors: *Automated Technology for Verification and Analysis, 5th International Symposium, ATVA 2007, Lecture Notes in Computer Science 4762*, Springer, pp. 425–440, doi:10.1007/978-3-540-75596-8\_30.
- [9] Werner Damm, Stefan Disch, Hardi Hungar, Jun Pang, Florian Pigorsch, Christoph Scholl, Uwe Waldmann & Boris Wirtz (2006): *Automatic verification of hybrid systems with large discrete state space*. In Susanne Graf & Wenhui Zhang, editors: *Proceedings of the 4th International Symposium on Automated Technology for Verification and Analysis, Lecture Notes in Computer Science 4218*, Springer-Verlag, pp. 276–291, doi:10.1007/11901914\_22.
- [10] Werner Damm & Bernd Finkbeiner (2011): *Does It Pay to Extend the Perimeter of a World Model?* In Michael Butler & Wolfram Schulte, editors: *FM 2011: Formal Methods - 17th International Symposium on Formal Methods, Limerick, Ireland, June 20-24, 2011. Proceedings, Lecture Notes in Computer Science 6664*, Springer Berlin Heidelberg, pp. 12–26, doi:10.1007/978-3-642-21437-0\_4.
- [11] Werner Damm & Bernd Finkbeiner (2014): *Automatic Compositional Synthesis of Distributed Systems*. In Cliff Jones, Pekka Pihlajasaari & Jun Sun, editors: *FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings, Lecture Notes in Computer Science 8442*, Springer International Publishing, pp. 179–193, doi:10.1007/978-3-319-06410-9\_13.

- [12] Werner Damm, Hans-Jörg Peter, Jan Rakow & Bernd Westphal (2013): *Can we build it: formal synthesis of control strategies for cooperative driver assistance systems*. *Mathematical Structures in Computer Science* 23, pp. 676–725, doi:10.1017/S0960129512000230.
- [13] Werner Damm, Guilherme Pinto & Stefan Ratschan (2007): *Guaranteed Termination in the Verification of LTL Properties of Non-linear Robust Discrete Time Hybrid Systems*. *Int. Journal of Foundations of Computer Science* 18(1), pp. 63–86, doi:10.1142/S0129054107004577.
- [14] Tomáš Dzetkulič & Stefan Ratschan (2011): *Incremental Computation of Succinct Abstractions for Hybrid Systems*. In Uli Fahrenberg & Stavros Tripakis, editors: *Formal Modeling and Analysis of Timed Systems, Lecture Notes in Computer Science* 6919, Springer Berlin Heidelberg, pp. 271–285, doi:10.1007/978-3-642-24310-3\_19.
- [15] Marco Faella (2007): *Games you cannot win*. Available at <http://wpage.unina.it/m.faella/Download/games07.pdf>. Workshop on Games and Automata for Synthesis and Validation, Lausanne, Switzerland.
- [16] Marco Faella (2009): *Admissible Strategies in Infinite Games over Graphs*. In Rastislav Kráľovič & Damian Niwiński, editors: *Mathematical Foundations of Computer Science 2009: 34th International Symposium, MFCS 2009, Novy Smokovec, High Tatras, Slovakia, August 24-28, 2009. Proceedings*, Springer Berlin Heidelberg, pp. 307–318, doi:10.1007/978-3-642-03816-7\_27.
- [17] Christian Frese (2010): *A Comparison of Algorithms for Planning Cooperative Motions of Cognitive Automobiles*. Technical Report IES-2010-06, Lehrstuhl für Interaktive Echtzeitsysteme, doi:10.5445/KSP/1000021181.
- [18] Christian Frese & Jürgen Beyerer (2010): *Planning Cooperative Motions of Cognitive Automobiles Using Tree Search Algorithms*. In Rüdiger Dillmann, Jürgen Beyerer, Uwe D. Hanebeck & Tanja Schultz, editors: *KI 2010: Advances in Artificial Intelligence, Lecture Notes in Computer Science* 6359, Springer Berlin Heidelberg, pp. 91–98, doi:10.1007/978-3-642-16111-7\_10.
