

Cooperative Reactive Synthesis*

Roderick Bloem

IAIK, Graz University of
Technology, Graz, Austria

Rüdiger Ehlers

University of Bremen and
DFKI GmbH, Bremen, Germany

Robert Könighofer

IAIK, Graz University of
Technology, Graz, Austria

A modern approach to engineering correct-by-construction systems is to synthesize them automatically from formal specifications. Oftentimes, a system can only satisfy its guarantees if certain environment assumptions hold, which motivates their inclusion in the system specification. Experience with modern synthesis approaches shows that synthesized systems tend to satisfy their specifications by actively working towards the violation of the assumptions rather than satisfying assumptions and guarantees together. Such uncooperative behavior is undesirable because it violates the aim of synthesis: the system should try to satisfy its guarantees and use the assumptions only when needed. Also, the assumptions often describe the valid behavior of other components in a bigger system, which should not be obstructed unnecessarily.

In this paper, we present a hierarchy of *cooperation levels* between system and environment. Each level describes how well the system enforces both the assumptions and guarantees. We show how to synthesize systems that achieve the highest possible cooperation level for a given specification in Linear Temporal Logic (LTL). The synthesized systems can also exploit cooperative environment behavior during operation to reach a higher cooperation level that is not enforceable by the system initially. The worst-case time complexity of our synthesis procedure is doubly-exponential, which matches the complexity of standard LTL synthesis. This is an extended version of [7] that features an additional appendix.

1 Introduction

When synthesizing reactive systems from their formal specifications, we typically start with a set of guarantees that the system should fulfill and a set of assumptions about the environment. A synthesis tool then computes an implementation that satisfies the guarantees in all environments that satisfy the assumptions. In many specifications, the system can influence whether the assumptions are satisfied; in particular, the system can actively force the environment to violate the assumptions. The resulting implementation is correct: it fulfills its guarantees if the assumptions are fulfilled. However, it does so in an undesirable way.

Take for example a flight stabilization system that adds small forces to choices of forces issued by the pilot. The system guarantees state that there is little jitter in the absolute forces applied. This can only work for “well-behaved” evolutions of the manually selected force values, which in turn depend on the forces currently applied. Without the well-behavedness assumption, the system has no way to stabilize the flight. However, if the system has the opportunity to add forces that make the pilot violate this assumption, it can do so without violating its overall specification (namely that the guarantees must hold whenever the assumptions hold). Clearly, such behavior is not covered by the specifier’s intent.

This observation leads to the question of how we can synthesize systems that cooperate with the environment whenever possible. Our main idea is that, at any point in time, the system should try to be

*This work was supported in part by the Austrian Science Fund (FWF) through the research network RiSE (S11406-N23), by the European Commission through the projects STANCE (317753) and IMMORTAL (644905), and by the Institutional Strategy of the University of Bremen, funded by the German Excellence Initiative.

as cooperative as possible while still ensuring correctness. Depending on the concrete situation, several levels of cooperation may be possible. Some examples are:

1. The system can enforce both the assumptions and guarantees to hold.
2. The system can enforce the guarantees if the assumptions hold and the system can give the environment the opportunity to satisfy the assumptions at the same time.
3. The system can neither enforce the assumptions nor the guarantees, but environment and system together can satisfy both.

The first of these levels is most beneficial: there is no need to rely on the environment, not even for satisfying the assumptions. On the second level we can satisfy the correctness objective even without allowing the system to enforce an assumption violation. From states of the third level, the system cannot enforce correctness. However, instead of resigning and behaving arbitrarily, the system still offers some executions along which both the assumptions and guarantees are fulfilled, thereby optimistically assuming that the environment is helpful.

In this paper, we perform a rigorous analysis of cooperation between the system and its environment in the setting of reactive synthesis. We generalize the three levels of cooperation from above to a fine-grained *cooperation hierarchy* that allows us to reason about how cooperative a controller for a given specification can be. A level in our hierarchy describes what objectives (in terms of assumptions and guarantees) the controller can achieve on its own and for what objectives it has to rely on the environment.

As a second contribution, we present a synthesis procedure to construct a controller that always picks the highest possible cooperation level in our hierarchy for a given linear-time temporal logic (LTL) specification. The synthesized controllers do not only enforce the highest possible level statically, but also exploit environment behavior that enables reaching a higher level that cannot be enforced initially. Thus, implementations synthesized with our approach are *maximally cooperative*, without any need to declare cooperation in the specification explicitly by enumerating scenarios in which a system can operate in a cooperative way. Our *maximally cooperative synthesis procedure* takes at most doubly-exponential time, which matches the complexity of standard LTL synthesis. Our techniques are applicable to all ω -regular word languages. For specifications given as deterministic Rabin word automata, the complexity is polynomial in the number of states and exponential in the number of acceptance condition pairs.

Outline. The presentation of this paper is structured as follows. Section 2 reviews related work and Section 3 introduces background and notation. Section 4 then presents our hierarchy of cooperation levels, while Section 5 describes our synthesis approach for obtaining cooperative systems. We conclude in Section 6.

2 Related Work

In earlier work [6], we already pointed out that existing synthesis approaches handle environment assumptions in an unsatisfactory way. We postulated that synthesized systems should (1) *be correct*, (2) *not be lazy* by satisfying guarantees even if assumptions are violated, (3) *never give up* by working towards the satisfaction of the guarantees even if this is not possible in the worst case, and (4) *cooperate* by helping the environment to satisfy the assumptions that we made about it [6]. Our current work addresses all of these challenges.

Besides correctness, our main focus is on cooperation, which is also addressed by Assume-Guarantee Synthesis [12] and synthesis under rationality assumptions [3, 10, 11, 18]. These approaches do not distinguish between system and environment, but synthesize implementations for both. Each component

works under certain assumptions about the other components not deviating from their synthesized implementations arbitrarily. In contrast, our work handles the system and its environment asymmetrically: we prefer the guarantees over the assumptions and prioritize correctness over supporting the environment.

Similar to existing work on synthesis of robust [4] and error-resilient [16] systems, our synthesis approach is also *not lazy* [6] in satisfying guarantees. The reason is that satisfying guarantees is more preferable in our cooperation hierarchy than satisfying guarantees only if the assumptions are satisfied. In contrast to the existing work, we do not only consider the worst case environment behavior for satisfying guarantees (even if assumptions are violated), but also the scenario where guarantees can only be satisfied with the help of the environment.

Like [17], we also address the *never give up* challenge [6] because our cooperation hierarchy also includes levels on which the system cannot enforce the guarantees any more. Still, our synthesis approach lets the system satisfy the guarantees for *some* environment behavior whenever that is possible.

In contrast to quantitative synthesis [1,5] where the synthesized system maximizes a certain pay-off, our approach is purely qualitative: a certain level of the hierarchy is either achieved or not. In this way, we not only avoid the computational blow-up induced by operating with numerical data, but also remove the need to assign meaningful quantities (such as probabilities) to the specifications.

Finally, [15] presents an extension of a synthesis algorithm for so-called GR(1) specifications [8] to produce mission plans for robots that always offer some execution on which both assumptions and guarantees are satisfied. Hence, [15] gives a solution for one level in our hierarchy, implemented for a subset of LTL.

3 Preliminaries

We denote the Boolean domain by $\mathbb{B} = \{\text{false}, \text{true}\}$ and the set of natural numbers (including 0) by \mathbb{N} .

Words: We consider synthesis of reactive systems with a finite set $I = \{i_1, \dots, i_m\}$ of Boolean input signals and a finite set $O = \{o_1, \dots, o_n\}$ of Boolean outputs. The input alphabet is $\mathcal{I} = 2^I$, the output alphabet is $\mathcal{O} = 2^O$, and $\Sigma = \mathcal{I} \times \mathcal{O}$. The set of finite (infinite) words over Σ is denoted by Σ^* (Σ^ω), where ε is the word of length 0. A set $L \subseteq \Sigma^\omega$ of infinite words is called a (*word*) *language*.

Specifications: A specification φ over Σ defines a language $L(\varphi)$ of allowed words. In this paper, specifications consist of two parts, the *environment assumptions* \mathcal{A} and the *system guarantees* \mathcal{G} . These parts can be combined by logical operators. For example, we write $\mathcal{A} \rightarrow \mathcal{G}$ to specify that the guarantees (only) need to hold if the assumptions are satisfied. We say that some finite word $w \in \Sigma^*$ is a *bad prefix* for φ if there does not exist a word $w' \in \Sigma^\omega$ such that $ww' \in L(\varphi)$. The set of infinite words that have no bad prefixes of φ will be called the *safety hull* of φ .

Reactive systems: A reactive system interacts with its environment in a synchronous way. In every time step j , the system first provides an output letter $y_j \in \mathcal{O}$, after which the environment responds with an input letter $x_j \in \mathcal{I}$. This is repeated in an infinite execution to produce the *trace* $w = (x_0, y_0)(x_1, y_1) \dots \in \Sigma^\omega$.

