Information Flow Guided Synthesis*

Bernd Finkbeiner[1], Niklas Metzger[1], and Yoram Moses[2]

[1] CISPA Hemlholtz Center of Information Security, Saarland, Germany
{finkbeiner, niklas.metzger}@cispa.de
[2] The Andrew and Erna Viterbi Faculty of Electrical and Computer Engineering, Technion, Israel
moses@ee.technion.ac.il

Abstract. Compositional synthesis relies on the discovery of assumptions, i.e., restrictions on the behavior of the remainder of the system that allow a component to realize its specification. In order to avoid losing valid solutions, these assumptions should be *necessary* conditions for realizability. However, because there are typically many different behaviors that realize the same specification, necessary behavioral restrictions often do not exist. In this paper, we introduce a new class of assumptions for compositional synthesis, which we call *information flow assumptions*. Such assumptions capture an essential aspect of distributed computing, because components often need to act upon information that is available only in other components. The presence of a certain flow of information is therefore often a necessary requirement, while the actual behavior that establishes the information flow is unconstrained. In contrast to behavioral assumptions, which are properties of individual computation traces, information flow assumptions are *hyperproperties*, i.e., properties of sets of traces. We present a method for the automatic derivation of information-flow assumptions from a temporal logic specification of the system. We then provide a technique for the automatic synthesis of component implementations based on information flow assumptions. This provides a new compositional approach to the synthesis of distributed systems. We report on encouraging first experiments with the approach, carried out with the BoSyHyper synthesis tool.

1 Introduction

In *distributed synthesis*, we are interested in the automatic translation of a formal specification of a distributed system's desired behavior into an implementation that satisfies the specification [26]. What makes distributed synthesis far more interesting than the standard synthesis of reactive systems, but also more challenging, is that the result consists of a set of implementations of subsystems, each of which operates based only on partial knowledge of the global system state. While algorithms for distributed synthesis have been studied since the

* This work was funded by DFG grant 389792660 as part of TRR 248 – CPEC, by the European Research Council (ERC) Grant OSARES (No. 683300), and by the German Israeli Foundation (GIF) Grant No. I-1513-407./2019.

1990s [12,20,26], their high complexity has resulted in applications of distributed synthesis being, so far, very limited.

One of the most promising approaches to making distributed synthesis more scalable is *compositional synthesis* [7,10,14,21,27]. The compositional synthesis of a distributed system with two processes, p and q, avoids the construction of the product of p and q and instead focuses on one process at a time. Typically, it is impossible to realize one process without making certain assumptions about the other process. Compositional synthesis therefore critically depends on finding the assumption that p must make about q, and vice versa: once the assumptions are known, one can build each individual process, relying on the fact that the assumption will be satisfied by the synthesized implementation of the other process. Ideally, the assumptions should be both *sufficient* (i.e., the processes are realizable under the assumptions) and *necessary* (i.e., any implementation that satisfies the specification would also satisfy the assumptions). Without sufficiency, the synthesis cannot find a compositional solution; without necessity, the synthesis loses valid solutions. While sufficiency is obviously checked as part of the synthesis process, it is often impossible to find necessary conditions, because there the specifications can be realized by many different behaviors. Any specific implementation would lead to a specific assumption; however, this implementation is only known once the synthesis is complete, and an assumption that is satisfied by *all* implementations often does not exist.

In this paper, we propose a way out of this chicken-and-egg type of situation. Previous work on generating assumptions for compositional synthesis has focused on *behavioral* restrictions on the environment of a subsystem. We introduce a new class of more abstract assumptions that, instead, focus on the *flow of information*. Consider a system architecture (depicted in Figure 1a) where two processes a and b are linked by a communication channel `c`, such that a can write to `c` and b can read from `c`. Suppose also that a reads a boolean input `in` from the environment that is, however, not directly visible to b. We are interested in a distributed implementation for a specification that demands that b should eventually output the value of input `in`. Since b cannot observe `in`, its synthesis must rely on the assumption that the value of `in` is communicated over the channel `c` by process a. Expressing this as a *behavioral assumption* is difficult, because there are many different behaviors that accomplish this. Process a could, for example, literally copy the value of `in` to `c`. It could also encode the value, for example by writing to `c` the negation of the value of `in`. Alternatively, it could delay the transmission of `in` by an arbitrary number of steps, and even use the length of the delay to encode information about the value of `in`. Fixing any such communication protocol, by a corresponding behavioral assumption on a, would unnecessarily eliminate potential implementations of b. The minimal assumption that subsystem a must rely on is in fact an information-flow assumption, namely that b will eventually learn the value of `in`.

We present a method that derives necessary information flow assumptions automatically. A fundamental difference between behavioral and information flow assumptions is that behavioral assumptions are *trace properties*, i.e., properties

of individual traces; by contrast, information flow assumptions are *hyperproperties*, i.e., properties of *sets* of traces. In our example, the assumption that a will eventually communicate the value of `in` to b is the hyperproperty that any two traces that differ in the value of `in` must eventually also differ in `c`. The precise difference between the two traces depends on the communication protocol chosen in the implementation of a; however, any correct implementation of a must ensure that some difference in b's input (on channel `c`) in the two traces occurs, so that b can then respond with a different output.

Once we have obtained information flow assumptions for all of the subsystems, we proceed to synthesize each subsystem under the assumption generated for its environment. It is important to note that, at this point, the implementation of the environment is not known yet; as a result, we only know *what* information will be provided to process b, but not *how*. This also means that we cannot yet construct an executable implementation of the process under consideration; after all, this implementation would need to correctly decode the information provided by its partner process. Clearly, we cannot determine how to *decode* the information before we know how the implementation of the sending process *encodes* the information!

Our solution to this quandary is to synthesize a prototype of an implementation for the process that works with *any* implementation of the sender, as long as the sender satisfies the information flow requirement. The prototype differs from the actual implementation in that it has access to the original (unencoded) information. Because of this information the prototype, which we call a *hyper implementation*, can determine the correct output that satisfies the specification. Later, in the actual implementation, the information is no longer available in its original, unencoded form, but must instead be decoded from the communication received from the environment. However, the information flow assumption guarantees that this is actually possible, and access to the original information is, therefore, no longer necessary.

In Section 2, we explain our approach in more detail, continuing the discussion of the bit transmission example mentioned above. The paper then proceeds to make the following contributions:

- We introduce the notion of *necessary information flow assumptions* (Section 4.1) for distributed systems with two processes and present a method for the automatic derivation of such assumptions from process specifications given in linear-time temporal logic (LTL).
- We strengthen information flow assumptions to the notion of *time-bounded* information flow assumptions (Section 4.2), which characterizes information that must be received in finite time. We introduce the notion of *uniform distinguishability* and prove that uniform distinguishability guarantees the necessity of the information flow assumption.
- We introduce the notion of *hyper implementations* (Section 5) and provide a synthesis method for their automatic construction. We also explain how to transform hyper implementations into actual process implementations.

- We present a *practical approach* (Section 6) that simplifies the synthesis for cases where the information flow assumption refers to a finite amount of information.
- We report on encouraging experimental results (Section 7).

Related work. Compositional synthesis is often studied in the setting of *complete information*, where all processes have access to all environment outputs [11, 15, 19, 22]. In the following, we focus on compositional approaches for the synthesis of distributed systems, where the processes have incomplete information about the environment outputs. Compositionality has been used to improve distributed synthesis in various domains, including reactive controllers [1,18]. Closest to our approach is assume-guarantee synthesis [3, 4], which relies on behavioral guarantees of the processs behaviour and assumptions about the behavior of the other processes. Recently, an extension of assume-guarantee synthesis for distributed systems was proposed [24], where the assumptions are iteratively refined. Using a weaker winning condition for synthesis, remorse-free dominance [8] avoids the explicit construction of assumptions and guarantees, resulting in implicit assumptions. A recent approach [16] uses behavioral guarantees in the form of certificates to guide the synthesis process. Certificates specify partial behaviour of each component and are iteratively synthesized. The fundamental difference between all these approaches to this work is that the assumptions are behavioral. To the best of our knowledge, this is the first synthesis approach based on information-flow assumptions. While there is a rich body of work on the verification of information-flow properties (cf. [9, 17, 29]), and the synthesis from information-flow properties and other hyperproperties has also been studied before (cf. [13]), the idea of utilizing hyperproperties as assumptions for compositional synthesis is new.

2 The Bit Transmission Problem

We use the *bit transmission* example from the introduction to motivate our approach. The example consists of two processes a and b that are combined into the distributed architecture shown in Figure 1a. Process a observes the (binary) input of the environment through variable `in` and can communicate with the second process b via a channel (modeled by the shared variable `c`). Process b observes its own local input from a and has a local output `out`. We are interested in synthesizing an implementation for our distributed system consisting of two strategies, one for each process, whose combined behavior satisfies the specification. In this example, the specification for process b is to transmit the initial value of `in`, an input of a, to b's own output; this is expressed by the linear-time temporal logic (LTL) formula $\varphi_b = \texttt{in} \leftrightarrow \Diamond \texttt{out}$. The specification does not restrict a's behavior, such that $\varphi_a = \mathit{true}$. Since the value of `out` is controlled by b, whereas `in` is determined by the environment and observed by a, this specification forces b to react to an input that b neither observes nor controls. To satisfy the goal, `out` must remain *false* forever if `in` is initially *false*,

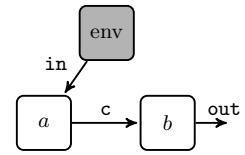

(a) The bit transmission architecture

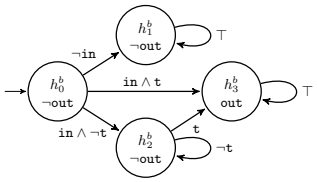

(b) The hyper implementation of b

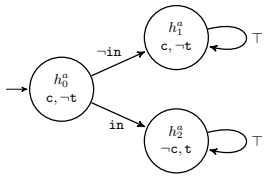

(c) The hyper implementation of a

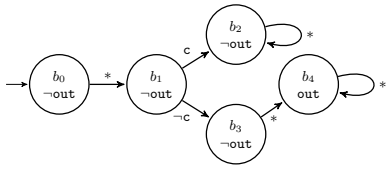

(d) The implementation of b

Fig. 1: The distributed system of the *bit transmission* protocol. The architecture is given in (a), the hyper-implementation of b in (b), the local implementation of a in(c), and the resulting local implementation of b in (d).

