

Logics and Algorithms for Hyperproperties

Bernd Finkbeiner, CISPA Helmholtz Center for Information Security

System requirements related to concepts like information flow, knowledge, and robustness cannot be judged in terms of individual system executions, but rather require an analysis of the relationship between multiple executions. Such requirements belong to the class of hyperproperties, which generalize classic trace properties to properties of sets of traces. During the past decade, a range of new specification logics has been introduced with the goal of providing a unified theory for reasoning about hyperproperties. This paper gives an overview on the current landscape of logics for the specification of hyperproperties and on algorithms for satisfiability checking, model checking, monitoring, and synthesis.

1. INTRODUCTION

In 2008, Clarkson and Schneider coined the term *hyperproperties* for the general class of properties that relate multiple executions of a computer system [Clarkson and Schneider 2008]. In many branches of computer science, however, the study of hyperproperties began long before that. The literature on *information flow security* contains many variations of the *noninterference* property, which requires that for all computations and all sequences of actions of a high-security agent A , the resulting observations made by a low-security observer B are identical to B 's observations that would result without A 's actions [Goguen and Meseguer 1982]. Noninterference has been adapted to a variety of scenarios, attackers, and system models (cf. [Sabelfeld and Sands 2009; Mantel and Reinhard 2007]). In the design of *distributed systems*, information flow and the resulting *knowledge* of agents plays a fundamental role [Fagin et al. 2003; Halpern and Moses 1990]. Another example of a hyperproperty that has been studied in many disciplines is the engineering notion of *robustness* [Dullerud and Paganini 2000], i.e., the guarantee that a system will not deviate strongly from its expected trajectory when disturbances enter the system. Variations of robustness have been applied to a wide range of software systems, from embedded controllers [Barthe et al. 2016] to sorting algorithms [Chaudhuri et al. 2012].

Despite the strong interest in hyperproperties across many different research areas, early proposals for specification languages were limited to hyperproperties from particular domains. A branch of hyperproperties that has seen active logic development are properties related to knowledge, such as the *temporal logic of knowledge* [Fagin et al. 2003]. Reasoning about *combinations of programs* is possible in *relational logics*, such as *relational Hoare Logic* (RHL), originally introduced by Benton [Benton 2004]. Generally, relational logics are restricted to properties that express a condition over a given set of programs or a k -fold self-composition of some program for some fixed k ; there are, however, extensions that are directed at specific properties such as differential privacy [Barthe et al. 2013] and sensitivity [Barthe et al. 2018].

A significant step towards a general logic for hyperproperties was the introduction of *HyperLTL* [Clarkson et al. 2014]. HyperLTL adds universal and existential quantification over traces to standard propositional linear-time temporal logic (LTL). HyperLTL can express complex information-flow properties like generalized noninterference, declassification, and quantitative noninterference. Unlike specialized techniques for particular hyperproperties, HyperLTL provides a unifying logic for expressing and verifying a wide range of policies. A key benefit of such a logic is that it can serve as a common input language and semantic framework for a range of algorithmic tools. In the paper, we provide an overview on algorithms for satisfiability checking, model checking, monitoring, and synthesis with HyperLTL as the common specification language.

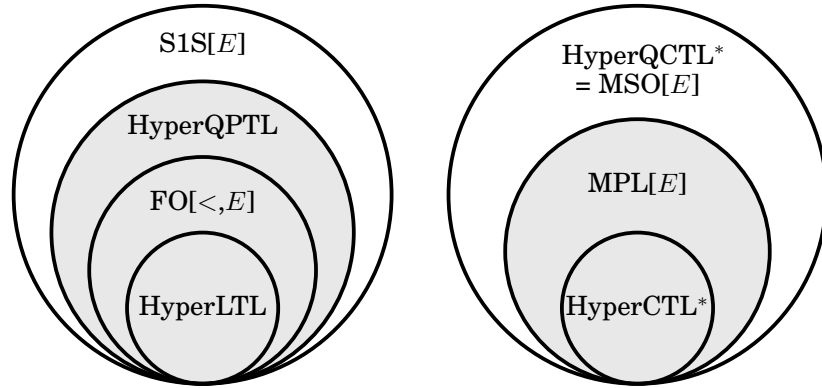


Fig. 1: The hierarchy of hyperlogics [Coenen et al. 2019]. Linear-time logics are shown on the left, branching-time logics on the right. The shaded areas indicate logics with decidable model-checking problem.