We can represent the complete behavior of a reactive system by a *computation tree* $\langle T, \tau \rangle$, where T is the set of nodes and a subset of \mathcal{I}^* , and τ assigns labels to the tree nodes. In such a tree, we have $\tau : T \rightarrow \mathcal{O}$, i.e., the tree nodes are labeled by the last output of the system. We call trees $\langle T, \tau \rangle$ with $T = \mathcal{I}^*$ *full trees* and consider only these henceforth, unless otherwise stated. We say that some trace $w = (x_0, y_0)(x_1, y_1) \dots \in \Sigma^\omega$ is *included* in $\langle T, \tau \rangle$ if for all $i \in \mathbb{N}$, $y_i = \tau(x_0 \dots x_{i-1})$. In such a case, we also say that w is a *branch* of $\langle T, \tau \rangle$.

We say that some specification φ is *realizable* if there exists a computation tree for the system all of whose traces are in $L(\varphi)$. The set of traces of $\langle T, \tau \rangle$ is denoted by $L(\langle T, \tau \rangle)$ and called the *word language* of $\langle T, \tau \rangle$. The set of all computation trees over \mathcal{I} and \mathcal{O} is denoted by \mathcal{T} . Since a computation tree represents the full behavior of a reactive system, we use these two terms interchangeably.

Rabin word and tree automata: In order to represent specifications over Σ in a finitary way, we employ *Rabin automata*. For word languages, we use deterministic Rabin word automata, which are defined as tuples $\mathcal{R} = (Q, \Sigma, \delta, q_0, \mathcal{F})$, where Q is the finite set of *states*, Σ is the finite alphabet, $\delta : Q \times \Sigma \rightarrow Q$ is the *transition function*, $q_0 \in Q$ is the initial state of \mathcal{R} , and $\mathcal{F} \subseteq 2^Q \times 2^Q$ is the *acceptance condition* of \mathcal{R} . Given a word $w \in \Sigma^\omega$, \mathcal{R} induces a *run* $\pi = \pi_0 \pi_1 \pi_2 \dots \in Q^\omega$, where $\pi_0 = q_0$ and for every $i \in \mathbb{N}$, we have $\pi_{i+1} = \delta(\pi_i, w_i)$. The run π is *accepting* if there exists some *acceptance condition pair* $(F, G) \in \mathcal{F}$ such that $\text{inf}(\pi) \cap F = \emptyset$, and $\text{inf}(\pi) \cap G \neq \emptyset$, where inf maps a sequence to the set of elements that appear infinitely often in the sequence. We say that w is accepted by \mathcal{R} if there exists a run for w that is accepting. The language of \mathcal{R} , denoted as $L(\mathcal{R})$, is defined to be the set of words accepted by \mathcal{R} . The language of a state $q \in Q$ is defined to be the language of the automaton $\mathcal{R}_q = (Q, \Sigma, \delta, q, \mathcal{F})$, which differs from \mathcal{R} only by its initial state. Without loss of generality, we assume that every Rabin automaton has a designated state \top that has the full language, i.e., for which $L(\mathcal{R}_\top) = \Sigma^\omega$.

In addition to deterministic Rabin word automata, we will later also need non-deterministic Rabin tree automata. A non-deterministic Rabin tree automaton $\mathcal{R} = (Q, \mathcal{I}, \mathcal{O}, \delta, q_0, \mathcal{F})$ is defined similarly to a word automaton, except that δ has a different structure, and we now have two alphabets, namely a *branching alphabet* \mathcal{I} and a *label alphabet* \mathcal{O} . The transition function δ is defined as $\delta : Q \times \mathcal{O} \rightarrow 2^{\mathcal{I} \rightarrow Q}$. A tree automaton accepts or rejects computation trees instead of words. Given a computation tree $\langle T, \tau \rangle$, we say that some tree $\langle T', \tau' \rangle$ is a *run tree* of \mathcal{R} for $\langle T, \tau \rangle$ if $T' = T$, $\tau' : T' \rightarrow Q$, $\tau'(\varepsilon) = q_0$, and for all $t' \in T'$, there exists a function $f \in \delta(\tau'(t'), \tau(t'))$ such that for all $i \in \mathcal{I}$, $f(i) = \tau'(t'i)$. We say that $\langle T', \tau' \rangle$ is *accepting* if every branch of $\langle T', \tau' \rangle$ is accepting, i.e., its sequence of labellings satisfies the Rabin acceptance condition \mathcal{F} . We say that a tree automaton accepts a computation tree if it has a corresponding accepting run tree.

Given a Rabin word automaton \mathcal{R} as specification over the alphabet $\Sigma = \mathcal{I} \times \mathcal{O}$, we can check if there exists a tree with branching alphabet \mathcal{I} and label alphabet \mathcal{O} all of whose traces are in the language of \mathcal{R} . It has been shown that this operation can be performed in time exponential in $|\mathcal{F}|$ and polynomial in $|Q|$ by first translating the Rabin word automaton to a non-deterministic Rabin tree automaton with the same set of states and the same acceptance condition (in linear time), and then checking the tree automaton's language for emptiness [19]. This approach gives rise to a reactive synthesis procedure for specifications in *linear temporal logic* (LTL) [19]: we can first translate the LTL specification to a deterministic Rabin word automaton with a number of states that is doubly-exponential in the length of the specification, and a number of acceptance condition pairs that is exponential in the length of the specification. Overall, this gives a doubly-exponential time procedure for LTL reactive synthesis, which matches the known complexity of the problem.

4 A Hierarchy of Cooperation Levels

In this section, we develop a hierarchy of cooperation levels that a system may achieve. We use a special logic to specify desired properties of the system to be synthesized. The syntax of a formula Φ in this logic is defined as

$$\Phi ::= \varphi \mid \langle E \rangle \varphi \mid G \langle E \rangle \varphi \mid \Phi \wedge \Phi, \quad (1)$$

where φ is a linear-time specification (e.g., an LTL formula) over an alphabet Σ . The Boolean connective \wedge has its expected semantics. The formula $\langle E \rangle \varphi$ is satisfied for a system $\langle T, \tau \rangle$ if there exists some trace of $\langle T, \tau \rangle$ on which φ holds. Similarly, the formula $G\langle E \rangle \varphi$ is satisfied if, from any point in a reactive system's trace that has been seen so far, there exists some suffix trace of the system such that φ holds. We call an instance of the grammar in Formula 1 a *cooperation level specification*. Our logic ranges over computation trees and has similarities to *strategy logic* [13] as well as *alternating-time temporal logic* (ATL) [2]. Yet, its semantics, to be given below, is very different. In particular, linear-time specifications φ are seen as atomic and even within the scope of a G operator, φ is only evaluated on *complete traces* of a system, always starting at the root of a computation tree.

More formally, we define the semantics as follows. Formulas are interpreted over a computation tree $\langle T, \tau \rangle \in \mathcal{S}$ using the following rules:

$$\begin{array}{lll}
\langle T, \tau \rangle & \models \varphi & \text{iff } \forall w \in L(\langle T, \tau \rangle) : w \in L(\varphi) \\
\langle T, \tau \rangle & \models \langle E \rangle \varphi & \text{iff } \exists w \in L(\langle T, \tau \rangle) : w \in L(\varphi) \\
\langle T, \tau \rangle & \models G\langle E \rangle \varphi & \text{iff } \forall w \in L(\langle T, \tau \rangle), i \in \mathbb{N}. \exists w' \in L(\langle T, \tau \rangle) \cap L(\varphi). \\
& & w_0 \dots w_{i-1} = w'_0 \dots w'_{i-1} \\
\langle T, \tau \rangle & \models \Phi_1 \wedge \Phi_2 & \text{iff } (\langle T, \tau \rangle \models \Phi_1) \wedge (\langle T, \tau \rangle \models \Phi_2)
\end{array}$$

Most cases are straightforward, but $G\langle E \rangle \varphi$ requires some attention. We follow an arbitrary trace of $\langle T, \tau \rangle$ up to step $i-1$. At step $i-1$, there must be some continuation of the trace that satisfies φ and that is part of the system's computation tree. If such a continuation exists in every step i , then $G\langle E \rangle \varphi$ is satisfied.

4.1 Defining the Interesting Levels of Cooperation

We start with the linear-time specification \mathcal{A} , which represents the assumptions about the environment, and \mathcal{G} , which represents the guarantees. We can combine them to the linear-time properties $\mathcal{A} \rightarrow \mathcal{G}$ and $\mathcal{A} \wedge \mathcal{G}$. The first of these represents the classical correctness requirement for reactive synthesis, while the latter represents the optimistic linear-time property that both assumptions and guarantees hold along a trace of the system. As generators for cooperation level specifications, we consider all rules from the grammar in Equation 1 except for the second one, as it is rather weak, and only considers what can happen from the initial state of a system onwards. Additionally, leaving out the conjuncts of the form $\langle E \rangle \varphi$ strengthens the semantic foundation for our *maximally cooperative synthesis* approach in Sect. 5.2. However, we discuss the consideration of conjuncts of the form $\langle E \rangle \varphi$ in Section 4.3.