- [19] Paul Gastin, Nathalie Sznajder & Marc Zeitoun (2006): *Distributed Synthesis for Well-Connected Architectures*. In S. Arun-Kumar & Naveen Garg, editors: *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science* 4337, Springer Berlin Heidelberg, pp. 321–332, doi:10.1007/11944836\_30.
- [20] Erich Grädel, Wolfgang Thomas & Thomas Wilke, editors (2002): *Automata, Logics, and Infinite Games*. LNCS 2500, Springer-Verlag, doi:10.1007/3-540-36387-4.
- [21] Lingyi Liu, David Sheridan, Viraj Athavale & Shobha Vasudevan (2011): *Automatic generation of assertions from system level design using data mining*. In: *Formal Methods and Models for Codesign (MEMOCODE), 2011 9th IEEE/ACM International Conference on*, pp. 191–200, doi:10.1109/MEMCOD.2011.5970526.
- [22] John Lygeros & Nancy Lynch (1998): *Strings of vehicles: Modeling and safety conditions*. In Thomas A. Henzinger & Shankar Sastry, editors: *Hybrid Systems: Computation and Control, Lecture Notes in Computer Science* 1386, Springer Berlin Heidelberg, pp. 273–288, doi:10.1007/3-540-64358-3\_45.
- [23] John Lygeros, Claire Tomlin & Shankar Sastry (1997): *Multiobjective hybrid controller synthesis*. In Oded Maler, editor: *Hybrid and Real-Time Systems, Lecture Notes in Computer Science* 1201, Springer Berlin Heidelberg, pp. 109–123, doi:10.1007/BFb0014720.
- [24] Oscar Mickelin, Necmiye Ozay & Richard M. Murray (2014): *Synthesis of Correct-by-construction Control Protocols for Hybrid Systems Using Partial State Information*. In: *2014 American Control Conference*, pp. 2305–2311, doi:10.1109/ACC.2014.6859229.
- [25] Amir Pnueli & Roni Rosner (1989): *On the Synthesis of a Reactive Module*. In: *Proceedings of the 16th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '89*, ACM, pp. 179–190, doi:10.1145/75277.75293.
- [26] Amir Pnueli & Roni Rosner (1990): *Distributed Reactive Systems Are Hard to Synthesize*. In: *Proceedings of the 31st Annual Symposium on Foundations of Computer Science, SFCS '90 2*, IEEE Computer Society, pp. 746–757, doi:10.1109/FSCS.1990.89597.

- [27] M.O. Rabin (1972): *Automata on Infinite Objects and Church's Problem*. Conference Board of the mathematical science, Conference Board of the Mathematical Sciences, doi:10.1090/cbms/013.
- [28] Arthur Richards & Jonathan P. How (2002): *Aircraft trajectory planning with collision avoidance using mixed integer linear programming*. In: *American Control Conference, 2002. Proceedings of the 2002*, 3, pp. 1936–1941, doi:10.1109/ACC.2002.1023918.
- [29] Claire J. Tomlin, George J. Pappas, Jana Košecká, John Lygeros & Shankar S. Sastry (1998): *Advanced air traffic automation: A case study in distributed decentralized control*. In Bruno Siciliano & Kimon P. Valavanis, editors: *Control Problems in Robotics and Automation*, Springer-Verlag, pp. 261–295, doi:10.1007/BFb0015088.
- [30] Tichakorn Wongpiromsarn, Ufuk Topcu & Richard M. Murray (2012): *Receding Horizon Temporal Logic Planning*. *Automatic Control, IEEE Transactions on* 57(11), pp. 2817–2830, doi:10.1109/TAC.2012.2195811.
- [31] Tichakorn Wongpiromsarn, Ufuk Topcu & Richard M. Murray (2013): *Synthesis of Control Protocols for Autonomous Systems*. *Unmanned Systems* 01(01), pp. 21–39, doi:10.1142/S2301385013500027.