While HyperLTL has found many applications, it is perhaps not surprising that there are also numerous hyperproperties that cannot be expressed in HyperLTL. An early observation was that the expressiveness of HyperLTL is incomparable to epistemic temporal logic [Bozzelli et al. 2015]. To address this lack of expressiveness in HyperLTL, a promising idea is to replace LTL with a more powerful logic as the underlying trace logic of HyperLTL.

In the realm of trace properties, LTL is the starting point of a well-known hierarchy of logics, where LTL is equivalent, by Kamp’s theorem [Kamp 1968], to monadic first-order logic of order $FO[<]$. Extending LTL with quantification over propositions leads to the strictly more expressive QPTL [Sistla 1983], which can express all ω -regular properties and is also expressively equivalent to monadic second-order logic of one successor (S1S) [Kesten and Pnueli 1995]. Extending LTL to branching time leads to the temporal logic CTL^* , which has been shown to be expressively equivalent to monadic path logic (MPL) [Abrahamson 1980]. Adding quantification over propositions leads to the strictly more expressive quantified computation tree logic ($QCTL^*$) [French 2001], which is expressively equivalent to monadic second-order logic (MSO) [Laroussinie and Markey 2014].

Using the extension of LTL to HyperLTL as a blueprint, we extend the temporal logics QPTL, CTL^* , and $QCTL^*$ with quantification over traces or paths to obtain the corresponding hyperlogics. A different method for the construction of hyperlogics has been introduced for first-order and second-order logics. Here, the extension consists of adding the *equal-level predicate* E (cf. [Thomas 2009; Finkbeiner and Zimmermann 2017]), which relates the same time points on different traces. Figure 1 shows the resulting hierarchy of hyperlogics (introduced in [Coenen et al. 2019]). For linear time, the most striking difference to the hierarchy of the standard logics is that, when lifted to hyperlogics, $FO[<]$ is no longer equivalent to LTL, and S1S is no longer equivalent to QPTL: $S1S[E]$ is strictly more expressive than HyperQPTL; HyperQPTL, in turn, is strictly more expressive than $FO[<, E]$, and $FO[<, E]$ is strictly more expressive than HyperLTL. For branching time, we have that $HyperQCTL^*$ is still expressively equivalent to $MSO[E]$. However, MPL, which is equivalent to CTL^* in the standard hierarchy, falls strictly between $HyperCTL^*$ and $HyperQCTL^*$ when equipped with the equal-level predicate ($MPL[E]$).

The logics in the hierarchy significantly extend the expressiveness of HyperLTL. Already the extension to QPTL suffices, for example, to subsume epistemic temporal

logic [Rabe 2016]. At the same time, a large part of the hierarchy is still supported by algorithmic analysis. The shaded areas of Fig. 1 indicate logics where the model-checking problem is decidable.

The focus of the next sections is on HyperLTL. We introduce the syntax and semantics of the logic in Section 2 and discuss the satisfiability, model checking, monitoring, and synthesis problems in Sections 3–6. We return to the hierarchy of hyperlogics in Section 7, where we also briefly discuss further extensions to probabilistic, infinite-state, and asynchronous hyperproperties.

2. HYPERLTL

HyperLTL [Clarkson et al. 2014] is an extension of linear-time temporal logic (LTL) [Pnueli 1977]. In the following, we briefly define the syntax and semantics of LTL and then describe the extension to HyperLTL. Let AP be a finite set of atomic propositions. A trace over AP is a map $t: \mathbb{N} \rightarrow 2^{AP}$, denoted by $t(0)t(1)t(2)\dots$. Let $(2^{AP})^\omega$ denote the set of all traces over AP .