Combining all four linear-time properties with the two chosen ways of lifting a linear-time property to our logic gives the following different conjuncts for cooperation level specifications:

$$D = \{ \mathcal{A} \rightarrow \mathcal{G}, \mathcal{G}, \mathcal{A}, G\langle E \rangle (\mathcal{A} \wedge \mathcal{G}), G\langle E \rangle \mathcal{G}, G\langle E \rangle \mathcal{A}, G\langle E \rangle (\mathcal{A} \rightarrow \mathcal{G}) \} \quad (2)$$

A *cooperation level specification* is a conjunction between elements from this set. So there are $2^7 = 128$ possible cooperation levels in this setting. Note that we removed the linear-time property $\mathcal{A} \wedge \mathcal{G}$ from D , as it can be simulated by a conjunction between \mathcal{A} and \mathcal{G} on the level of cooperation level specifications.

We can reduce the 128 possible cooperation levels substantially by using our knowledge of the semantics of cooperation level specifications, which can be expressed in *reduction rules*. For example, if in a cooperation level specification, for some linear-time property $\varphi \in \{ \mathcal{A}, \mathcal{G}, \mathcal{A} \rightarrow \mathcal{G}, \mathcal{A} \wedge \mathcal{G} \}$, φ is a conjunct along with $G\langle E \rangle \varphi$, then we can remove the latter from the cooperation level specification as

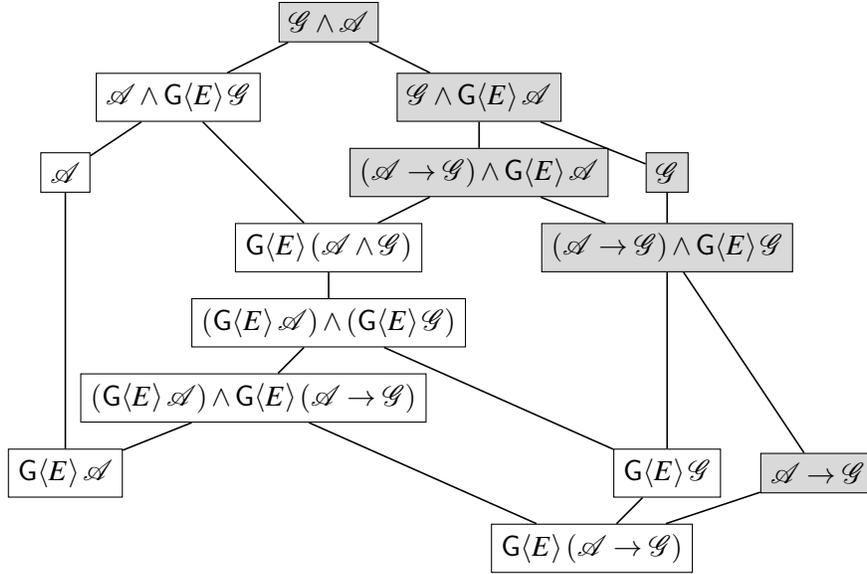


Figure 1: A Hasse diagram of the hierarchy of cooperation levels.

that conjunct is implied by the former. This is because if along every trace of a computation tree, φ holds, then for every node in the tree, we can find a trace containing the node along with φ holds. In a similar fashion, we can observe

- that \mathcal{G} implies $\mathcal{A} \rightarrow \mathcal{G}$,
- that $G\langle E \rangle (\mathcal{A} \wedge \mathcal{G})$ implies $G\langle E \rangle \mathcal{A}$ and $G\langle E \rangle \mathcal{G}$,
- that $G\langle E \rangle \mathcal{G}$ implies $G\langle E \rangle (\mathcal{A} \rightarrow \mathcal{G})$,
- that $\mathcal{A} \rightarrow \mathcal{G}$ and \mathcal{A} together imply \mathcal{G} , and
- that $\mathcal{A} \rightarrow \mathcal{G}$ and $G\langle E \rangle \mathcal{A}$ together imply $G\langle E \rangle (\mathcal{A} \wedge \mathcal{G})$,
- that \mathcal{A} and $G\langle E \rangle \mathcal{G}$ together imply $G\langle E \rangle (\mathcal{A} \wedge \mathcal{G})$,
- that $G\langle E \rangle (\mathcal{A} \rightarrow \mathcal{G})$ and \mathcal{A} together imply $G\langle E \rangle (\mathcal{G})$.

Only in the last three of these *rules*, conjuncts of the forms $G\langle E \rangle \varphi$ and φ interact. For example, if $\mathcal{A} \rightarrow \mathcal{G}$ holds along all traces of a computation tree, and we know that through every node in the tree, there is a trace on which \mathcal{A} holds, then along this trace, $\mathcal{A} \wedge \mathcal{G}$ holds as well. Thus, we also know that the computation tree satisfies $G\langle E \rangle (\mathcal{A} \wedge \mathcal{G})$. Note that $G\langle E \rangle (\mathcal{A} \wedge \mathcal{G})$ is not equal to $G\langle E \rangle (\mathcal{G}) \wedge G\langle E \rangle (\mathcal{A})$ because there exist computation trees for which a part of their traces satisfy $\mathcal{A} \wedge \neg \mathcal{G}$, the other traces satisfy $\mathcal{G} \wedge \neg \mathcal{A}$, but none of their traces satisfy $\mathcal{A} \wedge \mathcal{G}$.

After reducing the set of distinct cooperation levels by these reduction rules, we obtain only 15 semantically different subsets of D , which form a lattice with its partial order defined by implication. We leave out the true element of the lattice in the following (as it is trivially satisfied by all computation trees), and visualize the remaining cooperation levels hierarchically in Fig. 1. For every cooperation level, implied elements of D have also been left out in the vertex labeling. Every edge in the figure denotes a “stricter than”-relation, with the upper element being the stricter one. Vertices that enforce the traditional correctness criterion $\mathcal{A} \rightarrow \mathcal{G}$ are colored in gray.

In the next section, we discuss the different levels of the cooperation hierarchy in an example. After that, we will discuss design decisions that we made when constructing the hierarchy, as well as extensions and alternatives.

4.2 Example

This section presents an example from real life to illustrate some cooperation levels from Fig. 1. A more extensive example will be presented in Appendix A.

An IT provider makes a contract with a customer. The provider promises that whenever the computer breaks, it will eventually deliver a new computer, or one of its technicians will come to the customer to fix the computer while providing some free printing supplies as a bonus. The customer accepts the conditions, mainly because of the possibility of free printing supplies. The contract between the customer and the IT provider contains the following assumptions under which it has to be fulfilled: (1) every traffic blockade on the way to the customer is eventually cleared, and (2) if the road to the customer is clear and the IT provider technician is working on the computer, then it will eventually be fixed. The latter assumption is part of the contract to exclude problems with the customer’s software, which is not the responsibility of the IT provider. By simply replacing parts of the computer one-by-one, the problem is eventually fixed. We also have an additional assumption that traffic can be blocked on purpose by the IT provider, which is not written into the contract.

We model the setting with two Boolean input variables c, b and four Boolean output variables t, d, m, n for the IT service provider’s behavior. Variable c is true if the computer is currently working, b is true when the traffic is not blocked, $f = \text{true}$ indicates that the IT provider’s technician is trying to fix the computer, $m = \text{true}$ means that free printing supplies are delivered, $n = \text{true}$ means that the IT service provider delivers a new computer, and $d = \text{true}$ means that the provider currently blocks the traffic on purpose. In LTL syntax¹, the guarantee can be written as $\mathcal{G} = G(\neg c \rightarrow F(n \vee (c \wedge m)))$. The assumptions can be formalized as $\mathcal{A} = G(Fb) \wedge G((b \wedge f) \rightarrow Fc) \wedge G(d \rightarrow \neg b)$.

The following table summarizes some cooperation levels that can be achieved with different behavior of the IT provider, each expressed as an LTL property whose fulfillment completely determines the valuation of the output variables in all time steps. We focus on levels that enforce $\mathcal{A} \rightarrow \mathcal{G}$ (which are colored gray in Fig. 1).

Nr.	Behavior	Level in Fig. 1
1	$G(d \wedge \neg f \wedge \neg m \wedge \neg n)$	$\mathcal{A} \rightarrow \mathcal{G}$
2	$G(d \wedge m \wedge \neg f \wedge \neg n)$	$(\mathcal{A} \rightarrow \mathcal{G}) \wedge G\langle E \rangle \mathcal{G}$
3	$G(f \wedge m \wedge \neg n \wedge \neg d)$	$(\mathcal{A} \rightarrow \mathcal{G}) \wedge G\langle E \rangle \mathcal{A}$
4	$G(n \wedge d \wedge \neg f \wedge \neg m)$	\mathcal{G}
5	$G(n \wedge \neg d \wedge \neg f \wedge \neg m)$	$\mathcal{G} \wedge G\langle E \rangle \mathcal{A}$

Behavior 1 enforces an assumption violation by blocking the traffic. This is very uncooperative with the customers (and all other drivers on the streets). In this case, $G\langle E \rangle \mathcal{A}$ does not hold as the environment cannot satisfy \mathcal{A} from any point in any trace because it cannot set b to true at some point without violating \mathcal{A} , but not ever doing so violates \mathcal{A} as well. The cooperation level specification part $G\langle E \rangle \mathcal{G}$ does not hold either because since n and m are both false all of the time, \mathcal{G} is not fulfilled along any trace.