while `out` must eventually become *true* at least once if `in` starts with value *true*. Indeed, in order to set `out` to *true*, process b must *know* that `in` is initially *true*, which can only be satisfied via information flow from a to b. We can capture this information flow requirement as the following hyperproperty: For every pair of traces that disagree on the initial value of `in`, process a must (eventually) behave differently on `c`. The requirement can be expressed in HyperLTL by the formula $\Psi = \forall\pi,\pi'.(\mathtt{in}_\pi \not\leftrightarrow \mathtt{in}_{\pi'}) \rightarrow \Diamond(\mathtt{c}_\pi \not\leftrightarrow \mathtt{c}_{\pi'})$. The information flow requirement does not restrict a to behave in a particular manner; the *encoding* of the information about `in` on the channel `c` depends on a's behavior. Under the assumption that a will behave according to the information flow requirement Ψ, one can synthesize a solution of b that is correct for every implementation of a. Given its generality, we call such a solution a *hyper implementation*, shown in Figure 1b. Since the point in time when the information is received by b is unknown during the local synthesis process, an additional boolean variable `t` is added to the specification of b. This variable signals that the information has been transmitted and is later derived by a's implementation. Setting `out` to *true* is only allowed after `t` is observed by process b. When the hyper implementation is composed with the actual implementation of a, as shown in Figure 1c, both local specifications are satisfied. The resulting local implementation of b, depicted in Figure 1d, branches only on local inputs and, together with a, satisfies the specification. While changing state b_0 to b_1, process b cannot distinguish `in` from ¬`in`. It has to wait for one time step, i.e., the first difference in outputs of process a, to observe the difference in the shared communication channel. The value of `t` is obtained from a's implementation and set to *true* with the first difference in `c`, forbidding the edge from h_0^b to h_3^bin the local implementation of b.

3 Preliminaries

Architectures. For ease of exposition we focus in this paper on systems with two processes. Let $\mathcal{V}$ be a set of variables. An architecture with two black-box processes p and q is given as a tuple $(I_p, I_q, O_p, O_q, I_e)$, where I_p, I_q, O_p, O_q, and I_e are all subsets of $\mathcal{V}$. O_p and O_q are the *output variables* of p and q. O_e are the output variables of the uncontrollable environment. The three sets O_p, O_q and O_e form a partition of $\mathcal{V}$. I_p and I_q are the *input variables* of processes p and q, respectively. For each black-box process, the inputs and outputs are disjoint, i.e., $I_p \cap O_p = \emptyset$ and $I_q \cap O_q = \emptyset$. The inputs I_p and I_q of the black-box processes are all either outputs of the environment or outputs of the other black-box process, i.e., $I_p \subseteq O_q \cup O_e$ and $I_q \subseteq O_p \cup O_e$. We assume that all variables are of boolean type. For a set $V \subseteq \mathcal{V}$, every subset $V' \subseteq V$ defines a *valuation* of V, where the variables in $V \cap V'$ have value *true* and the variables in $V \setminus V'$ have value *false*.

Implementations. An implementation of an architecture $(I_p, I_q, O_p, O_q, I_e)$ is a pair (s_p, s_q), consisting of a strategy for each of the two black-box processes. A *strategy* for a black-box process p is a function $s_p : (2^{I_p})^* \to (2^{O_p})$ that maps finite sequences of valuations of p's input variables (i.e., *histories* of inputs) to a valuation of p's output variables. The (synchronous) *composition* $s_p || s_q$ of the two strategies is the function $s : (2^{O_e})^* \to (2^{\mathcal{V}})$ that maps finite sequences of valuations of the environment's output variables to valuations of all variables: we define $s(\epsilon) = s_p(\epsilon) \cup s_q(\epsilon)$ and, for $v \in (2^{O_e})^*, x \in 2^{O_e}$, $s(v \cdot x) = (s_p(f_p(v)) \cup s_q(f_q(v)) \cup x)$, where f_p and f_q map sequences of environment outputs to sequences of process inputs with $f_p(\epsilon) = \epsilon, f_p(v \cdot x) = f_p(v) \cdot ((x \cup s_q(f_q(v))) \cap I_p)$ and $f_q(\epsilon) = \epsilon, f_q(v \cdot x) = f_p(v) \cdot ((x \cup s_p(f_p(v))) \cap I_q)$.

Specifications. Our specifications refer to traces over the set $\mathcal{V}$ of all variables. In general, for a set $V \subseteq \mathcal{V}$ of variables, a *trace* over V is an infinite sequence $x_0 x_1 x_2 \ldots \in (2^{\mathcal{V}})^\omega$ of valuations of V. A *specification* $\varphi \subseteq (2^{\mathcal{V}})^\omega$ is a set of traces over $\mathcal{V}$. Two traces of disjoint sets $V, V' \subset \mathcal{V}$ can be *combined* by forming the union of their valuations at each position, i.e., $x_0 x_1 x_2 \ldots \sqcup y_0 y_1 y_2 \ldots = (x_0 \cup y_0)(x_1 \cup y_1)(x_2 \cup y_2) \ldots$. Likewise, the *projection* of a trace onto a set of variables $V' \subseteq \mathcal{V}$ is formed by intersecting the valuations with V' at each position: $x_0 x_1 x_2 \ldots \downarrow_{V'} = (x_0 \cap V')(x_1 \cap V')(x_2 \cap V') \ldots$.

For our specification language, we use propositional linear-time temporal logic (LTL) [25], with the set $\mathcal{V}$ of variables as atomic propositions and the usual temporal operators Next $\bigcirc$, Until $\mathcal{U}$, Globally $\Box$, and Eventually $\Diamond$. System specifications are given as a conjunction $\varphi_p \wedge \varphi_q$ of two LTL formulas, where φ_p refers only to variables in $O_p \cup O_e$, i.e., the formula relates the outputs of process p to the outputs of the environment, and φ_q refers only to variables in $O_q \cup O_e$. The two formulas represent the *local specifications* for the two black-box processes. An implementation $s = (s_p, s_q)$ defines a set of traces

$$\mathit{Traces}(s_p, s_q) = \{x_0 x_1 \ldots \in (2^O)^\omega \mid x_k = s(i_0 i_1 \ldots i_{k-1}) \text{ for all } k \in \mathbb{N} \text{ for some } i_o i_1 i_2 \ldots \in (2^{O_e})^\omega\}.$$

We say that the implementation *satisfies* the specification if the traces of the implementation are contained in the specification, i.e., $\mathit{Traces}(s_p, s_q) \subseteq \varphi$.
The synthesis problem. Given an architechture and a specification φ, the synthesis problem is to find an implementation s that satisfies φ. We say that a specification φ is *realizable* in a given architecture if such an implementation exists, and *unrealizable* if not.
Hyperproperties. We capture information-flow assumptions as hyperproperties. A *hyperproperty over* $\mathcal{V}$ is a set $H \subseteq 2^{(2^{\mathcal{V}})^\omega}$ of sets of traces over $\mathcal{V}$ [6]. An implementation (s_p, s_q) satisfies the hyperproperty H iff its traces are an element of H, i.e., $\mathit{Traces}(s_p, s_q) \in H$. A specification language for hyperproperties is the temporal logic HyperLTL [5]. HyperLTL extends LTL with quantification over trace variables. The syntax of HyperLTL is given by the following grammar $\varphi ::= \forall \pi. \varphi \mid \exists \pi. \varphi \mid \psi$ and $\psi ::= v_\pi \mid \neg\psi \mid \psi \wedge \psi \mid \bigcirc \psi \mid \psi \mathcal{U} \psi$ where $v_\pi \in \mathcal{V}$ is a variable and $\pi \in \mathcal{T}$ is a trace variable. Note that the output variables are indexed by trace variables. The quantification over traces makes it possible to express properties like "ψ *must hold on all traces*", which is expressed by $\forall \pi.\ \psi$. Dually, one can express that "*there exists a trace on which* ψ *holds*", denoted by $\exists \pi.\ \psi$. The temporal operators are defined as in LTL.

In some cases, a hyperproperty can be expressed in terms of a binary relation on traces. A relation $R \subseteq (2^{\mathcal{V}})^\omega \times (2^{\mathcal{V}})^\omega$ of pairs of traces defines the hyperproperty H, where a set T of traces is an element of H iff for all pairs $\pi, \pi' \in T$ of traces in T it holds that $(\pi, \pi') \in R$. We call a hyperproperty defined in this way a *2-hyperproperty*. In HyperLTL, 2-hyperproperties are expressed as formulas with two universal quantifiers and no existential quantifiers. A 2-hyperproperty can equivalently be represented as a set of infinite sequences over the product alphabet Σ^2: for a given 2-hyperproperty $R \subseteq \Sigma^\omega \times \Sigma^\omega$, let $R' = \{(\sigma_0, \sigma'_0)(\sigma_1, \sigma'_1) \ldots \mid (\sigma_0 \sigma_1 \ldots, \sigma'_0 \sigma'_1 \ldots) \in R\}$. This representation is convenient for the use of automata to recognize 2-hyperproperties.

4 Necessary Information Flow in Distributed Systems

In reactive synthesis it is natural that the synthesized process reacts to different environment outputs. This is also the case for distributed synthesis, where some outputs of the environment are not observable by a local process and the hidden values must be communicated to the process. In the following we show when such information flow is necessary.

4.1 Necessary Information Flow

Our analysis focuses on pairs of situations for which the specification dictates a *different* reaction from a given black-box process p. Such pairs imply the need for information flow that will enable p to distinguish the two situations: if p cannot distinguish the two situations, it will behave in the same manner in both. Consequently, the specification will be violated, no matter how p is implemented, in at least one of the two situations. A process p needs to satisfy

a local specification φ_p, which relates its outputs O_p to the outputs O_e of the environment. (Recall that O_e may contain inputs to the other black-box process.) We are therefore interested in pairs of traces over O_e for which φ_p does *not* admit a common valuation of O_p. We collect such pairs of traces in a *distinguishability relation*, denoted by Δ_p:

Definition 1 (Distinguishability). *Given a local specification φ_p for process p, the* distinguishability relation *Δ_p is the set of pairs of traces over O_e (environment outputs) such that no trace over O_p satisfies φ_p in combination with both traces in the pair. Formally:*

$$\Delta_p = \{(\pi_e, \pi'_e) \in (2^{O_e})^\omega \times (2^{O_e})^\omega \mid \\ \forall \pi_p \in (2^{O_p})^\omega.\ \textit{if}\ \ \pi_e \sqcup \pi_p \vDash \varphi_p\ \ \textit{then}\ \ \pi'_e \sqcup \pi_p \nvDash \varphi_p\ \}$$

By definition of Δ_p, process p must distinguish π_e from π'_e, because it cannot respond to both in the same manner. In our running example, Δ_b consists of all pairs of sequences of values of `in` that differ in the first value of `in`. Process b must act differently in such situations: if `in` is initially *true* then b must eventually set `out` to *true*, while if it starts as *false*, then b must keep `out` always set to *false*.