2.1. LTL

The formulas of LTL are generated by the following grammar:

$$\varphi ::= a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi$$

where $a \in AP$ is an *atomic proposition*, the Boolean connectives \neg and \wedge have the usual meaning, \mathbf{X} is the temporal *next* operator, and \mathbf{U} is the temporal *until* operator. We also consider the usual derived Boolean connectives, such as \vee , \rightarrow , and \leftrightarrow , and the derived temporal operators *eventually* $\mathbf{F}\varphi \equiv \text{true} \mathbf{U}\varphi$, *globally* $\mathbf{G}\varphi \equiv \neg\mathbf{F}\neg\varphi$, *weak until* $\varphi \mathbf{W}\psi \equiv (\varphi \mathbf{U}\psi) \vee \mathbf{G}\varphi$, and *release* $\varphi \mathbf{R}\psi \equiv \neg(\neg\varphi \mathbf{U}\neg\psi)$. The satisfaction of an LTL formula φ over a trace t at a position $i \in \mathbb{N}$, denoted by $t, i \models \varphi$, is defined as follows:

$$\begin{aligned} t, i \models a & \quad \text{iff } a \in t(i), \\ t, i \models \neg\varphi & \quad \text{iff } t, i \not\models \varphi, \\ t, i \models \varphi_1 \wedge \varphi_2 & \quad \text{iff } t, i \models \varphi_1 \text{ and } t, i \models \varphi_2, \\ t, i \models \mathbf{X}\varphi & \quad \text{iff } t, i+1 \models \varphi, \\ t, i \models \varphi_1 \mathbf{U}\varphi_2 & \quad \text{iff } \exists k \geq i : t, k \models \varphi_2 \wedge \forall i \leq j < k : t, j \models \varphi_1. \end{aligned}$$

We say that a trace t satisfies a sentence ϕ , denoted by $t \models \phi$, if $t, 0 \models \phi$. For example, the LTL formula $\mathbf{G}(a \rightarrow \mathbf{F}b)$ specifies that every position in which a is true must eventually be followed by a position where b is true.

2.2. HyperLTL

The formulas of HyperLTL are generated by the grammar

$$\begin{aligned} \phi & ::= \exists\pi. \phi \mid \forall\pi. \phi \mid \psi \\ \psi & ::= a_\pi \mid \neg\psi \mid \psi \wedge \psi \mid \mathbf{X}\psi \mid \psi \mathbf{U}\psi \end{aligned}$$

where a is an atomic proposition from a set AP and π is a trace variable from a set \mathcal{V} . Further Boolean connectives and the temporal operators \mathbf{F} , \mathbf{G} , \mathbf{W} , and \mathbf{R} are derived as for LTL.

The semantics of HyperLTL is defined with respect to a trace assignment, a partial mapping $\Pi: \mathcal{V} \rightarrow (2^{AP})^\omega$. The assignment with empty domain is denoted by Π_\emptyset . Given a trace assignment Π , a trace variable π , and a trace t , we denote by $\Pi[\pi \rightarrow t]$ the assignment that coincides with Π everywhere but at π , which is mapped to t . The satisfaction of a HyperLTL formula ϕ over a trace assignment Π and a set of traces T at a position $i \in \mathbb{N}$, denoted by $T, \Pi, i \models \phi$, is defined as follows:

$T, \Pi, i \models a_\pi$	iff	$a \in \Pi(\pi)(i),$
$T, \Pi, i \models \neg\psi$	iff	$T, \Pi, i \not\models \psi,$
$T, \Pi, i \models \psi_1 \wedge \psi_2$	iff	$T, \Pi, i \models \psi_1$ and $T, \Pi, i \models \psi_2,$
$T, \Pi, i \models \mathbf{X}\psi$	iff	$T, \Pi, i + 1 \models \psi,$
$T, \Pi, i \models \psi_1 \mathbf{U} \psi_2$	iff	$\exists k \geq i. T, \Pi, k \models \psi_2$ $\wedge \forall i \leq j < k. T, \Pi, j \models \psi_1,$
$T, \Pi, i \models \exists\pi. \phi$	iff	$\exists t \in T. T, \Pi[\pi \rightarrow t], i \models \phi,$
$T, \Pi, i \models \forall\pi. \phi$	iff	$\forall t \in T. T, \Pi[\pi \rightarrow t], i \models \phi.$