Behavior 2 is better because printing supplies are always delivered. This satisfies $G\langle E \rangle \mathcal{G}$ because, at any point, the environment can make the computer work again (which occasionally happens with computers). $G\langle E \rangle \mathcal{A}$ is still not satisfied for the same reason as before.

¹In LTL syntax, G means “always” and F means “eventually”.

Behavior 3 is even less destructive. The technician tries to fix the computer ($f = \text{true}$) and brings along free printing supplies ($m = \text{true}$) without blocking traffic. This satisfies the guarantees if the assumptions are satisfied. On top of that, the assumptions can always be satisfied by setting both b and c to true. Since d is always false, the last assumption (which was the problem before) always holds.

Behavior 4 is better than Behavior 2 but incomparable with Behavior 3. By always (or repeatedly) providing a new computer, \mathcal{G} is enforced independent of the assumptions. However, $G\langle E \rangle \mathcal{A}$ does not hold because, since the traffic is blocked ($d = \text{true}$), the assumptions cannot hold.

Behavior 5 is similar but without enforcing a traffic blockade. It thus satisfies \mathcal{G} and $G\langle E \rangle \mathcal{A}$ simultaneously. The last two of the four assumptions are even enforced. However, since the first assumption cannot be enforced by any behavior of the IT provider, $\mathcal{G} \wedge \mathcal{A}$ cannot be achieved.

4.3 Discussion

Before solving the *cooperative synthesis* problem in the next section, let us discuss some interesting aspects of our hierarchy of cooperation levels.

Incomparable levels: As already discussed in the example, our hierarchy of cooperation levels in Fig. 1 contains levels that are incomparable, i.e., where neither of the levels is stricter or clearly more desirable. The relative preferences between such levels may depend on the application. Among the levels that enforce $\mathcal{A} \rightarrow \mathcal{G}$ (gray in Fig. 1), there is only one incomparability, namely between \mathcal{G} and $(\mathcal{A} \rightarrow \mathcal{G}) \wedge G\langle E \rangle \mathcal{A}$. The former favors the guarantees, the latter the assumptions. For incomparabilities between a level that enforces $\mathcal{A} \rightarrow \mathcal{G}$ and one that does not, we suggest to prefer the former unless there are good reasons not to.

Symmetry in Fig. 1: Our hierarchy is asymmetric because we included $\mathcal{A} \rightarrow \mathcal{G}$ but not $\mathcal{G} \rightarrow \mathcal{A}$ in D . Including $\mathcal{A} \rightarrow \mathcal{G}$ contradicts the philosophy of cooperation to some extent, but is justified by the fact that we always take the point of view of the system in this paper, and try to synthesize a *correct implementation* that also helps the environment to satisfy its assumptions whenever this is reasonable, but without assuming that the environment has its own goals and behaves *rationally* (as in [3, 10–12, 18]). The guarantees often cannot be enforced unconditionally by the system, so the system has to rely on the assumptions to hold. However, the combination of $\mathcal{A} \rightarrow \mathcal{G}$ with $G\langle E \rangle \mathcal{A}$ (or $G\langle E \rangle \mathcal{G}$) in a cooperation level specification eliminates the possibility for the system to achieve correctness by simply enforcing a violation of \mathcal{A} (and \mathcal{G}).

Additional operators: We did not include any properties with a plain $\langle E \rangle$ -operator in our default hierarchy. The reason is that $\langle E \rangle \varphi$ is quite a weak goal. The system could have exactly one trace on which φ is satisfied. As soon as the environment deviates from the input along this trace, the system can behave arbitrarily. Our default hierarchy also does not contain any occurrences of $\mathcal{A} \vee \mathcal{G}$, but this is mainly to keep the presentation simple. Including $\mathcal{A} \vee \mathcal{G}$ and $G\langle E \rangle (\mathcal{A} \vee \mathcal{G})$ would extend the hierarchy to 23 levels. Additionally including $\langle E \rangle$ wherever $G\langle E \rangle$ is applied results in 77 levels. Thus, even with extensions in the applied operators, the size of our hierarchy remains manageable. Details to these refined hierarchies and the reduction rules to obtain them can be found in Appendix B.

Fine-grainedness: Our hierarchy considers two dimensions: the goals (in terms of assumptions and guarantees) to achieve, and the certainty with which the goal can be achieved: enforced, always reachable ($G\langle E \rangle$), and initially reachable ($\langle E \rangle$, if considered). In contrast, most existing synthesis approaches [4, 5, 16] (see Section 2) that go beyond plain correctness only focus on *enforcing* a maximum in the dimension of goals. Yet, in this dimension they are often more fine-grained, e.g., by considering individual assumptions and guarantees or how often some property is violated. In principle, we could also increase

the granularity in our goal dimension. However, this also comes at a price: it would increase the size of the hierarchy and induce more incomparabilities, which makes it more difficult to define the preference between the incomparable levels for a concrete application.

5 Synthesizing Desirable Systems

After defining our hierarchy of cooperation levels, we turn towards the synthesis of implementations that maximize the possible cooperation level. We start by describing how we can synthesize an implementation for a single cooperation level, and then show in Sect. 5.2 how to synthesize *maximally cooperative implementations*, which move upwards in the cooperation level hierarchy whenever possible during the execution.

5.1 Implementing a Single Cooperation Level

The simple reduction of the cooperative synthesis problem for LTL to synthesis from a logic such as *strategy logic* [13] is obstructed by the fact that in our semantics, we always evaluate traces from the start when evaluating a subformula of the shape $G\langle E \rangle \varphi$. Strategy logic lacks a *rewind* operator that would allow to jump back to the start of a trace. We could, however, encode a cooperation level specification into CTL* with linear past [9], and use a corresponding synthesis procedure.

Instead of following this route, we give a direct automata-theoretic approach to the synthesis of reactive systems that implement some level of cooperation with the environment on a specification $(\mathcal{A}, \mathcal{G})$. The automata built in this way keep information about the states of the assumptions and guarantees explicit, which is needed to synthesize maximally cooperative implementations in the next subsection. Also, it makes the synthesis approach applicable to general ω -regular word language specifications.

Starting from the linear-time properties \mathcal{A} and \mathcal{G} , we show how to build a non-deterministic Rabin tree automaton that encodes the synthesis problem for some cooperation level specification H . Such a tree automaton can be checked for language emptiness in order to perform synthesis for one level of cooperation.

Let $\mathcal{R}_{\mathcal{A}}$, $\mathcal{R}_{\mathcal{G}}$, $\mathcal{R}_{\mathcal{A} \rightarrow \mathcal{G}}$, and $\mathcal{R}_{\mathcal{A} \wedge \mathcal{G}}$ be deterministic Rabin word automata that encode \mathcal{A} , \mathcal{G} , $\mathcal{A} \rightarrow \mathcal{G}$, and $\mathcal{A} \wedge \mathcal{G}$. We translate the conjuncts (elements of D from Equation 2) of some expression H in the grammar from Equation 1 to non-deterministic Rabin tree automata \mathcal{T}^D individually and then build the product \mathcal{T}^H between these tree automata, which encodes that all elements of H have to be satisfied in a candidate computation tree. As cooperation level specifications have only three types of conjuncts, namely φ , $\langle E \rangle \varphi$, and $G\langle E \rangle \varphi$ for some linear-time property φ , we simply give the translations for these types separately.

Case φ : If φ is a linear-time property, the occurrence of φ in a cooperation level specification indicates that φ should hold on every trace of a synthesized implementation. Translating $\mathcal{R}_{\varphi} = (Q, \mathcal{I} \times \mathcal{O}, \delta, q_0, \mathcal{F})$ to a Rabin tree automaton $\mathcal{T}^D = (Q, \mathcal{I}, \mathcal{O}, \delta', q_0, \mathcal{F})$ that enforces φ to hold along all traces in the tree is a standard construction, where we set:

$$\forall q \in Q, o \in \mathcal{O} : \delta'(q, o) = \{\{i \mapsto \delta(q, (i, o)) \mid i \in \mathcal{I}\}\}$$

Case $\langle E \rangle \varphi$: Here, we require the synthesized system to offer one path along which φ holds. In the tree automaton, we non-deterministically choose this path. Starting with $\mathcal{R}_{\varphi} = (Q, \mathcal{I} \times \mathcal{O}, \delta, q_0, \mathcal{F})$, we

obtain $\mathcal{T}^D = (Q, \mathcal{I}, \mathcal{O}, \delta', q_0, \mathcal{F})$ with:

$$\forall q \in Q, o \in \mathcal{O} : \delta'(q, o) = \bigcup_{i \in \mathcal{I}} \{ \{i \mapsto \delta(q, (i, o))\} \cup \{i' \mapsto \top \mid i' \in \mathcal{I}, i' \neq i\} \}$$