In general, a black-box process p must satisfy its specification φ_p despite having only partial access to O_e. The distinguishability relation therefore directly defines an *information flow* requirement: In order to satisfy φ_p, enough information about O_e must be communicated to p via its local inputs I_p to ensure that p can distinguish any pair of traces in Δ_p. We formalize this information flow assumption as a 2-hyperproperty, which states that if the outputs of the environment in the two traces must be distinguished, i.e, the projection on O_e is in Δ_p, then there must be a difference in the local inputs I_p:

Definition 2 (Information flow assumption). *The* information flow assumption *ψ_p induced by Δ_p is the 2-hyperproperty defined by the relation*

$$R = \{(\pi, \pi') \in (2^{\mathcal{V}})^\omega \times (2^{\mathcal{V}})^\omega \mid (\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \in \Delta_p \textit{ then } \pi{\downarrow_{I_p}} \neq \pi'{\downarrow_{I_p}}\}$$

In our running example, the information flow assumption for process b requires that on any two executions that disagree on the initial value of `in`, the values communicated to b over the channel `c` must differ at some point. Observe that the information flow assumption ψ_p specifies neither how the information is to be encoded on `c` nor the point in time when the different communication occurs. However, ψ_p requires that the communication differs eventually if the initial values of `in` are different. Moreover, notice that both Δ_p and ψ_p are determined by p's specification φ_p. The following theorem shows that the information flow assumption ψ_p is a necessary condition, the proof can be found in Appendix D.1.

Theorem 1. *Every implementation that satisfies the local specification φ_p for p also satisfies the information flow assumption ψ_p.*

4.2 Time-bounded Information Flow

We now introduce a strengthened version of the information flow assumption. As shown in Theorem D.1, the information flow assumption is a necessary condition for the existence of an implementation that satisfies the specification. Often, however, the information flow assumption is not strong enough to allow for the separate synthesis of individual components in a compositional approach.

Consider again process b in our motivating example. The information flow assumption guarantees that any pair of traces that differ in the initial value of the global input `in` will differ at some point in the value of the channel `c`. This assumption is not strong enough to allow process b to satisfy the specification that b must eventually set `out` to *true* iff the initial value of `in` is *true*. Suppose that `in` is *true* initially. Then b must at some point set `out` to *true*. Process b can only do so when it *knows* that the initial value of `in` is *true*. The information flow assumption is, however, too weak to guarantee that process b will eventually obtain this knowledge. To see this, consider a hypothetical behavior of process a that sets `c` forever to *true*, if `in` is *true* in the first position, and if `in` is *true* then a keeps `c` true for $n-1$ steps, where $n > 0$ is some fixed natural number, before it sets `in` to *false* at the n^{th} step. This behavior of process a satisfies the information flow assumption for any number n; however, without knowing n, process b does not know how many steps it should wait for `in` to become *false*. If, at any point in time t, the channel `c` has not yet been set to *false*, process b can never rule out the possibility that the initial value of `in` is *true*; it might simply be the case that $t < n$ and, hence, the time when `c` will be set to *false* still lies in the future of t! Hence, process b can never actually set `out` to *true*.

We begin by presenting a finer version of the distinguishability relation from Definition 1 that we call *time-bounded distinguishability*. Recall that by Definition 1, a pair (π_e, π'_e) is in the distinguishability relation Δ_p if every output sequence π_p for p violates p's specification φ_p when combined with at least one of the input sequences π_e or π'_e. Equivalently, if φ_p is satisfied by π_p combined with π_e, then it is violated when π_p is combined with π'_e. Observe that for p to behave differently in two scenarios, a difference must occur at a finite time t. Clearly, this will only happen if p's input shows a difference in finite time. To capture this, we say that a pair (π_e, π'_e) of environment output sequences is in the *time-bounded* distinguishability relation if the violation with π'_e is guaranteed to happen in finite time. In order to avoid this violation, process p must act in finite time, before the violation occurs on π'_e. We say that a trace π *finitely violates* an LTL formula φ, denoted by $\pi \models_f \varphi$, if there exists a finite prefix w of π such that every (infinite) trace extending w violates φ.

Definition 3 (Time-bounded distinguishability). *Given a local specification φ_p for process p, the* time-bounded distinguishability relation Λ_p *is the set of pairs $(\pi_e, \pi'_e) \in (2^{O_e})^\omega \times (2^{O_e})^\omega$ of traces of global inputs such that every trace of local outputs $\pi_p \in O_p$ either violates the specification φ_p when combined*

with π_e, or finitely violates p's local specification φ_p when combined with π'_e:

$$\begin{aligned}\Lambda_p = \{(\pi_e, \pi'_e) \in (2^{O_e})^\omega \times (2^{O_e})^\omega \mid \\ \forall \pi_p \in (2^{O_p})^\omega.\ \textit{if}\ \ \pi_e \sqcup \pi_p \vDash \varphi_p \ \ \textit{then}\ \ \pi'_e \sqcup \pi_p \nvDash_f \varphi_p\ \}\end{aligned}$$

Note that, unlike the distinguishability relation Δ_p, the *time-bounded* distinguishability relation Λ_p is not symmetric: For (π_e, π'_e), the trace $\pi'_e \sqcup \pi_p$ has to finitely violate φ_p, while the trace $\pi_e \sqcup \pi_p$ only needs to violate φ_p in the infinite evaluation. As a result, the corresponding *time-bounded* information flow assumption will also be asymmetric: we require that on input π_e, process p eventually obtains the knowledge that the input is different from π'_e. For input π'_e we do not pose such a requirement. The intuition behind this definition is that on environment output π'_e, process p must definitely produce some output that does *not* finitely violate φ_p. This output can safely be produced without ever knowing that the input is π'_e. However, on input π_e, it becomes necessary for process p to eventually deviate from the output that would work for π'_e. In order to safely do so, p needs to realize after some finite time that the input is not π'_e. In our running example, π_e would be an input in which `in` is initially *true*, while π'_e will be one in which it starts out being *false*.

Suppose we have a function $t : (2^{O_e})^\omega \to \mathbb{N}$ that identifies, for each environment output π_e, the time $t(\pi_e)$ by which process p is guaranteed to know that the environment output is not π'_e. We define the information flow assumption for this particular function t as a 2-hyperproperty. Since we do not know t in advance, the time-bounded information flow assumption is the (infinite) union of all 2-hyperproperties corresponding to the different possible functions t.

Definition 4 (Time-bounded information flow assumption). *Given the time-bounded distinguishability relation Λ_p for process p, the* time-bounded information flow assumption *χ_p for p is the (infinite) union over the 2-hyperproperties induced by the following relations R_t, for all possible functions $t : (2^{O_e})^\omega \to \mathbb{N}$:*

$$\begin{aligned}R_t = \{(\pi, \pi') \in (2^{\mathcal{V}})^\omega \times (2^{\mathcal{V}})^\omega \mid \\ \textit{if}\ (\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \in \Lambda_p,\ \textit{then}\ \pi[0..t(\pi{\downarrow_{O_e}})]{\downarrow_{I_p}} \neq \pi'[0..t(\pi{\downarrow_{O_e}})]{\downarrow_{I_p}}\}\end{aligned}$$

Unlike the information flow assumption (cf. Theorem D.1), the *time-bounded* information flow assumption is not in general a necessary assumption. Consider a modification of our motivating example, where there is an additional environment output `start`, which is only visible to process a, not to process b. The previous specification φ_b is modified so that if `in` is *true* initially, then `out` must be *true* two steps after `start` becomes *true* for the first time; if `in` is *false* initially, then `out` must become *false* after two positions have passed since the first time `start` has become *true*. The specification φ_a ensures that the channel `c` is set to *true* until `start` becomes *true*. Clearly, this is realizable: if `in` is *false* initially, process a sets `c` to *false* once `start` becomes *true*, otherwise `c` stays *true* forever. Process b starts by setting `out` to *true*. It then waits for `c` to become *false*, and, if and when that happens, sets `out` to *false*. In this way, process b accomplishes

the correct reaction within two steps after `start` has occurred. However, the function t required by the time-bounded information flow assumption does not exist, because the time of the communication depends on the environment: the prefix needed to distinguish an environment output π_e, where `in` is *true* initially from an environment output π'_e, where `in` is *false* initially, depends on the time when `start` becomes *true* on π'_e.

We now characterize a set of situations in which the time-bounded information flow requirement is still a necessary requirement. For this purpose we consider time-bounded distinguishability relations where the safety violation occurs after a bounded number of steps. We call such time-bounded distinguishability relations *uniform*; the formal definition follows below.

Definition 5 (Uniform distinguishability). *A time-bounded distinguishability relation Λ_p is* uniform *if for every trace $\pi_e \in (2^{O_e})^\omega$ of global inputs, and every trace $\pi_p \in (2^{O_p})^\omega$ of local outputs of p, there exists a natural number $n \in N$ such that for all $\pi'_e \in (2^{O_e})^\omega$ s.t. $(\pi_e, \pi'_e) \in \Lambda_p$ if $\pi_e \sqcup \pi_p \vDash \varphi_p$ then $\pi'_e \sqcup \pi_p \nvDash_n \varphi_p$.*

Theorem 2. *Let Λ_p be a uniform time-bounded distinguishability relation derived from process p's local specification φ_p. Every computation tree that satisfies φ_p also satisfies the time-bounded information flow assumption χ_p.*

The proof of Theorem 2 can be found in Appendix D.1. The relations presented in this section as well as the uniformity check can be represented by and verified with automata. A detailed explanation of the automata and their constructions can be found in Appendix C.

5 Compositional Synthesis

We now use the time-bounded information flow assumptions to split the distributed synthesis problem for an architecture $(I_p, I_q, O_p, O_q, I_e)$ into two separate synthesis problems. The local implementations are then composed and form a correct system, whose decomposition returns the solution for each process.

5.1 Constructing the Hyper Implementations

We begin with the synthesis of local processes. Let Λ_p and Λ_q be the time-bounded distinguishability relations for p and q, and let χ_p and χ_q be the resulting time-bounded information flow assumptions. In the individual synthesis problems, we ensure that process p provides the information needed by process q, i.e., that the implementation of p satisfies χ_q, and, similarly, that q provides the information needed by p, i.e., q's implementation satisfies χ_p.