We say that a set T of traces satisfies a sentence ϕ , denoted by $T \models \phi$, if $T, \Pi_\emptyset, 0 \models \phi$.

A *Kripke structure* is a tuple $\mathcal{K} = (S, s_0, \delta, \text{AP}, L)$ consisting of a set of states S , an initial state s_0 , a transition function $\delta : S \rightarrow 2^S$, a set of *atomic propositions* AP, and a *labeling function* $L : S \rightarrow 2^{\text{AP}}$ that assigns to each state a set of atomic propositions that are true in the state. We require that each state s has a successor, that is $\delta(s) \neq \emptyset$, to ensure that every execution of a Kripke structure can always be continued to infinity. In a *finite* Kripke structure, S is a finite set.

A *path* of a Kripke structure is an infinite sequence $s_0 s_1 \dots \in S^\omega$ such that s_0 is the initial state of \mathcal{K} and $s_{i+1} \in \delta(s_i)$ for all $i \in \mathbb{N}$. By $\text{Paths}(\mathcal{K}, s)$, we denote the set of all paths of \mathcal{K} starting in state $s \in S$. A *trace* of a path $\sigma = s_0 s_1 \dots$ is a sequence of labels $l_0 l_1 \dots$ with $l_i = L(s_i)$ for all $i \in \mathbb{N}$. $\text{Tr}(\mathcal{K}, s)$ is the set of all traces of paths of a Kripke structure \mathcal{K} starting in state s . A Kripke structure \mathcal{K} with initial state s_0 satisfies an LTL formula φ , denoted by $\mathcal{K} \models \varphi$ iff for all traces $\pi \in \text{Tr}(\mathcal{K}, s_0)$, it holds that $\pi \models \varphi$. Likewise, the Kripke structure satisfies a HyperLTL formula ϕ , also denoted by $\mathcal{K} \models \phi$, iff $\text{Tr}(\mathcal{K}, s_0) \models \phi$.

2.3. Examples

We illustrate the specification of hyperproperties with HyperLTL using some prominent examples from information-flow security.

- A system satisfies *observational determinism* [Zdancewic and Myers 2003] if every pair of traces with the same low-security input remains indistinguishable for an observer; i.e., the program appears to be deterministic to low-security users. For a system with low-security input l and output o , observational determinism can be expressed in HyperLTL as follows:

$$\forall\pi. \forall\pi'. (l_\pi \leftrightarrow l_{\pi'}) \rightarrow \mathbf{G} (o_\pi \leftrightarrow o_{\pi'})$$

- *Noninference* [McLean 1994] specifies that the behavior observable by an observer must not change when all high-security inputs are replaced by a dummy input. In the following HyperLTL formula, we express noninference for a high-security input h , which we set to *false* as dummy input, low-security input l , and output o :

$$\forall\pi. \exists\pi'. \mathbf{G} (\neg h_{\pi'} \wedge (l_\pi \leftrightarrow l_{\pi'}) \wedge (o_\pi \leftrightarrow o_{\pi'}))$$

- *Generalized noninterference* [McCullough 1988] is similar to observational determinism in that it requires that the output (o) should not be influenced by high-security input (h). Unlike observational determinism, it allows, however, for nondeterministic behavior in the low-security input (l):

$$\forall\pi. \forall\pi'. \exists\pi''. \mathbf{G} ((h_\pi \leftrightarrow h_{\pi''}) \wedge (l_{\pi'} \leftrightarrow l_{\pi''}) \wedge (o_{\pi'} \leftrightarrow o_{\pi''}))$$

The existentially quantified trace π'' combines the high-security input of the universally quantified trace π with the low-security input and output of the universally quantified trace π' .