Case $G\langle E \rangle \varphi$: In this case, for every node in a computation tree, we require the synthesized system to have a path on which φ holds that includes the selected node. Given $\mathcal{R}_\varphi = (Q, \mathcal{I} \times \mathcal{O}, \delta, q_0, \mathcal{F})$, we can implement this requirement as a non-deterministic Rabin tree automaton $\mathcal{T}^D = (Q', \mathcal{I}, \mathcal{O}, \delta', q'_0, \mathcal{F}')$ with:

$$\begin{aligned} Q' &= Q \times \mathbb{B} \\ \forall (q, b) \in Q', o \in \mathcal{O} : \delta'((q, b), o) &= \bigcup_{i \in \mathcal{I}} \{ \{i \mapsto (\delta(q, (i, o)), \text{true})\} \\ &\quad \cup \{i' \mapsto (\delta(q, (i', o)), \text{false}) \mid i' \in \mathcal{I}, i' \neq i\} \} \\ q'_0 &= (q_0, \text{true}) \\ \mathcal{F}' &= \{ (F \times \{\text{true}\}, G \times \{\text{true}\}) \mid (F, G) \in \mathcal{F} \} \\ &\quad \cup \{ (\emptyset, Q \times \{\text{false}\}) \} \end{aligned}$$

The automaton augments the states in Q by a Boolean flag. From every node in a computation tree accepted by \mathcal{T}^D , regardless of whether it is flagged by true or false, there must exist a branch consisting only of true-labeled nodes. The original acceptance condition \mathcal{F} must hold along this branch. However, not all branches of a tree accepted by \mathcal{T}^D have to satisfy φ , as those branches along which the flag is false infinitely often are trivially accepting. Intuitively, \mathcal{T}^D also enforces the safety hull of φ along all branches in the computation tree as once a bad prefix has been seen on the way to a computation tree node, there cannot exist a branch containing the node along which φ holds, which is enforced by the (required) successor branch that is always labeled by true.

In order to obtain a non-deterministic Rabin tree automaton \mathcal{T}^H for a complete cooperation level specification H , we can compute a product automaton from the Rabin tree automata \mathcal{T}^D for the individual conjuncts. Computing such a product automaton is a standard operation in automata theory. Given a set $\{(Q_1, \mathcal{I}, \mathcal{O}, \delta_1, q_{0,1}, \mathcal{F}_1), \dots, (Q_n, \mathcal{I}, \mathcal{O}, \delta_n, q_{0,n}, \mathcal{F}_n)\}$ of non-deterministic Rabin tree automata, their product is defined as the non-deterministic Rabin tree automaton $\mathcal{T}^H = (Q, \mathcal{I}, \mathcal{O}, \delta, q_0, \mathcal{F})$ with:

$$\begin{aligned} Q &= Q_1 \times \dots \times Q_n \\ \delta((q_1, \dots, q_n), o) &= \bigotimes_{i \in \{1, \dots, n\}} \delta_i(q_i, o) \\ q_0 &= (q_{0,1}, \dots, q_{0,n}) \\ \mathcal{F} &= \{ (F_1 \times \dots \times F_n, G_1 \times \dots \times G_n) \mid \\ &\quad (F_1, G_1) \in \mathcal{F}_1, \dots, (F_n, G_n) \in \mathcal{F}_n \} \end{aligned}$$

The second line of this equation holds for all $(q_1, \dots, q_n) \in Q$ and $o \in \mathcal{O}$. Also, we used the special operator \bigotimes that maps sets $\{M_i \subseteq 2^{\mathcal{Q}_i \rightarrow \mathcal{I}}\}_{i \in \{1, \dots, n\}}$ to the set

$$\{ \{ ((q_1, \dots, q_n), i) \in Q \times \mathcal{I} \mid \bigwedge_{j \in \{1, \dots, n\}} f_j(i) = q_j \} \mid f_1 \in M_1, \dots, f_n \in M_n \}$$

By performing reactive synthesis using the product Rabin tree automaton as specification, we can obtain an implementation that falls into the chosen cooperation level.

5.2 Maximally Cooperative Synthesis

For linear-time specifications \mathcal{A} and \mathcal{G} , there are typically some cooperation level specifications that are not realizable (such as $\mathcal{A} \wedge \mathcal{G}$) and some that are realizable (such as $G\langle E \rangle(\mathcal{A} \rightarrow \mathcal{G})$). By iterating over all cooperation level specifications and applying reactive synthesis for the product Rabin tree automata \mathcal{T}' computed by the construction from above, we can check which is the highest cooperation level that can be realized and compute a respective implementation.

In some cases, there may however be the possibility to switch to a higher level of cooperation during the system's execution. Take for example the case that initially, the cooperation level $G\langle E \rangle \mathcal{A} \wedge (\mathcal{A} \rightarrow \mathcal{G})$ is not realizable, but $(\mathcal{A} \rightarrow \mathcal{G})$ is. This may be the case if \mathcal{A} represents the constraint that the environment has to perform a certain finite sequence of actions in reaction to the system's output, which is representable as a safety property. If the actions are triggered by the system, then the system cannot ensure that the environment succeeds with performing them correctly, hence violating \mathcal{A} . If $\mathcal{A} \rightarrow \mathcal{G}$ can only be realized by the system by triggering these actions at least once, then after these actions have been performed by the environment, the cooperation level specification $G\langle E \rangle \mathcal{A} \wedge (\mathcal{A} \rightarrow \mathcal{G})$ can however be enforced by the system by not triggering them again.

This observation motivates the search for *maximally cooperative implementations*, which at any point in time realize the highest possible cooperation level. Before describing how to synthesize such implementations, let us first formally define what this means.

When determining the cooperation level during the execution of a system, we only look at the part of its computation tree that is consistent with the input obtained from the environment so far. Given a computation tree $\langle T, \tau \rangle$ and the input part of a prefix trace $w^{\mathcal{J}} = w_0^{\mathcal{J}} w_1^{\mathcal{J}} \dots w_n^{\mathcal{J}} \in \mathcal{J}^*$, we define the *bobble tree* [14] of $\langle T, \tau \rangle$ for $w^{\mathcal{J}}$ to be the tree $\langle T', \tau' \rangle$, where $T' = \{\varepsilon\} \cup \{w_0^{\mathcal{J}} w_1^{\mathcal{J}} \dots w_k^{\mathcal{J}} \mid k \leq n\} \cup \{w^{\mathcal{J}} t \mid t \in \mathcal{J}^*\}$, and $\tau'(t) = \tau(t)$ for all $t \in T'$. We call $w^{\mathcal{J}}$ the *split node* of $\langle T', \tau' \rangle$ and $(\tau(\varepsilon), w_0^{\mathcal{J}})(\tau(w_0^{\mathcal{J}}), w_1^{\mathcal{J}}) \dots (\tau(w_0^{\mathcal{J}} \dots w_{n-1}^{\mathcal{J}}), w_n^{\mathcal{J}})$ the *split word* of $\langle T', \tau' \rangle$. Intuitively, the bobble tree has a single path to the split node $w^{\mathcal{J}}$ and is full from that point onwards. Cutting a full computation tree into a bobble tree does not reduce the cooperation level that the tree fulfills for the specification types in the classes that we can build from the conjuncts in D (from Eqn. 2):

Lemma 1. *Let $\langle T, \tau \rangle$ be a computation tree and $\langle T', \tau' \rangle$ be a bobble tree built from $\langle T, \tau \rangle$. If for some cooperation level specification H consisting of conjuncts in D , we have that $\langle T, \tau \rangle$ fulfills H , then $\langle T', \tau' \rangle$ also fulfills H .*

Proof. Proof by induction over the structure of H , using the semantics given on page 5. All conjuncts in D have an outermost universal quantification over the elements in $L(\langle T, \tau \rangle)$. Reducing the number of elements in $L(\langle T, \tau \rangle)$ does not make these constraints harder to fulfill. \square

Bobble trees provide us with a semantical basis for switching between cooperation levels: if a reactive system $\langle T, \tau \rangle$ for a cooperation level H executes, and after some prefix trace w , there exists a bobble tree with split word w that allows a strictly higher cooperation level H' , then it makes sense to continue the execution of the system according to cooperation level H' . We thus define:

Definition 1. *Let $(\mathcal{A}, \mathcal{G})$ be a linear-time specification. We call a computation tree $\langle T, \tau \rangle$ maximally cooperative if for every split node $t \in T$, the bobble tree induced by t and $\langle T, \tau \rangle$ implements a highest possible cooperation level for \mathcal{A} and \mathcal{G} among the bobble trees with the same split word.*

Note that since for some bobble tree, there may be multiple highest cooperation levels, it makes sense to define a preference order for the cooperation level specifications, so that when synthesizing an

implementation, the implementation can always pick the most desired one when the possibility to move up in the hierarchy arises.

In order to synthesize maximally cooperative implementations, we first need a way to check, for every split word, for the existence of a bobble tree for cooperation level specifications. The special structure of the tree automata \mathcal{T}^H built according to the construction in Sect. 5.1 offers such a way.