We carry out the individual synthesis of a process implementation on trees that branch according to the input of the process (including $\mathtt{t}_p$) *and* the environment's output. In such a tree, the synthesized process thus has access to full information. We call this tree a *hyper implementation*, rather than an implementation, because the hyper implementation describes how the process will react to

certain information, without specifying *how* the process will receive information. This detail is left open until we know the other process' hyper implementation: at that point, both hyper implementations can be turned into standard strategies, which are trees that branch according to the process' own inputs.

Definition 6 (Hyper implementation). *Let p and q be processes and e be the environment. A $2^{O_e \cup I_p \cup \{\mathtt{t}_p\}}$-branching $2^{O_p \cup \{\mathtt{t}_q\}}$-labeled tree h_p is a* hyper implementation *of p.*

Since the hyper implementation has access to the full information, while the time-bounded information flow assumption only guarantees that the relevant information arrives after some bounded time, the strategy has "too much" information. We compensate for this by introducing a *locality condition*: on two traces $(\pi_e, \pi'_e) \in \Lambda_p$ in the distinguishability relation of process p, as long as the input to the process from the external environment is identical, process p's output must be identical until $\mathtt{t}_p$ happens (which signals that the bound for the transmission of the information has been reached). For traces $(\pi_e, \pi'_e) \notin \Lambda_p$ outside the distinguishability relation, process p's output must be identical until there is a difference in the input to process p or in the value of $\mathtt{t}_p$.

Definition 7 (Locality condition). *Given the time-bounded distinguishability relation Λ_p for process p, the* locality condition η_p *for p is the 2-hyperproperty induced by the following relation R:*

$$\begin{aligned} R = \{(\pi, \pi') \in (2^{O_e \cup I_p \cup \{t_p\}})^\omega \times (2^{O_e \cup I_p \cup \{t_p\}})^\omega \mid \\ \text{if } (\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \in \Lambda_p, \text{ then } \pi[0..t]{\downarrow_{O_p}} = \pi'[0..t]{\downarrow_{O_p}} \text{ and} \\ \text{if } (\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \notin \Lambda_p, \text{ then } \pi[0..t']{\downarrow_{O_p}} = \pi'[0..t']{\downarrow_{O_p}}\} \end{aligned}$$

where t is the smallest natural number such that $\mathtt{t}_p \in \pi[0..t]$ or $\pi[0..t] \downarrow_{I_p} \neq \pi'[t] \downarrow_{I_p}$ (and ∞ if no such t exists), and t' is the smallest natural number such that $\pi[0..t'] \downarrow_{I_p} \neq \pi'[0..t'] \downarrow_{I_p}$ or $\pi[0..t'] \downarrow_{\{\mathtt{t}_p\}} \neq \pi'[0..t'] \downarrow_{\{\mathtt{t}_p\}}$ (and ∞ if no such t' exists).

We use HyperLTL to formulate the locality condition for process b in our running example. Based on the time-bounded distinguishability relation Λ_b, which relates every trace with $\mathtt{in}$ in the first step to all traces on which $\neg\mathtt{in}$ holds there, we can write the locality condition:

$$\begin{aligned} \forall \pi, \pi'. (\mathtt{in}_\pi \wedge \neg \mathtt{in}_{\pi'}) \rightarrow ((\mathtt{t}_\pi \vee \mathtt{c}_\pi \nleftrightarrow \mathtt{c}_{\pi'}) \, \mathcal{R} (\mathtt{out}_\pi \leftrightarrow \mathtt{out}_{\pi'})) \\ \wedge (\neg(\mathtt{in}_\pi \wedge \neg \mathtt{in}_{\pi'})) \rightarrow (\mathtt{t}_\pi \nleftrightarrow \mathtt{t}_{\pi'} \vee \mathtt{c}_\pi \nleftrightarrow \mathtt{c}_{\pi'}) \, \mathcal{R} (\mathtt{out}_\pi \leftrightarrow \mathtt{out}_{\pi'})) \end{aligned}$$

The order in the formula is analogous to the order in Definition 7. For all pairs of traces that are in the distinguishability relation, i.e., $\mathtt{in}$ is *true* on π and *false* on π', the outputs being equivalent on both traces can only be released by $\mathtt{t}$ on trace π or by a difference in the local inputs ($\mathtt{c}$). Moreover, if the traces are not in the distinguishability relation, i.e., $\neg(\mathtt{in}_\pi \wedge \neg \mathtt{in}_{\pi'})$, then only a difference in $\mathtt{t}$ or $\mathtt{c}$ can release $\mathtt{out}$ to be equivalent on both traces. With the locality condition at hand, we define when a hyper implementation is locally correct:

Definition 8 (Local correctness of hyper implementations). *Let p and q be processes, let φ_p be the local specification of p, let η_p be its locality condition, and let χ_q be the information flow assumption of q. The hyper implementation h_p of p is* locally correct *if it satisfies φ_p, η_p, and χ_q.*

The specification φ_p is a trace property, while η_p and χ_q are hyperproperties. Since all properties that need to be satisfied by the process are guarantees, it is not necessary to assume explicit behaviour of process q to realize process p. Local correctness relies on the guarantee that the other process satisfies the current process' own information flow assumption. Note that both the locality condition and the information flow assumption for p build on the time-bounded distinguishability relation of p.

5.2 Composition of Hyper Implementations

The hyper implementations of each of the processes are locally correct and satisfy the information flow assumptions of the other process respectively. However, the hyper implementations have full information of the inputs and are dependent on the additional variables $\mathtt{t}_p$ and $\mathtt{t}_q$. To construct practically executable local implementations, we first compose the hyper implementations into one strategy.

Definition 9 (Composition of hyper implementations). *Let p and q be two processes with hyper implementations given as infinite $2^{O_e \cup I_p \cup \{\mathtt{t}_p\} \cup IC_p}$-branching $2^{O_p \cup \{\mathtt{t}_q\}}$-labeled tree h_p for process p, and an infinite $2^{O_e \cup I_q \cup \{\mathtt{t}_q\} \cup IC_q}$-branching $2^{O_p \cup \{\mathtt{t}_p\}}$-labeled tree h_q for process q.*

Given two hyper implementations h_p and h_q, we define the composition $h = h_p || h_q$ to be a 2^{O_e}-branching $2^{O_p \cup O_q}$-labeled tree, where $h(v) = (h_p(f_p(v)) \cup h_q(f_q(v))) \cap (O_p \cup O_q)$ and f_p, f_q are defined as follows:

$$f_p(\epsilon) = \epsilon \qquad f_p(v \cdot x) = f_p(v) \cdot ((x \cap I_p) \cup (h_q(f_q(v)) \cap (I_p \cup \{\mathtt{t}_p\}))$$
$$f_q(\epsilon) = \epsilon \qquad f_q(v \cdot x) = f_q(v) \cdot ((x \cap I_q) \cup (h_p(f_p(v)) \cap (I_q \cup \{\mathtt{t}_q\}))$$

If each hyper implementation satisfies the time-bounded information flow assumption of the other process, then there exists a strategy for each process (given as a tree that branches according to the local inputs of the process), such that the combined behavior of the two strategies corresponds exactly to the composition of the hyper implementations.

The composition of the hyper implementations of the bit transmission protocol is shown in Figure 2. The initial state is the combination of both processes initial states with the corresponding outputs. We change the state after the value of `in` is received. While process a directly reacts to `in`, process b cannot observe its value, and the composition can either be in h_0^b or h_1^b. Booth states have the same output. In the next step, process a communicates the value of `in` by setting `c` to true or false, such that the loop states h_1^a, h_1^a and h_2^a, h_3^b are reached.

The local strategies of the processes are constructed from the composed hyper implementations. As an auxiliary notion we introduce the *knowledge set*: the set of finite traces in the composition that cannot be distinguished by a process.

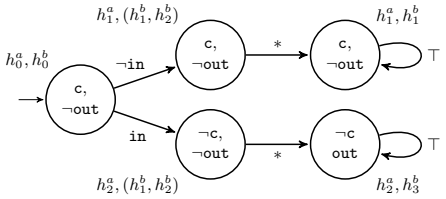

Fig. 2: The composition of the hyper implementations of a in Figure 1c and b in Figure 1d. The states are labeled with the combination of states reached for both processes, and multiple, if they cannot be distinguished.

Definition 10 (Knowledge set). *Let p and q be two processes with composed hyper implementations $h = h_p||h_q$. For a finite trace $v \in (2^{I_p})^*$ of inputs to p, we define the knowledge set $K_p(v)$ to be*

$$K_p(v) \triangleq \{w \mid w \text{ is a finite trace of } (2^{O_e})^* \text{ and } f_p(w) = v\}.$$

Lemma 1. *For all s $v, v' \in (2^{I_p})^*$, if $K_p(v) = K_p(v')$ then $h(v) \downarrow_{O_p} = h(v') \downarrow_{O_p}$.*

The proof of Lemma 1 can be found in Appendix D.2. The local strategies from the composed hyper implementations are then defined as follows:

Definition 11 (Local strategies from hyper implementations). *Let p and q be two processes with time-bounded information flow assumptions χ_p and χ_q, and $h = h_p||h_q$ be the composition of their hyper implementations. For $j \in \{p, q\}$ the strategy s_j, represented as a 2^{I_j}-branching 2^{O_j}-labeled tree for process j, is defined as follows:*

$$s_j(\epsilon) = \epsilon \qquad s_j(v) = \begin{cases} \emptyset & \text{if } |K_j(v)| = 0 \\ h(min(K_j(v))) \downarrow_{O_j} & \text{if } |K_j(v)| > 0 \end{cases}$$

where $min(K_j(v))$ is the smallest trace based on an arbitrary order over $K_j(v)$.

The base case of the definition inserts a label for unreachable traces in the composed hyper implementation. For example, the local inputs $I_p \backslash O_e$ are determined by s_q, and not all input words in $(2^{I_q})^*$ are possible. Process p's local strategy s_p can discard these input words. The second case of the definition picks the smallest trace in the knowledge set and computes the outputs from h that are local to a process. Intuitively, the outputs of h have to be the same for every trace that a process considers possible in the composed hyper implementations. We therefore pick one of them, compute the output of the composed hyper-strategy, and restrict the output to the local outputs of the process. The following theorem states the correctness of the construction in Definition 11.