Table I: Complexity of HyperLTL satisfiability for formulas from the full logic, and for the temporal safety and liveness fragments. A HyperLTL formula is temporal safety (resp. temporal liveness) if its LTL body describes a safety (resp. liveness) property. All results denote completeness.

Full HyperLTL		Temporal Safety		Temporal Liveness	
full logic	Σ_1^1 ¹	full fragment	coRE ³	full fragment	Σ_1^1 ³
$\exists^*\forall^*$	EXPSpace ²	$\forall^*\exists^* \mathbf{G}(\mathbf{X}^*)$	coRE ³	$\forall\exists^* \mathbf{F}(\mathbf{X}^*)$	NP ³
\exists^*	PSPACE ²	$\forall\exists^* \mathbf{G}$	NEXP ³	$\forall\exists^* \mathbf{F} \wedge \dots \wedge \mathbf{F}$	NP ³
\forall^*	PSPACE ²	$\forall^*\exists^* \mathbf{X}^*$	NEXP ³		

¹ [Fortin et al. 2021]

² [Finkbeiner and Hahn 2016]

³ [Beutner et al. 2022]

As discussed in the introduction, hyperproperties have many applications beyond security. An illustrative example is the specification of *error resistant codes*, which transmit data over noisy channels (cf. [Finkbeiner et al. 2015]). A typical correctness condition for such a code is that all code words have a certain minimal Hamming distance. The following HyperLTL formula specifies that all code words produced by an encoder with input i and output o have a Hamming distance of at least d :

$$\forall\pi.\forall\pi'. (\mathbf{F}(i_\pi \leftrightarrow \neg i_{\pi'})) \rightarrow \neg\text{Ham}(d-1, \pi, \pi')$$

where the subformula $\text{Ham}(d, \pi, \pi')$ is defined recursively as follows:

$$\begin{aligned} \text{Ham}(-1, \pi, \pi') &= \text{false} \\ \text{Ham}(d, \pi, \pi') &= (o_\pi \leftrightarrow o_{\pi'}) \mathbf{W}((o_\pi \leftrightarrow \neg o_{\pi'}) \wedge \mathbf{X}\text{Ham}(d-1, \pi, \pi')). \end{aligned}$$

The subformula $\text{Ham}(d, \pi, \pi')$ expresses that the Hamming distance between the outputs on π and π' is *less than* d : for $d \geq 0$ this means that either π and π' agree forever on the output o , or there is a point in time where they differ and the suffixes from that point onward have a Hamming distance of less than $d-1$. The full specification thus means that on any pair of traces π, π' where the input eventually differs, the Hamming distance must not be less than d .

3. SATISFIABILITY

In the HyperLTL satisfiability problem, we decide, for a given HyperLTL formula ϕ , whether or not there exists a set T of traces such that $T \models \phi$. Table I gives an overview on the complexity of the satisfiability problem. Satisfiability of HyperLTL formulas without quantifier alternation is PSPACE-complete [Finkbeiner and Hahn 2016], which is the same complexity as LTL satisfiability [Sistla and Clarke 1985]. For formulas from the $\exists^*\forall^*$ fragment, the complexity increases to EXPSpace [Finkbeiner and Hahn 2016]. Formulas of this fragment can be translated into equisatisfiable (but exponentially larger) formulas with only existential formulas by explicitly enumerating all possible interactions between the universal and existential quantifiers. Satisfiability of hyperproperties with $\forall^*\exists^*$ trace quantifier alternation is undecidable [Finkbeiner and Hahn 2016; Fortin et al. 2021].