Definition 2. Let $\mathcal{T}^H = (Q, \mathcal{I}, \mathcal{O}, \delta, q_0, \mathcal{F})$ be a non-deterministic tree automaton for a cooperation level specification built according to the (product) construction from Sect. 5.1. We have that Q is of the shape $C_1 \times \dots \times C_n$, where for every $i \in \{1, \dots, n\}$, we either have $C_i = Q'$ or $C_i = Q' \times \mathbb{B}$ for some Rabin word automaton state set Q' . For a state $q = (c_1, \dots, c_n) \in Q$, we define $\text{unpack}(q) = \text{unpack}(c_1) \cup \dots \cup \text{unpack}(c_n)$, where we concretize $\text{unpack}(q') = q'$ and $\text{unpack}((q', b)) = q'$ for some word automaton state $q' \in Q'$ and $b \in \mathbb{B}$.

Lemma 2. Let \mathcal{T}^H be a tree automaton built according to the (product) construction from Sect. 5.1 from a cooperation level specification with conjuncts in D over the linear-time specifications $\mathcal{A}, \mathcal{G}, \mathcal{A} \rightarrow \mathcal{G}$, and $\mathcal{A} \wedge \mathcal{G}$. Let those linear-time specifications be represented by Rabin word automata with the state sets $Q_{\mathcal{A}}, Q_{\mathcal{G}}, Q_{\mathcal{A} \rightarrow \mathcal{G}}$, and $Q_{\mathcal{A} \wedge \mathcal{G}}$, respectively. Without loss of generality, let these state sets be disjoint. We have that:

1. For all reachable states q in \mathcal{T}^H , there is at most one state in $\text{unpack}(q)$ from each of $Q_{\mathcal{A}}, Q_{\mathcal{G}}, Q_{\mathcal{A} \rightarrow \mathcal{G}}$, and $Q_{\mathcal{A} \wedge \mathcal{G}}$.
2. All states in \mathcal{T}^H with the same set $\text{unpack}(q)$ have the same languages.

Proof. The first claim follows directly from the constructions from Sec. 5.1: for every state in the individual Rabin tree automata \mathcal{T}^D built from cooperation level specification conjuncts of the shapes $G(E)\varphi$ and φ , the automata always track the state of the corresponding word automata for a branch of the tree.

For the second claim, we decompose the states in \mathcal{T}^H into their factors and prove the claim for each factor individually. For factors originating from cooperation level specifications of the shape φ for some linear-time property φ , this fact is trivial. For factors originating from specification conjuncts of the shape $G(E)\varphi$, the claim follows from the fact that the tree automaton states that only differ in their Boolean flag have the same successor state functions. \square

Lemma 2 tells us how we can switch between cooperation levels. Assume that we can always read off *all* current states of $Q_{\mathcal{A}}, Q_{\mathcal{G}}, Q_{\mathcal{A} \rightarrow \mathcal{G}}$, and $Q_{\mathcal{A} \wedge \mathcal{G}}$ from $\text{unpack}(q)$ for any state q of a product tree automaton \mathcal{T}^H . This assumption can be made satisfied by letting \mathcal{T}^H be the product of *all* elements in the set of considered cooperation level specification conjuncts D , but only using the ones in the current cooperation level H when building the acceptance condition of \mathcal{T}^H . Now consider a second cooperation level specification H' that is higher in the hierarchy than H and its associated tree automaton $\mathcal{T}^{H'}$. Let W be the states in \mathcal{T}^H with a non-empty language and W' be the states of $\mathcal{T}^{H'}$ with a non-empty language. If we find a state q' in $\mathcal{T}^{H'}$ for which $\text{unpack}(q') = \text{unpack}(q)$, and state q' has a non-empty language, then we can simply re-route every transition to q to q' and obtain a new tree automaton with the states in \mathcal{T}^H and $\mathcal{T}^{H'}$ that enforces a higher cooperation level on a bobble tree along all branches in run trees that lead to q . If we now identify the non-empty tree automaton states for all cooperation level specifications in our hierarchy, and apply this approach to all pairs of the corresponding tree automata and all of their states, we end up with a tree automaton that accepts maximally cooperative computation trees. More formally, this line of reasoning shows the correctness of the following construction:

Definition 3. Let H_1, \dots, H_{14} be the cooperation level specifications of our hierarchy, ordered by preference and respecting the hierarchy's partial order \leq_H , and let $\mathcal{T}_1^H, \dots, \mathcal{T}_{14}^H$ be the non-deterministic

Rabin tree automata for them. Let us furthermore rename states q in an automaton \mathcal{T}_j^H to (q, j) to make their names unique, and let every tree automaton \mathcal{T}_j^H be given as a tuple $(Q_j, \mathcal{I}, \mathcal{O}, \delta_j, q_{0,j}, \mathcal{F}_j)$. Let $W \subseteq \bigcup_j Q_j$ be the states in the tree automata with a non-empty language. We define the Rabin tree automaton $\mathcal{T} = (Q', \mathcal{I}, \mathcal{O}, \delta', q'_0, \mathcal{F}')$ encoding the maximally cooperative synthesis problem as follows:

$$\begin{aligned} Q' &= \bigcup_{j \in \{1, \dots, 14\}} Q_j \\ q'_0 &= q_{0,j} \text{ for } j = \max\{j \in \{1, \dots, 14\} \mid q_{0,j} \in W\} \\ \mathcal{F}' &= \mathcal{F}_1 \cup \dots \cup \mathcal{F}_{14} \\ \delta'((q, j), o) &= \{ \{i \mapsto (q'', j') \mid j' = \max\{k \in \{1, \dots, 14\} \mid \exists (q'', k) \in Q_k : \\ &\quad q'' \in W, \text{unpack}(q'') = \text{unpack}(f(i)), H_j \leq_H H_k\}, \\ &\quad \text{unpack}(q'') = \text{unpack}(f(i)), \\ &\quad ((j = j') \rightarrow q'' = f(i))\} \mid f \in \delta_j(q, o) \} \\ &\text{for all } (q, j) \in Q' \text{ and } i \in \mathcal{I} \end{aligned}$$

Theorem 1. *A Rabin tree automaton built from linear-time specifications \mathcal{A} and \mathcal{G} according to Def. 3 encodes the maximally cooperative synthesis problem for the specification $(\mathcal{A}, \mathcal{G})$. Building the tree automaton and checking it for emptiness can be performed in doubly-exponential time for specifications in LTL.*

For specifications given as deterministic Rabin word automata, the time complexity is polynomial in the number of states and exponential in the number of acceptance pairs.

Proof. For the correctness, note that the tree automaton can switch between cooperation levels only finitely often, and whenever it switches, it only does so to strictly higher levels of cooperation.

To obtain doubly-exponential time complexity of maximally cooperative synthesis from LTL specifications, we first translate \mathcal{A} , \mathcal{G} , $\mathcal{A} \rightarrow \mathcal{G}$, and $\mathcal{A} \wedge \mathcal{G}$ to deterministic Rabin word automata with a doubly-exponential number of states and a singly-exponential number of acceptance condition pairs, which takes doubly-exponential time. The overall sizes of the tree automata built for the cooperation levels are then polynomial in the sizes of the Rabin word automata. We can compute W in time exponential in the number of acceptance pairs and polynomial in the number of tree automaton states, which sums up to doubly-exponential time (in the lengths of \mathcal{A} and \mathcal{G}) for LTL. When building \mathcal{T} and computing the accepted computation tree, the same argument applies. If the specification is given in form of deterministic Rabin word automata, then all the product automata computed in the process have a number of states that is polynomial in the number of states of the input automata and a number of acceptance pairs that is polynomial in the number of acceptance pairs of the input automata. By the complexity of checking Rabin tree automata for emptiness and computing W , the second claim follows as well. \square

Theorem 1 states that synthesizing maximally cooperative implementations from LTL specifications does not have a higher complexity than LTL synthesis in general. Note, however, that both the synthesis time and the size of the resulting systems can increase in practice as the Rabin automata built in cooperative synthesis are larger than in standard LTL synthesis.

Also note that we can extend the theory from this subsection to also include cooperation level specification conjuncts of the shape $\langle E \rangle \varphi$. However, we would need to add flags to the tree automata to keep track of whether the current branch in a computation tree is the one on which φ should hold. As

these flags need to be tracked along changes between the cooperation levels, the definitions from this subsection would become substantially more complicated. Thus, we refrained from doing so here.

6 Conclusion

Conventional synthesis algorithms often produce systems that actively work towards violating environment assumptions rather than satisfying assumptions and guarantees together. In this paper, we worked out a fine-grained hierarchy of cooperation levels between the system and the environment for satisfying both guarantees *and* assumptions as far as possible. We also presented a synthesis procedure that maximizes the cooperation level in the hierarchy for linear-time specifications, such as Linear Temporal Logic (LTL). The worst-case complexity of this procedure for LTL is the same as of conventional LTL synthesis. Our approach relieves the user from requiring cooperation in the specification explicitly, which helps to keep the specification clean and abstract.

In the future, we plan to work out cooperative synthesis procedures for other specification languages, and evaluate the results on industrial applications.