Theorem 3. *Let p and q be two processes with time-bounded information flow assumptions χ_p and χ_q, let $h = h_p||h_q$ be the composition of the hyper implementations, and s_p and s_q be the local strategies. Then, for all $v \in (2^{O_e})^*$ it holds that $h(v) = s_p(g_p(v)) \cup s_q(g_q(v))$ where g_p, g_q are defined as follows:*

$$g_p(\epsilon) = \epsilon \qquad g_p(v \cdot x) = g_p(v) \cdot ((x \cap I_p) \cup (s_q(g_q(v)) \cap I_p)$$
$$g_q(\epsilon) = \epsilon \qquad g_q(v \cdot x) = g_q(v) \cdot ((x \cap I_q) \cup (s_p(g_p(v)) \cap I_q)$$

The proof is inductive over the words $v \in (2^{O_e})^*$ and can be found in Appendix D.2 Combining all definitions and theorems of the previous sections, we conclude with the following corollary.

Corollary 1. *Let $(I_p, I_q, O_p, O_q, I_e)$ be an architecture and $\varphi = \varphi_p \wedge \varphi_q$ be a specification. If the hyper-strategies h_p and h_q are locally correct, then the implementation (s_p, s_q) satisfies φ.*

6 A More Practical Approach

A major disadvantage of the synthesis approach of the preceding sections is that the hyper implementations are based on the full set of environment outputs; as a result, hyper implementations branch according to inputs that are not actually available; this, in turn, necessitates the introduction of the locality condition.

In this section, we develop a more practical approach, where the branching is limited to the information that is actually available to the process: this includes any environment output directly visible to the process and, additionally, the information the process is guaranteed to receive according to the information flow assumption. As a result, the synthesis of the process is sound without need for a locality condition. We develop this approach under two assumptions: First, we assume that the time-bounded information flow assumption only depends on environment outputs the sending process can actually see; second, we assume that the time-bounded information flow assumption can be decomposed into a finite set of classes in the following sense: For a trace π of environment outputs, the information class $[\pi]_p$ describes that, on the trace π, the process p eventually needs to become aware that the current trace is in the set $[\pi]$. The information class is obtained by collecting all traces that are *not* related to π in the time-bounded distinguishability relation.

Definition 12 (Information classes). *Given a time-bounded distinguishability relation Λ_p for process p, the* information class $[\pi]_p$ *of a trace π over O_e is the following set of traces:* $[\pi]_p = (2^{O_e})^\omega \setminus \{\pi' \in (2^{O_e})^\omega \mid (\pi, \pi') \in \Lambda_p\}$

The next definition relativizes the specification of the processes for a particular information class, reflecting the fact that the process does not know the actual environment output, but only its information class; hence, the process output needs to be correct for all environment outputs in the information class.

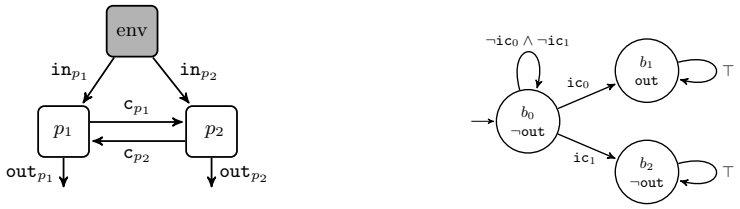

(a) (b)

Fig. 3: The architecture used for our experiments in (a) where the number outputs, inputs, and communication channels can vary. Figure 3b shows the implementation of process b for its bit transmission component specification.

Definition 13 (Relativized specification). *For a process p with specification φ_p and an information class c, the relativized specification $\varphi_{p,c}$ is the following trace property over $(I_p \cap O_e) \cup O_p$:*

$$\varphi_{p,c} = \{\pi_e \sqcup \pi_p \mid \pi_e \in (2^{I_p \cap O_e})^\omega, \pi_p \in (2^{O_p})^\omega \text{ s.t. } \forall \pi'_e \in c.\ \pi'_e \sqcup \pi_p \vDash \varphi_p\}$$

The component specification, which is the basis for the synthesis of the process, must take into account that the process does not know the information class in advance; the behavior of the other process will only eventually reveal the information class. Let IC be the set of information classes for process p. Assume that this set is finite. We now replace the inputs of the process that come from the other process with new input channels IC as new inputs. In the hyper implementation, receiving such an input reveals the information class to the process. In the actual implementation, the information class will be revealed by the actual outputs of the other process that are observable for p. The component specification requires that the processes satisfy the relativized specification under the assumption that the information class is eventually received. We encode this assumption as a trace condition ψ, which requires that exactly one of the elements of IC eventually occurs.

Definition 14 (Component specification). *For process p with specification φ_p, the component specification $\langle\varphi_p\rangle$ over $(I_p \cap O_e) \cup IC \cup O_p$ is defined as*

$$\langle\varphi_p\rangle = \{\pi \in (2^{(I_p \cap O_e) \cup IC \cup O_p})^\omega \mid \text{ if } \pi \vDash \psi \text{ then } \pi \vDash \bigwedge_{c \in IC} (\Diamond c \rightarrow \varphi_{p,c}\}$$

where ψ is the following trace property over $(I_p \cap O_e) \cup IC \cup O_p$:

$$\psi = \{\pi \in (2^{(I_p \cap O_e) \cup IC \cup O_p})^\omega \mid \exists \pi' \in (2^{O_e})^\omega.\ \pi \downarrow_{I_p \cap O_e} = \pi' \downarrow_{I_p \cap O_e} \text{ and } \pi \vDash \Diamond[\pi'] \text{ and exactly one element of } IC \text{ occurs on } \pi\}$$

The component specification allows us to replace the locality condition (Def. 7), which is a hyperproperty, with a trace property. Note, however, that the process additionally needs to satisfy the information flow assumption of the other

process, which may in general depend on the full set O_e of environment outputs. This would require us to synthesize the process on the full set O_e, and to re-introduce the locality condition. In practice, however, the information flow assumption of one process often only depends on the information of the other process. In this case, it suffices to synthesize each process based only on the locally visible environment outputs.

Figure 3b shows the implementation of b for its component specification $\langle\varphi_b\rangle$. In contrast to its hyper implementation (cf. Figure 1b), it does not branch according to $\texttt{in}$ and $\texttt{t}_p$, but only variables in IC. The specification is encoded as the following LTL formula:

$$\langle\varphi_b\rangle = \big(\square\neg\texttt{ic}_0 \vee \square\neg\texttt{ic}_1\big) \wedge \diamondsuit\big((\texttt{ic}_0 \vee \texttt{ic}_1)\big) \rightarrow \big((\diamondsuit\,\texttt{ic}_0 \rightarrow \diamondsuit\,\texttt{out}) \wedge (\diamondsuit\,\texttt{ic}_1 \rightarrow \square\neg\texttt{out})\big)$$

The left hand side of the implication represents the assumption ψ, the right hand side specifies the guarantee for each information class. The composition and decomposition can be performed analogously to the hyper implementations, where we map the value of $\texttt{ic}$ to the values of the communication variables. We construct the automata for component specifications in Appendix D.3.

7 Experiments

The focus of our experiments is on the performance of the compositional synthesis approach compared to non-compositional synthesis methods for distributed systems. While the time-bounded information flow assumptions and the component specification can be computed automatically by automata constructions, we have, for the purpose of these experiments, built them manually and encoded them as formulas in HyperLTL or LTL, which were then entered to the BoSy/BoSyHyper [13] synthesis tool. Our experiments are based on the following benchmarks:

- **AC.** *Atomic commit.* The atomic commitment protocol specifies that the output of a local process is set to true iff the observable input and the unobservable inputs are true as well. We only consider one round of communication, the initial input determines all values. The parameter shows how many input variables each process receives, Par. = 1 for the running example.
- **EC.** *Eventual commit.* The atomic commit benchmark extended to eventual inputs - if all inputs (independently of each other) eventually will be true, then there needs to be information flow.
- **SA.** *Send all.* Every input of the sender is relevant for the receiver, so it will eventually be sent if it it set to true. The parameter represents the number of input values and therefore the number of information classes.

Table 1 shows the performance of the compositional synthesis approach. The column architecture (Arch.) signalizes for each benchmark if the information flow is directional (dir.) or bidirectional (bidir.). Column (Inflow send) indicates the running time for the sending process; where applicable, column (Inflow rec.) indicates the running time for the synthesis of the process that only receives information. We compare the compositional approach to BoSyHyper, based

Table 1: The results of the experiments with execution times given in seconds. The cell is highlighted if it was faster than the other approaches, where the sum of sender and receiver is taken as reference.

Bench.	Arch.	Par.	Inflow send.	Inflow rec.	Distr.BoSy	Inc. BoSy
AC	dir	1	0.92	0.70	**1.41**	2.31
	dir	2	**0.36**	**1.28**	2.86	2.30
	dir	3	**0.92**	**0.68**	2.46	2.55
	dir	4	**0.92**	**0.79**	720.60	3.41
	dir	5	**0.92**	**0.68**	TO	9.27
	bidir	1	1.45	-	**0.96**	9.27
	bidir	2	**2.49**	-	TO	TO
	bidir	3	**79.18**	-	TO	TO
	bidir	4	TO	-	TO	TO
EC	dir	1	0.68	1.87	**0.92**	2.556
	dir	2	0.94	1.85	**0.96**	3.90
	dir	3	**202.09**	**1.78**	TO	TO
	dir	4	TO	TO	TO	TO
	bidir	1	**3.77**	-	4.63	147.46
	bidir	2	TO	-	TO	TO
SA	dir	1	1.31	0.92	2.21	**1.579**
	dir	2	**1.78**	**0.92**	27.47	TO
	dir	3	TO	1.08	TO	TO

on a standard encoding of distributed synthesis in HyperLTL (Inc. BoSy), and a specialized tool for distributed synthesis [2] (Distr. BoSy). All experiments were performed on a MacBook Pro with a 2,8 GHz Intel Quad Core processor and 16 GB of RAM. The timeout was 30 minutes.

Information flow guided synthesis outperforms the standard approaches, especially for more complex components. For example, in the atomic commitment benchmark, scaling in the number of inputs does not impact the synthesis of the local processes, while Distr. BoSy eventually times out, and the running time of Inc. BoSy increases faster than for the information flow synthesis. For all approaches, the Send All benchmark is the hardest one to solve. Here, each input that will eventually be set needs to be eventually sent, which leads to non-trivial communication over the shared variables and an increased state space to memorize the individual inputs. Nevertheless, the information flow guided synthesis outperforms the other approaches and times out with parameter 3 because BoSyHyper cannot cope with the number of states needed. Synthesizing a receiver that does not satisfy an information flow assumption is close to irrelevant for every benchmark run. Since these processes are synthesized with local LTL specifications, scaling only in the number of local inputs or information that will eventually be received is easily possible. Notably, these receivers are compatible with any implementation of the sender, whereas the solutions of the other approaches are only compatible for the same synthesis run.