References

- [1] S. Almagor, U. Boker, and O. Kupferman. Formalizing and reasoning about quality. In *ICALP*, pages 15–27. Springer, 2013.
- [2] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
- [3] D. Berwanger. Admissibility in infinite games. In *STACS*, pages 188–199, 2007.
- [4] R. Bloem, K. Chatterjee, K. Greimel, T. A. Henzinger, G. Hofferek, B. Jobstmann, B. Könighofer, and R. Könighofer. Synthesizing robust systems. *Acta Inf.*, 51(3-4):193–220, 2014.
- [5] R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In *CAV*, pages 140–156, 2009.
- [6] R. Bloem, R. Ehlers, S. Jacobs, and R. Könighofer. How to handle assumptions in synthesis. In *SYNT*, pages 34–50, 2014.
- [7] R. Bloem, R. Ehlers, and R. Könighofer. Cooperative reactive synthesis. In *ATVA*. Springer, 2015. To appear.
- [8] R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa’ar. Synthesis of reactive(1) designs. *J. Comput. Syst. Sci.*, 78(3):911–938, 2012.
- [9] Laura Bozzelli. The complexity of CTL* + linear past. In *Foundations of Software Science and Computational Structures (FOSSACS)*, pages 186–200, 2008.
- [10] R. Brenguier, J.-F. Raskin, and M. Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS*, page 23, 2014.
- [11] K. Chatterjee, L. Doyen, E. Filiot, and J.-F. Raskin. Doomsday equilibria for omega-regular games. In *VMCAI*, pages 78–97, 2014.
- [12] K. Chatterjee and T. A. Henzinger. Assume-guarantee synthesis. In *TACAS*, 2007.
- [13] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.
- [14] R. Ehlers and B. Finkbeiner. Monitoring realizability. In *RV*, pages 427–441, 2011.
- [15] R. Ehlers, R. Könighofer, and R. Bloem. Synthesizing cooperative reactive mission plans. In *IROS*. IEEE, 2015.
- [16] R. Ehlers and U. Topcu. Resilience to intermittent assumption violations in reactive synthesis. In *HSCC*, pages 203–212, 2014.

- [17] M. Faella. Admissible strategies in infinite games over graphs. In *MFCS*, 2009.
- [18] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *TACAS*, 2010.
- [19] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL*, 1989.

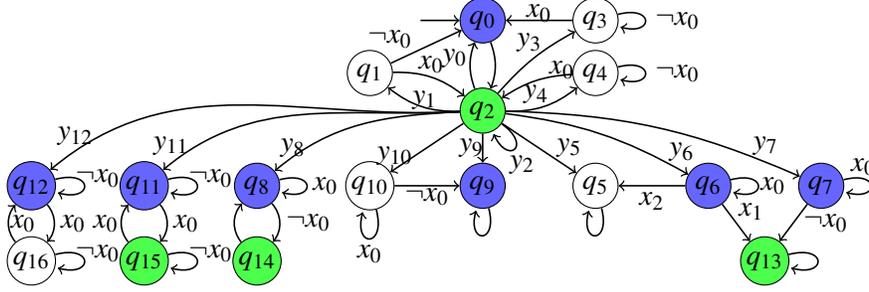


Figure 2: Example specification to illustrate all cooperation levels from Fig. 1.

A Another Example

To complement Section 4.2, this section illustrates *all* cooperation levels from Fig. 1 on a more technical example. The assumptions \mathcal{A} are defined using the Rabin word automaton $\mathcal{R}_{\mathcal{A}} = (Q, \mathcal{I} \times \mathcal{O}, \delta, q_0, \{(\emptyset, G_{\mathcal{A}})\})$ shown in Fig. 2, where $Q = \{q_0, \dots, q_{16}\}$, the input alphabet is $\mathcal{I} = \{x_0, x_1, x_2\}$, and the output alphabet is $\mathcal{O} = \{y_0, \dots, y_{12}\}$. Fig. 2 labels edges with conditions over the input and output letters to define the transition function δ . For instance, the condition $\neg x_0$ means that the transition is taken for all letters $(x, y) \in \mathcal{I} \times \mathcal{O}$ where $x \neq x_0$. Edges without a label are always taken. Note that the system can only influence the next state from q_2 : with output y_i , the next state will be q_i . The acceptance condition is defined with $G_{\mathcal{A}} = \{q_0, q_6, q_7, q_8, q_9, q_{11}, q_{12}\}$. That is, the blue states in Fig. 2 must be visited infinitely often for \mathcal{A} to be satisfied. The guarantees \mathcal{G} are defined using the automaton $\mathcal{R}_{\mathcal{G}} = (Q, \mathcal{I} \times \mathcal{O}, \delta, q_0, \{(\emptyset, G_{\mathcal{G}})\})$, which differs from $\mathcal{R}_{\mathcal{A}}$ only in the acceptance condition: with $G_{\mathcal{G}} = \{q_2, q_{13}, q_{14}, q_{15}\}$, \mathcal{G} is satisfied if the green states in Fig. 2 are visited infinitely often.

The following paragraphs present and discuss one system behavior (defined in the form of a computation tree) per cooperation level in Fig. 1.

$\mathcal{A} \rightarrow \mathcal{G}$: The computation tree $\langle \mathcal{I}^*, \tau_5 \rangle$ with $\tau_5(w^{\mathcal{I}}) = y_5$ for all $w^{\mathcal{I}} \in \mathcal{I}^*$ represents a system that always outputs y_5 . It satisfies $\mathcal{A} \rightarrow \mathcal{G}$ because q_5 is a trap where neither \mathcal{A} nor \mathcal{G} is satisfied. This is “correct” but rather unsatisfactory.

$(\mathcal{A} \rightarrow \mathcal{G}) \wedge \mathbf{G}\langle E \rangle \mathcal{G}$: is satisfied by the computation tree $\langle \mathcal{I}^*, \tau_4 \rangle$ with $\tau_4(w^{\mathcal{I}}) = y_4$. Again, $\mathcal{A} \rightarrow \mathcal{G}$ holds because staying in $\{q_2, q_4\}$ violates \mathcal{A} . Still, at any point in time, there exists some future environment behavior that satisfies \mathcal{G} (namely one that chooses x_0 infinitely often). Thus, τ_4 is slightly better than τ_5 .

$(\mathcal{A} \rightarrow \mathcal{G}) \wedge \mathbf{G}\langle E \rangle \mathcal{A}$: is satisfied by $\langle \mathcal{I}^*, \tau_3 \rangle$ with $\tau_3(w^{\mathcal{I}}) = y_3$: At any point in time, \mathcal{A} can be satisfied for some future environment behavior (namely if x_0 is chosen in q_3 infinitely often). Since $\mathcal{A} \rightarrow \mathcal{G}$ (because of the unconditional edge from q_0 to q_2), this implies that also \mathcal{G} can be reached. Yet, neither \mathcal{A} nor \mathcal{G} are enforced because the environment could always give x_1 to get stuck in q_3 . However, assuming that the environment does not behave totally self-destructive, both \mathcal{A} and \mathcal{G} will be satisfied. Thus, τ_3 is even better than τ_4 .

\mathcal{G} : The computation tree with $\tau_2(w^{\mathcal{I}}) = y_2$ enforces \mathcal{G} , but defeats any hope of reaching \mathcal{A} . It is thus better than τ_2 but incomparable with τ_3 : Enforcing \mathcal{G} is better than having \mathcal{G} reachable, but \mathcal{A} being unreachable is worse than having \mathcal{A} reachable. Hence, the choice between τ_2 and τ_3 is a question of selfishness.

$\mathcal{G} \wedge \mathbf{G}\langle E \rangle \mathcal{A}$: is satisfied by $\tau_1(w^{\mathcal{I}}) = y_1$: No matter if the environment picks x_0 or not in q_1 , the state q_2 will always be visited because of the unconditional edge from q_0 to q_2 . Visiting q_0 infinitely often

(and thus satisfying \mathcal{A}) is always possible, but not enforced. Hence, τ_1 dominates both τ_2 and τ_3 .

$\mathcal{G} \wedge \mathcal{A}$: is enforced by $\tau_0(w^{\mathcal{E}}) = y_0$: independent of the environment, both q_0 and q_2 are visited repeatedly. Thus, τ_0 is the most desirable system behavior.

The remaining cooperation levels from the hierarchy in Fig. 1 do not satisfy the traditional correctness criterion $\mathcal{A} \rightarrow \mathcal{G}$. Still, a computation tree for such a level behaves better than arbitrarily in various ways, and are useful from states where $\mathcal{A} \rightarrow \mathcal{G}$ cannot be enforced.