8 Conclusion

The approach of the paper provides the foundation for a new class of distributed synthesis algorithms, where the assumptions refer to the flow of information and are represented as hyperproperties. In many situations, necessary information flow assumptions exist even if there are no necessary behavioral assumptions. There are at least two major directions for future work. The first direction concerns the insight that compositional synthesis profits from the generality of hyperproperties; at the same time, synthesis from hyperproperties is much more challenging than synthesis from trace properties. To address this issue, we have introduced the more practical method in Section 6, which replaces locality, a hyperproperty, with the component specification, a trace property. However, this method is limited to information flow assumptions that refer to a finite amount of information. It is very common that the required amount of information is infinite in the sense that the same type of information must be transmitted again and again. We conjecture that our method can be extended to such situations.

A second major direction is the extension to distributed systems with more than two processes. The two-process case has the advantage that the assumptions of one process must be guaranteed by the other. With more than two processes, the localization of the assumptions becomes more difficult or even impossible, if multiple processes have (partial) access to the required information.

References

1. Alur, R., Moarref, S., Topcu, U.: Compositional synthesis of reactive controllers for multi-agent systems. In: Chaudhuri, S., Farzan, A. (eds.) Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9780, pp. 251–269. Springer (2016). https://doi.org/10.1007/978-3-319-41540-6_14, https://doi.org/10.1007/978-3-319-41540-6_14
2. Baumeister, J.E.: Encodings of Bounded Synthesis of Distributed Systems. B.Sc. Thesis, Saarland University (2017)
3. Bloem, R., Chatterjee, K., Jacobs, S., Könighofer, R.: Assume-guarantee synthesis for concurrent reactive programs with partial information. In: Baier, C., Tinelli, C. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings. Lecture Notes in Computer Science, vol. 9035, pp. 517–532. Springer (2015). https://doi.org/10.1007/978-3-662-46681-0_50, https://doi.org/10.1007/978-3-662-46681-0_50
4. Chatterjee, K., Henzinger, T.A.: Assume-guarantee synthesis. In: Grumberg, O., Huth, M. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 13th International Conference, TACAS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007 Braga, Portugal, March 24 - April 1, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4424, pp. 261–275. Springer (2007). https://doi.org/10.1007/978-3-540-71209-1_21, https://doi.org/10.1007/978-3-540-71209-1_21

5. Clarkson, M.R., Finkbeiner, B., Koleini, M., Micinski, K.K., Rabe, M.N., Sánchez, C.: Temporal logics for hyperproperties. In: Abadi, M., Kremer, S. (eds.) Principles of Security and Trust - Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings. Lecture Notes in Computer Science, vol. 8414, pp. 265–284. Springer (2014). https://doi.org/10.1007/978-3-642-54792-8_15, https://doi.org/10.1007/978-3-642-54792-8_15
6. Clarkson, M.R., Schneider, F.B.: Hyperproperties. Journal of Computer Security **18**(6), 1157–1210 (2010)
7. Damm, W., Finkbeiner, B.: Automatic Compositional Synthesis of Distributed Systems. In: Jones, C.B., Pihlajasaari, P., Sun, J. (eds.) FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8442, pp. 179–193. Springer (2014). https://doi.org/10.1007/978-3-319-06410-9_13
8. Damm, W., Finkbeiner, B.: Automatic compositional synthesis of distributed systems. In: Jones, C.B., Pihlajasaari, P., Sun, J. (eds.) FM 2014: Formal Methods - 19th International Symposium, Singapore, May 12-16, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8442, pp. 179–193. Springer (2014). https://doi.org/10.1007/978-3-319-06410-9_13, https://doi.org/10.1007/978-3-319-06410-9_13
9. Dimitrova, R., Finkbeiner, B., Kovács, M., Rabe, M.N., Seidl, H.: Model checking information flow in reactive systems. In: Kuncak, V., Rybalchenko, A. (eds.) Verification, Model Checking, and Abstract Interpretation - 13th International Conference, VMCAI 2012, Philadelphia, PA, USA, January 22-24, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7148, pp. 169–185. Springer (2012). https://doi.org/10.1007/978-3-642-27940-9_12, https://doi.org/10.1007/978-3-642-27940-9_12
10. Filiot, E., Jin, N., Raskin, J.: Compositional Algorithms for LTL Synthesis. In: Bouajjani, A., Chin, W. (eds.) Automated Technology for Verification and Analysis - 8th International Symposium, ATVA 2010, Singapore, September 21-24, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6252, pp. 112–127. Springer (2010). https://doi.org/10.1007/978-3-642-15643-4_10
11. Filiot, E., Jin, N., Raskin, J.: Compositional algorithms for LTL synthesis. In: Bouajjani, A., Chin, W. (eds.) Automated Technology for Verification and Analysis - 8th International Symposium, ATVA 2010, Singapore, September 21-24, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6252, pp. 112–127. Springer (2010). https://doi.org/10.1007/978-3-642-15643-4_10, https://doi.org/10.1007/978-3-642-15643-4_10
12. Finkbeiner, B., Schewe, S.: Uniform distributed synthesis. In: Proceedings of the 20th ACM/IEEE Symposium on Logic in Computer Science (LICS). pp. 321–330 (2005)
13. Finkbeiner, B., Hahn, C., Lukert, P., Stenger, M., Tentrup, L.: Synthesizing reactive systems from hyperproperties. In: Chockler, H., Weissenbacher, G. (eds.) Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10981, pp. 289–306. Springer (2018). https://doi.org/10.1007/978-3-319-96145-3_16, https://doi.org/10.1007/978-3-319-96145-3_16
14. Finkbeiner, B., Passing, N.: Dependency-Based Compositional Synthesis. In: Hung, D.V., Sokolsky, O. (eds.) Automated Technology for Verification and Analysis -

18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12302, pp. 447–463. Springer (2020). https://doi.org/10.1007/978-3-030-59152-6_25

15. Finkbeiner, B., Passing, N.: Dependency-based compositional synthesis. In: Hung, D.V., Sokolsky, O. (eds.) Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings. Lecture Notes in Computer Science, vol. 12302, pp. 447–463. Springer (2020). https://doi.org/10.1007/978-3-030-59152-6_25, https://doi.org/10.1007/978-3-030-59152-6_25
16. Finkbeiner, B., Passing, N.: Compositional synthesis of modular systems. In: Hou, Z., Ganesh, V. (eds.) Automated Technology for Verification and Analysis - 19th International Symposium, ATVA 2021, Gold Coast, QLD, Australia, October 18-22, 2021, Proceedings. Lecture Notes in Computer Science, vol. 12971, pp. 303–319. Springer (2021). https://doi.org/10.1007/978-3-030-88885-5_20, https://doi.org/10.1007/978-3-030-88885-5_20
17. Finkbeiner, B., Rabe, M.N., Sánchez, C.: Algorithms for model checking HyperLTL and HyperCTL *. In: Kroening, D., Pasareanu, C.S. (eds.) Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9206, pp. 30–48. Springer (2015). https://doi.org/10.1007/978-3-319-21690-4_3, https://doi.org/10.1007/978-3-319-21690-4_3
18. Hecking-Harbusch, J., Metzger, N.O.: Efficient trace encodings of bounded synthesis for asynchronous distributed systems. In: Chen, Y., Cheng, C., Esparza, J. (eds.) Automated Technology for Verification and Analysis - 17th International Symposium, ATVA 2019, Taipei, Taiwan, October 28-31, 2019, Proceedings. Lecture Notes in Computer Science, vol. 11781, pp. 369–386. Springer (2019). https://doi.org/10.1007/978-3-030-31784-3_22, https://doi.org/10.1007/978-3-030-31784-3_22
19. Kugler, H., Segall, I.: Compositional synthesis of reactive systems from live sequence chart specifications. In: Kowalewski, S., Philippou, A. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 15th International Conference, TACAS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5505, pp. 77–91. Springer (2009). https://doi.org/10.1007/978-3-642-00768-2_9
20. Kupferman, O., Vardi, M.Y.: Synthesizing distributed systems. In: Logic in Computer Science (LICS) (2001)
21. Kupferman, O., Piterman, N., Vardi, M.Y.: Safraless Compositional Synthesis. In: Ball, T., Jones, R.B. (eds.) Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4144, pp. 31–44. Springer (2006). https://doi.org/10.1007/11817963_6
22. Kupferman, O., Piterman, N., Vardi, M.Y.: Safraless compositional synthesis. In: Ball, T., Jones, R.B. (eds.) Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4144, pp. 31–44. Springer (2006). https://doi.org/10.1007/11817963_6, https://doi.org/10.1007/11817963_6
23. Kupferman, O., Vardi, M.: Model checking of safety properties. Formal Methods in System Design **19** (09 1999). https://doi.org/10.1023/A:1011254632723

24. Majumdar, R., Mallik, K., Schmuck, A., Zufferey, D.: Assume-guarantee distributed synthesis. IEEE Trans. Comput. Aided Des. Integr. Circuits Syst. **39**(11), 3215–3226 (2020). https://doi.org/10.1109/TCAD.2020.3012641, https://doi.org/10.1109/TCAD.2020.3012641
25. Pnueli, A.: The temporal logic of programs. In: 18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977. pp. 46–57. IEEE Computer Society (1977). https://doi.org/10.1109/SFCS.1977.32, https://doi.org/10.1109/SFCS.1977.32
26. Pnueli, A., Rosner, R.: Distributed Reactive Systems Are Hard to Synthesize. In: 31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume II. pp. 746–757. IEEE Computer Society (1990). https://doi.org/10.1109/FSCS.1990.89597
27. Schewe, S., Finkbeiner, B.: Semi-automatic distributed synthesis. Int. J. Found. Comput. Sci. **18**(1), 113–138 (2007)
28. Wolper, P., Vardi, M.Y., Sistla, A.P.: Reasoning about infinite computation paths (extended abstract). In: 24th Annual Symposium on Foundations of Computer Science, Tucson, Arizona, USA, 7-9 November 1983. pp. 185–194. IEEE Computer Society (1983). https://doi.org/10.1109/SFCS.1983.51, https://doi.org/10.1109/SFCS.1983.51
29. Yasuoka, H., Terauchi, T.: Quantitative information flow - verification hardness and possibilities. In: Proceedings of the 23rd IEEE Computer Security Foundations Symposium, CSF 2010, Edinburgh, United Kingdom, July 17-19, 2010. pp. 15–27. IEEE Computer Society (2010). https://doi.org/10.1109/CSF.2010.9, https://doi.org/10.1109/CSF.2010.9

A Automata Prelimnaries

A *nondeterministic finite automaton* (NFA) is a tuple $\mathcal{A} = (Q, \Sigma, q_0, F, \delta)$, where Q denotes a finite set of states, Σ is a finite alphabet, q_0 is a designated initial state, $F \subseteq Q$ is the set of accepting states, and $\delta : Q \times \Sigma \to \mathcal{P}(Q)$ is the transition relation that maps a state and a letter to a set of possible successor states. A run of $\mathcal{A}$ on a finite word $w = w_0 \dots w_n \in \Sigma^*$ is a sequence of states $r = q_0 \dots q_{n+1} \in Q^*$ with $q_{i+1} \in \delta(q_i, w_i)$ for all $0 \leq i \leq n$. The run r is accepting if $q_{n+1} \in F$. The set of all accepted words by an automaton $\mathcal{A}$ is called its *language*, denoted by $\mathcal{L}(\mathcal{A})$.

A *Büchi automaton* $\mathcal{B} = (Q, \Sigma, q_0, F, \Delta)$ is an automaton over infinite words. A run of $\mathcal{B}$ on an infinite word $w = w_1 w_2 \dots \in \Sigma^\omega$ is an infinite sequence $r = q_0 q_1 \dots \in Q^\omega$ with $q_{i+1} \in \delta(q_i, w_i)$ for all $i \in \mathbb{N}$. A run r is accepting if there exist infinitely many $i \in \mathbb{N}$ such that $q_i \in F$. We use a Büchi automaton $\mathcal{A}$ over the alphabet Σ^2 to represent the 2-hyperproperty $R' \subseteq (\Sigma^2)^\omega$ with $\mathcal{L}(\mathcal{A}) = R'$.

B Proofs of Section 4.1

Theorem. *Let Λ_p be a uniform time-bounded distinguishability relation derived from process p's local specification φ_p. Every computation tree that satisfies φ_p also satisfies the time-bounded information flow assumption χ_p.*

Proof. Let (s_a, s_b) be an implementation that satisfies φ_p. We show that the time-bounded information-flow assumption χ_p is satisfied by defining a function $t : (2^{O_e})^\omega \to \mathbb{N}$ such that the 2-hyperproperty given by R_t is satisfied. To compute $t(\pi'_e)$ for some trace of inputs $\pi'_e \in (2^{O_e})^\omega$, we consider the trace of outputs $\pi'_p \in (2^{O_p})^\omega$ obtained by applying the implementation to the prefixes of π'_e. Since Λ_p is uniform, there is a natural number $n \in \mathbb{N}$ such that for all π_e with $(\pi_e, \pi'_e) \in \Lambda_p$, we have that $\pi'_e \sqcup \pi_p \not\models_n \varphi_p$. We set $t(\pi_e)$ to n.

To convince yourself that χ_p is satisfied, suppose, by way of contradiction, that R_t is violated on some pair $(\pi_e, \pi'_e) \in \Lambda_p$ of input traces, i.e., the projection on I_p is the same for π_e and π'_e on the entire prefix of length $t(\pi'_e)$. But then, also the output of process p must be the same along the entire prefix; this, however, means that π'_e will violate φ_p after $n = t(\pi'_e)$ steps, contradicting our assumption that the implementation satisfies φ_p. □

C Computing Information Flow Assumptions

C.1 Automata for Information Flow Assumptions

We first give an explicit construction of an automaton that recognizes the information flow assumption ψ_p that is induced by φ_p. The local specification φ_p is given as an LTL formula, which can be translated into an equivalent Büchi automaton $\mathcal{A}_\varphi$ over alphabet $2^{O_e \cup O_p}$ [28]. We self-compose $\mathcal{A}_\varphi$ into an automaton $\mathcal{B}$ over the alphabet $2^{O_e \cup O_p} \times 2^{O_e \cup O_p}$ such that $\mathcal{B}$ accepts a sequence of pairs iff both the projection on the first components and the projection on the second components are accepted by $\mathcal{A}_\varphi$ and, additionally, both components always agree on the values of O_p. We then construct a Büchi automaton $\mathcal{C}$ over the alphabet $2^{O_e} \times 2^{O_e}$ that guesses the values of O_p nondeterministically so that a pair of sequences is accepted by $\mathcal{C}$ iff there exists a valuation of O_p such that the extended sequences are accepted by $\mathcal{B}$. The automaton $\mathcal{C}$ thus accepts all sequences of global inputs that process p does *not* need to distinguish, because there is a sequence of outputs that satisfies the specification in both cases. We construct another Büchi automaton $\mathcal{D}$ over alphabet $2^{I_p} \times 2^{I_p}$ that recognizes a sequence of pairs of local input values iff they differ at some point. Finally, we construct a Büchi automaton $\mathcal{E}$ over the alphabet $2^{O_e \cup I_p} \times 2^{O_e \cup I_p}$ that accepts a sequence of pairs iff the sequence of projections on O_e is accepted by $\mathcal{C}$ or the sequence of projections on I_p is accepted by $\mathcal{D}$. The automaton $\mathcal{E}$ recognizes the information flow assumption ψ_p of process p.

Theorem 4. *For a process p with local specification φ_p, there exists a Büchi automaton with an exponential number of states in the length of φ_p that recognizes the information flow assumption ψ_p induced by φ_p.*

Proof. The automaton $\mathcal{E}$ described above recognizes ψ_p. We now claim that it has the stated size. The number of states of $\mathcal{A}_\varphi$ is exponential in the length of φ_p. By construction, the number of states of $\mathcal{B}$ is quadratic in the number of states of A_φ, and $\mathcal{C}$ has the same number of states as $\mathcal{B}$. The automaton $\mathcal{D}$

needs only two states. Hence, $\mathcal{E}$ has only two more states than $\mathcal{B}$ and so its total number of states is exponential in the length of φ_p, as claimed. □

C.2 Checking Uniformity

We begin with the construction of an automaton $\mathcal{A}_{\Lambda_p}$ over alphabet $2^{O_e} \times 2^{O_e}$ that recognizes the time-bounded distinguishability relation Λ_p. Let $\mathcal{A}_{\neg\varphi_p}$ be a deterministic ω-automaton over alphabet $2^{O_e \cup O_p}$ that recognizes all traces that violate the local specification φ_p. Let $\mathcal{B}_{\neg\varphi_p}$ be a deterministic finite-word automaton over alphabet $2^{O_e \cup O_p}$ that recognizes the bad prefixes of φ_p. We combine $\mathcal{A}_{\neg\varphi_p}$ and $\mathcal{B}_{\neg\varphi_p}$ into a deterministic ω-automaton $\mathcal{C}$ over alphabet $2^{O_e} \times 2^{O_e} \times 2^{O_p}$ that accepts traces of two inputs π_e, π'_e and an output π_p such that $\pi_e \sqcup \pi_p$ violates φ_p or $\pi'_e \sqcup \pi_e$ *finitely* violates φ_p. We obtain the universal automaton $\mathcal{D}_{\Lambda_p}$ with alphabet $2^{O_e} \times 2^{O_e}$ as the universal projection of $\mathcal{C}$ with respect to the outputs π_p.

Theorem 5. *For a process p with local specification φ_p, there exists a universal ω-automaton $\mathcal{A}_{\Lambda_p}$ over alphabet $2^{O_e} \times 2^{O_e}$ that recognizes the time-bounded distinguishability relation Λ_p. The number of states of $\mathcal{A}_{\Lambda_p}$ is doubly-exponential in the length of φ_p.*

Proof. Both $\mathcal{A}_{\neg\varphi_p}$ and $\mathcal{B}_{\neg\varphi_p}$ have doubly-exponentially many states in the length of φ_p [23]. The size of $\mathcal{C}$ is the product of the sizes of $\mathcal{A}_{\neg\varphi_p}$ and $\mathcal{B}_{\neg\varphi_p}$. Because of the universal projection, $\mathcal{D}_{\Lambda_p}$ is universal, rather than deterministic, but still of doubly-exponential size. □

Next, we check whether the time-bounded distinguishability relation is uniform. We construct an automaton that recognizes all traces of inputs and local outputs where no uniform bound exists. Let $\mathcal{A}_{\varphi_p}$ be a universal ω-automaton over alphabet $2^{O_e \cup O_p}$ that recognizes all traces that satisfy the local specification φ_p. We combine $\mathcal{A}_{\varphi}$ with $\mathcal{D}_{\Lambda_p}$ to a universal ω-automaton $\mathcal{E}$ over alphabet $2^{O_e} \times 2^{O_e} \times 2^{O_p}$ that accepts traces of two inputs π_e, π'_e and an output π_p when $(\pi_e, \pi'_e) \in \Lambda_p$ and $\pi_e \sqcup \pi_p \vDash \varphi_p$. From $\mathcal{E}$ we construct a universal automaton $\mathcal{F}$ over alphabet $2^f \times 2^{O_e} \times 2^{O_p}$ that accepts π_e and π_p if there exists an π'_e such that the bad prefix is reached on π'_e after f becomes *true* for the first time. Finally, we obtain a universal automaton $\mathcal{G}$ over alphabet $2^{O_e} \times 2^{O_p}$ that accepts those π_e and π_p that are accepted by $\mathcal{F}$ for all traces of f that set f to *true* at least once.

Theorem 6. *For a process p with local specification φ_p, whether the time-bounded distinguishability relation is uniform can be checked in quadruply exponential running time.*

Proof. $\mathcal{A}_{\varphi_p}$ is exponential in the length of φ_p, $\mathcal{D}_{\Lambda_p}$ is doubly exponential; hence, $\mathcal{E}$ is also doubly exponential. Because of the projection in O_e, $\mathcal{F}$ is triply exponential. Because $\mathcal{F}$ is universal, the universal projection in f does not cause a further increase in the number of states, the size of $\mathcal{G}$ is thus triply exponential. Emptiness of a universal automaton can be checked in exponential time, resulting in an overall quadruply exponential running time. □

C.3 Computing Time-bounded Information Flow Assumptions

Our final goal in this section is to compute *time-bounded* information flow assumptions. Time boundedness introduces a new difficulty, because an unbounded number of traces is required to satisfy the same bound; hence, the time-bounded information flow assumption is not a k-hyperproperty for any value $k \in \mathbb{N}$. In the following, we nonetheless represent the time-bounded information flow property as a 2-hyperproperty, by employing the following trick: We introduce a fresh atomic proposition t, which is to be read by process p as a new input and is to be computed by process p's environment. The first occurrence of t indicates that the time bound has been reached. This extra proposition allows us to express the time-bounded information flow assumption as a 2-hyperproperty: we first require that t occurs on every trace that appears as a left trace in Λ_p (condition 1). Furthermore, process p must observe a difference between any pair of traces in Λ_p before t occurs on the left trace (condition 2).