$\mathbf{G}\langle E \rangle(\mathcal{A} \rightarrow \mathcal{G})$: is satisfied by the computation tree $\langle \mathcal{I}^*, \tau_6 \rangle$ with $\tau_6(w^{\mathcal{E}}) = y_6$ for all $w^{\mathcal{E}} \in \mathcal{I}^*$. At any point in time, there exists some environment behavior that either violates \mathcal{A} (by going to q_5) or that satisfies \mathcal{G} (by going to q_{13}). However, \mathcal{G} alone is not reachable from every point in every execution: if the environment already traversed to q_5 , \mathcal{G} is lost. Similarly, \mathcal{A} is not reachable at every point in time because the environment may already have traversed to q_{13} . Finally, $\mathcal{A} \rightarrow \mathcal{G}$ is not enforced because with input x_0 the environment may stay in q_6 forever, thereby satisfying \mathcal{A} but not \mathcal{G} . Nevertheless, having $\mathcal{A} \rightarrow \mathcal{G}$ feasible at any point in time is better than nothing.

$\mathbf{G}\langle E \rangle \mathcal{G}$: The computation tree with $\tau_7(w^{\mathcal{E}}) = y_7$ satisfies $\mathbf{G}\langle E \rangle \mathcal{G}$, and is thus better than τ_6 . $\mathbf{G}\langle E \rangle \mathcal{A}$ does not hold because from q_{13} , \mathcal{A} is unreachable. $\mathcal{A} \rightarrow \mathcal{G}$ is not enforced because the environment could stay in q_7 forever. Yet, τ_7 is still incomparable with τ_5 (enforcing $\mathcal{A} \rightarrow \mathcal{G}$), because with τ_5 , there is not even a hope of reaching \mathcal{G} .

$\mathbf{G}\langle E \rangle \mathcal{A}$: is similar to the previous case. The computation tree with $\tau_{10}(w^{\mathcal{E}}) = y_{10}$ satisfies $\mathbf{G}\langle E \rangle \mathcal{A}$, but \mathcal{A} is not enforced because the environment can stay in q_{10} . Moreover, $\mathbf{G}\langle E \rangle(\mathcal{A} \rightarrow \mathcal{G})$ does not hold because it is unreachable from q_9 .

\mathcal{A} : The computation tree with $\tau_9(w^{\mathcal{E}}) = y_9$ enforces \mathcal{A} , and thus dominates τ_{10} . Yet all hopes for satisfying the guarantees are lost.

$\mathcal{A} \wedge \mathbf{G}\langle E \rangle \mathcal{G}$: The system behavior $\tau_8(w^{\mathcal{E}}) = y_8$ enforces $\mathcal{A} \wedge \mathbf{G}\langle E \rangle \mathcal{G}$ and eliminates this defect of τ_9 : if the environment is not hostile and produces x_1 or x_2 infinitely often, \mathcal{G} is achievable.

$(\mathbf{G}\langle E \rangle \mathcal{A}) \wedge \mathbf{G}\langle E \rangle(\mathcal{A} \rightarrow \mathcal{G})$: is satisfied by the system behavior $\tau_{12}(w^{\mathcal{E}}) = y_{12}$. The guarantee \mathcal{G} is lost, but at any point in time, it is still possible that \mathcal{A} is satisfied (by visiting s_{12} infinitely often) and that \mathcal{A} is violated (by only visiting s_{16} from some point on). This is slightly better than τ_{10} (which satisfies $\mathbf{G}\langle E \rangle \mathcal{A}$ alone) because the correctness property $\mathcal{A} \rightarrow \mathcal{G}$ is reachable at any point. With τ_{10} , it may happen that the correctness property is lost once and for all.

$(\mathbf{G}\langle E \rangle \mathcal{A}) \wedge (\mathbf{G}\langle E \rangle \mathcal{G})$ and $\mathbf{G}\langle E \rangle(\mathcal{A} \wedge \mathcal{G})$: The system behavior $\tau_{11}(w^{\mathcal{E}}) = y_{11}$ satisfies $(\mathbf{G}\langle E \rangle \mathcal{A}) \wedge (\mathbf{G}\langle E \rangle \mathcal{G})$ as well as $\mathbf{G}\langle E \rangle(\mathcal{A} \wedge \mathcal{G})$. In our example, there is no way to distinguish these two levels because we chose objectives defined by visiting some states infinitely often. For other specification classes, there can be a difference, though. \mathcal{A} is not enforced by τ_{11} because the environment could always stay in q_{15} . $\mathcal{A} \rightarrow \mathcal{G}$ is not enforced either because the environment could stick to q_{11} . Still, this is better than a prospect of reaching only \mathcal{A} or only \mathcal{G} .

B Extended Cooperation Hierarchies

As discussed in Section 4.3, we can extend our hierarchy in various ways. When extending the conjuncts from Eqn. 2 with $\mathcal{A} \vee \mathcal{G}$ and $\mathbf{G}\langle E \rangle(\mathcal{A} \vee \mathcal{G})$, we obtain the hierarchy shown in Figure 3, which has 23 levels. We applied the following reduction rules in addition to those mentioned in Section 4.1:

- \mathcal{G} implies $\mathcal{A} \vee \mathcal{G}$,
- \mathcal{A} implies $\mathcal{A} \vee \mathcal{G}$,
- $\mathbf{G}\langle E \rangle \mathcal{G}$ implies $\mathbf{G}\langle E \rangle(\mathcal{A} \vee \mathcal{G})$,

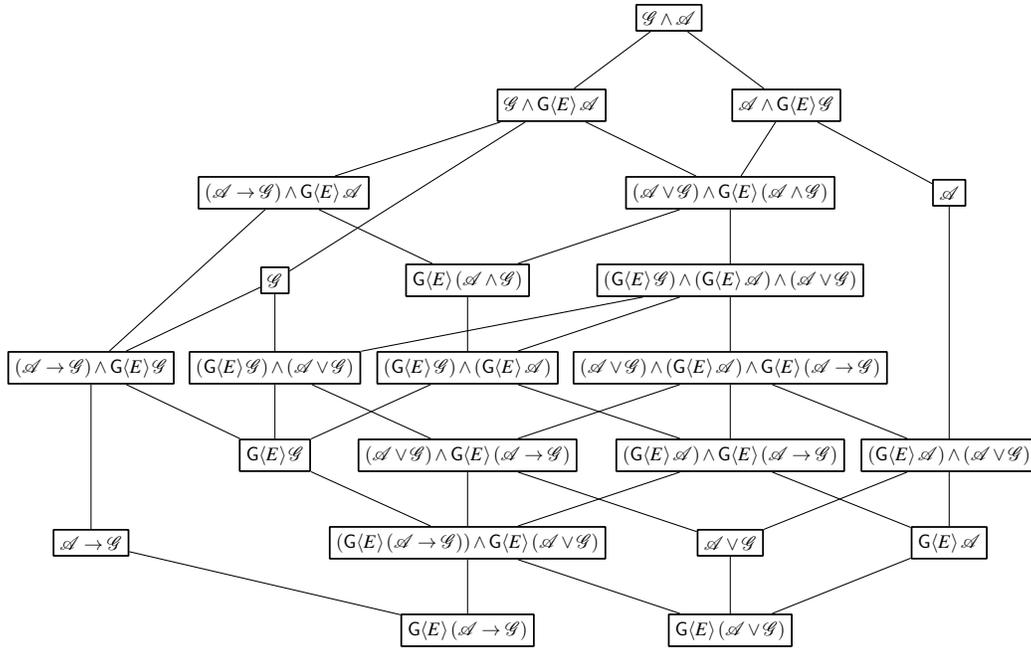


Figure 3: A refined hierarchy of cooperation levels.

- $G\langle E \rangle A$ implies $G\langle E \rangle (A \vee G)$,
- $A \rightarrow G$ and $A \vee G$ together imply G , and
- $A \rightarrow G$ and $G\langle E \rangle (A \vee G)$ together imply $G\langle E \rangle G$.

When further extending the conjuncts from Eqn. 2 with $\langle E \rangle (G)$, $\langle E \rangle (A)$, $\langle E \rangle (A \wedge G)$, $\langle E \rangle (A \rightarrow G)$ and $\langle E \rangle (A \vee G)$ we obtain a hierarchy of 77 cooperation levels using the following additional reduction rules:

- $G\langle E \rangle \varphi$ implies $\langle E \rangle \varphi$ for all $\varphi \in \{G, A, A \wedge G, A \rightarrow G, A \vee G\}$
- $\langle E \rangle G$ implies $\langle E \rangle (A \rightarrow G)$,
- $\langle E \rangle (A \wedge G)$ implies $\langle E \rangle A$ and $G\langle E \rangle G$,
- $A \rightarrow G$ and $\langle E \rangle A$ together imply $\langle E \rangle (A \wedge G)$,
- A and $\langle E \rangle G$ together imply $\langle E \rangle (A \wedge G)$,
- $\langle E \rangle (A \rightarrow G)$ and A together imply $\langle E \rangle (G)$.
- $\langle E \rangle G$ implies $\langle E \rangle (A \vee G)$,
- $\langle E \rangle A$ implies $\langle E \rangle (A \vee G)$,
- $A \rightarrow G$ and $\langle E \rangle (A \vee G)$ together imply $\langle E \rangle G$.

The first six reduction rules directly correspond to the rules from Section 4.1 but using $\langle E \rangle$ instead of $G\langle E \rangle$. The latter three rules correspond to rules from the previous paragraph, again instantiated with $G\langle E \rangle$ instead of $\langle E \rangle$.