We begin with the universal automaton $\mathcal{A}_{\Lambda_p}$ over alphabet $2^{O_e} \times 2^{O_e}$ from Theorem 5, which recognizes the time-bounded distinguishability relation Λ_p. We dualize $\mathcal{A}_{\Lambda_p}$ to obtain the nondeterministic automaton $\overline{\mathcal{A}_{\Lambda_p}}$ that recognizes all pairs of traces *not* in Λ_p. For condition 1, we construct a nondeterministic automaton $\mathcal{H}_1$ that checks that t occurs on the left trace; for condition 2, we construct a nondeterministic automaton $\mathcal{H}_2$ that ensures that the traces differ in the local inputs before t occurs. Combining $\overline{\mathcal{A}_{\Lambda_p}}$ with $\mathcal{H}_1$ and $\mathcal{H}_2$, we obtain a nondeterministic automaton $\mathcal{I}$ over the alphabet $2^{O_e \cup I_p \cup O_p \cup \{t\}} \times 2^{O_e \cup I_p \cup O_p \cup \{t\}}$ that represents the time-bounded information flow assumption.

Theorem 7. *For a process p with local specification φ_p, there exists a nondeterministic ω-automaton with a doubly-exponential number of states in the length of φ_p that recognizes the time-bounded information flow assumption χ_p induced by φ_p.*

Proof. The automaton $\mathcal{I}$ described above recognizes χ_p. We now claim that it has the stated size. By Theorem 5, the number of states of $\mathcal{A}_{\Lambda_p}$ is doubly-exponential in the length of φ_p. The dual $\overline{\mathcal{A}_{\Lambda_p}}$ has the same size as $\mathcal{A}_{\Lambda_p}$; finally, $\mathcal{H}_1$ and $\mathcal{H}_2$ each has a constant number of states. Thus, the number of states of $\mathcal{I}$ is also doubly-exponential in the length of φ_p. □

D Proofs

D.1 Proofs of Section 4.1

Theorem. *Every implementation that satisfies the local specification φ_p for p also satisfies the information flow assumption ψ_p.*

Proof. Assume that there exists an implementation (s_p, s_q) that satisfies φ_p but not ψ_p. We show that this leads to a contradiction. Since ψ_p is not satisfied, there exists a pair of traces π, π' such that $(\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \in \Delta_p$ and $\pi{\downarrow_{I_p}} = \pi'{\downarrow_{I_p}}$. Let $\pi_e = \pi{\downarrow_{O_e}}$, and $\pi'_e = \pi'{\downarrow_{O_e}}$. Since the inputs to process p are the same on π and

π', and since the strategies s_p and s_q are deterministic, the sequence of outputs is also the same. Let $x_0x_1x_2\ldots = \pi{\downarrow_{I_p}} = \pi'{\downarrow_{I_p}}$ be the sequence of inputs. We construct the sequence of outputs $o_0o_1o_2\ldots$ generated by the implementation as follows: $o_k = s_p(x_0x_1\ldots x_{k-1})$ for all $k \in \mathbb{N}$. Given that the implementation satisfies φ_p, we have that both $\pi_e \sqcup o$ and $\pi'_e \sqcup o$ satisfy φ_p. This, however, contradicts the assumption that $(\pi{\downarrow_{O_e}}, \pi'{\downarrow_{O_e}}) \in \Delta_p$. □

Theorem. *Let Λ_p be a uniform time-bounded distinguishability relation derived from process p's local specification φ_p. Every computation tree that satisfies φ_p also satisfies the time-bounded information flow assumption χ_p.*

Proof. Let (s_a, s_b) be an implementation that satisfies φ_p. We show that the time-bounded information-flow assumption χ_p is satisfied by defining a function $t : (2^{O_e})^\omega \to \mathbb{N}$ such that the 2-hyperproperty given by R_t is satisfied. To compute $t(\pi'_e)$ for some trace of inputs $\pi'_e \in (2^{O_e})^\omega$, we consider the trace of outputs $\pi'_p \in (2^{O_p})^\omega$ obtained by applying the implementation to the prefixes of π'_e. Since Λ_p is uniform, there is a natural number $n \in \mathbb{N}$ such that for all π_e with $(\pi_e, \pi'_e) \in \Lambda_p$, we have that $\pi'_e \sqcup \pi_p \not\models_n \varphi_p$. We set $t(\pi_e)$ to n.

To convince yourself that χ_p is satisfied, suppose, by way of contradiction, that R_t is violated on some pair $(\pi_e, \pi'_e) \in \Lambda_p$ of input traces, i.e., the projection on I_p is the same for π_e and π'_e on the entire prefix of length $t(\pi'_e)$. But then, also the output of process p must be the same along the entire prefix; this, however, means that π'_e will violate φ_p after $n = t(\pi'_e)$ steps, contradicting our assumption that the implementation satisfies φ_p. □

D.2 Proofs of Section 5.2

Lemma. *For all pairs of finite traces $v, v' \in (2^{I_p})^*$, if $K_p(v) = K_p(v')$ then $h(v){\downarrow_{O_p}} = h(v'){\downarrow_{O_p}}$.*

Proof. If $K_p(v)$ is a singleton or empty, then the lemma is trivially true. Assume $|K_p(v)| \geq 1$ and there exists $w, w' \in K_p(v)$ s.t. $h(w){\downarrow_{O_p}} \neq h(w'){\downarrow_{O_p}}$. Since w and w' agree on the local inputs to p, there exists at least one $a \in O_e \backslash I_p$ s.t. $w{\downarrow_{O_a}} \neq w'{\downarrow_{O_a}}$. Then, $h_p(w) \neq h_p(w')$ has to hold following the function f_p of Definition 9. Given the locality from Definition 7, this is only possible if t_p was observed in the input to h_p, which is replaced by the output of h_q in Definition 9. Since h_q satisfies the time bounded information flow assumption χ_p from Definition 4, h_p observes a difference in I_p before it reacts to the global inputs. Therefore, $h(w){\downarrow_{O_p}} = h(w'){\downarrow_{O_p}}$ which contradicts the assumption. □

Theorem. *Let p and q be two processes with time-bounded information flow assumptions χ_p and χ_q, let $h = h_p || h_q$ be the composition of the hyper implementations, and s_p and s_q be the local strategies. Then, for all $v \in (2^{O_e})^*$ it holds that $h(v) = s_p(g_p(v)) \cup s_q(g_q(v))$ where g_p, g_q are defined as follows:*

$$g_p(\epsilon) = \epsilon \qquad g_p(v \cdot x) = g_p(v) \cdot ((x \cap I_p) \cup (s_q(g_q(v)) \cap I_p))$$
$$g_q(\epsilon) = \epsilon \qquad g_q(v \cdot x) = g_q(v) \cdot ((x \cap I_q) \cup (s_p(g_p(v)) \cap I_q))$$

Proof. Proof by induction over $v \in (2^{O_e})^*$. *Base case:* Let $v = \epsilon$, then $s_p(g_p(\epsilon)) \cup s_q(g_q(\epsilon)) = h(\epsilon) = \epsilon$. *Induction Step:* The induction step is shown from $v \in (2^{O_e})^*$ to $v \cdot x \in (2^{O_e})^*$, with $x \in 2^{O_e}$. Inserting g_p from Theorem 3 we obtain $g_p(v \cdot x) = g_p(v) \cdot ((x \cap I_p) \cup (s_q(g_q(v) \cap I_p))$. Since $s_q(g_q(v))$ and $g_p(v)$ is assumed correct, we show that the input trace returned by g_p and given to s_p is correct: The input is local to p because $x \cap I_p$ and $s_q(g_q(v)) \cap I_p)$ remove unobservable inputs, and all outputs of the previous step from q are added to the current input. It remains to show that the outputs of the local strategies combined are equal to the output of h: $h(v \cdot x) = s_p(g_p(v \cdot x)) \cup s_q(g_q(v \cdot x))$. Let $v' \cdot x' = g_p(v \cdot x)$. Given Definition 11 and Definition 10, we know that $s_p(v' \cdot x') = h(K(w)) \downarrow_{O_p}$, with $w \in K_p(v' \cdot x')$. Since $K_p(v')$ is assumed correct, we show that adding x' to v' still results in correctness of $h(K(w))$. Following Lemma 1, all elements in $K(v' \cdot x')$ and therefore all corresponding paths in h have the same label and picking any with $min(K_p(v' \cdot x'))$ is correct. It follows that $h(v \cdot x) \downarrow_{O_p} = s_p(g_p(v \cdot x))$. Using the same argument for s_q by interchanging p and q in every index yields the correctness of the theorem, i.e., for all $v \in (2^{O_e})^*$ it holds that $h(v) = s_p(g_p(v)) \cup s_q(g_q(v))$. □

D.3 Automaton for the Component Specification

We implement the definitions of Section 6 with corresponding automata constructions.

1. By complementing the automaton for the time-bounded distinguishability relation, we obtain an automaton $\mathcal{A}_{IC}$ that associates each trace over O_e with its information class: i.e., the pair (v, w) of traces over O_e is accepted by the complement automaton iff (v, w) is not in the time-bounded distinguishability relation.
2. We obtain the information classes in the following iterative process (under the assumption that the number of information classes is finite):
   (a) We identify some trace v such that there is a pair (v, w) in the language of $\mathcal{A}_{IC}$;
   (b) for each such trace v, we compute an automaton $\mathcal{A}_{[v]}$ for the information class $[v]$, i.e., an automaton that accepts all traces v' with $(v', w') \in \mathcal{L}(\mathcal{A}_{IC})$ iff $(v, w') \in \mathcal{L}(\mathcal{A}_{IC})$ for all w';
   (c) we eliminate all (v, w) with $v \in \mathcal{L}(\mathcal{R})$ from $\mathcal{A}_{IC}$;
   (d) repeat until the language of $\mathcal{A}_{IC}$ is empty.
3. We build an automaton $\mathcal{A}_{\phi_{p,c}}$ for the relativized specification. The automaton uses universal branching to guess the trace from the information class and applies φ_p to each guess.
4. Using the automata $\mathcal{A}_{[v]}$ for the information classes we build an automaton $\mathcal{A}_\psi$ for condition ψ from Definition 14.
5. Using $\mathcal{A}_\psi$ and $\mathcal{A}_{\phi_{p,c}}$, we build an automaton $\mathcal{A}_{\varphi_p}$ for the component specification.