3.1. Satisfiability of the $\forall^*\exists^*$ fragment

The $\forall^*\exists^*$ fragment contains many interesting hyperproperties, such as noninterference and generalized noninterference. Despite the general undecidability, some positive re-

sults have been obtained by classifying hyperproperties into temporal safety and temporal liveness [Beutner et al. 2022]. A HyperLTL formula is *temporal safety* (resp. *temporal liveness*) if its LTL body describes a safety (resp. liveness) property.¹ The restriction to temporal safety reduces the complexity from Σ_1^1 to coRE. While still undecidable, this enables using common first-order techniques such as resolution, tableaux, and related methods [Robinson and Voronkov 2001]. In contrast to temporal safety properties, the class of temporal liveness formulas is of analytical complexity. As shown in Table I, both temporal safety and temporal liveness contain decidable fragments. A notable difference is that while formulas from the $\forall^*\exists^* G(X^*)$ fragment are undecidable, formulas from the $\forall\exists^* F(X^*)$ are decidable [Beutner et al. 2022].

3.2. Practical algorithms

Despite the general undecidability, there is significant practical interest in algorithms that detect satisfiability and unsatisfiability of HyperLTL formulas.

Decision procedures. The decidability of the $\exists^*\forall^*$ fragment is particularly useful in practice because it allows us to check implications between alternation-free formulas. The tool EAHyper implements the decision procedure as a reduction to LTL satisfiability [Finkbeiner et al. 2017].

Finding largest models. For formulas in the $\forall\exists^*$ fragment, an incomplete method to detect (un)satisfiability is to search for the largest model that satisfies the formula. For satisfiable $\forall\exists^*$ formulas, this model is unique. The algorithm iteratively eliminates choices for the \exists^* quantifiers that admit no witness trace when chosen for the \forall quantifier [Beutner et al. 2022].

Finding smallest models. A method for general HyperLTL formulas is to search for finite models of bounded size and then iteratively increase the bound [Finkbeiner et al. 2018; Mascle and Zimmermann 2020]. Such an approach finds smallest models, but cannot determine unsatisfiability.

4. MODEL CHECKING

In model checking, we decide, for a given finite Kripke structure \mathcal{K} and a given HyperLTL formula ϕ , whether or not $\mathcal{K} \models \phi$. The model-checking problem is decidable. As shown in the right-most column of Table II, the complexity is non-elementary in the number of quantifier alternations; each quantifier alternation causes a further exponential cost. Fortunately, many hyperproperties of interest have only a small number of quantifier alternations. Observational determinism, for example, is a universal formula and, therefore, only requires nondeterministic logarithmic space in the size of the Kripke structure; noninference and generalized noninterference both have a single quantifier alternation and, hence, can be checked in PSPACE. The table furthermore shows that quantifier alternation is much less costly for Kripke structures with restricted structures, such as acyclic graphs and trees.² For acyclic graphs, the complexity of all model-checking problems is in the polynomial hierarchy within PSPACE, for trees even in deterministic logarithmic space. In the following, we first sketch a basic automata-theoretic model-checking procedure, and then discuss more practical approaches.

¹Temporal safety and temporal liveness differ from the notions of hypersafety and hyperliveness [Clarkson and Schneider 2008].

²To satisfy the requirement that every state of a Kripke structure has a successor state, we assume implicit self-loops on terminal states in tree-shaped and acyclic Kripke structures [Bonakdarpour and Finkbeiner 2018].

Table II: Complexity of the HyperLTL model-checking problem for tree-shaped, acyclic, and general Kripke structures. The complexity is given in the size of the Kripke structure, where k is the number of quantifier alternations in $(\forall^*\exists^*)^*$. All results denote completeness.

	Trees	Acyclic graphs	General graphs
\forall^+/\exists^+	L ¹	NL ¹	NL ²
$\exists^+\forall^+/\forall^+\exists^+$		NP/coNP ¹	PSPACE ³
$(\forall^*\exists^*)^*$		Π_k^p ¹ Σ_k^p ¹	$(k-1)$ -EXPSPACE ⁴

¹ [Bonakdarpour and Finkbeiner 2018]

² [Finkbeiner et al. 2015]

³ [Clarkson et al. 2014]

⁴ [Rabe 2016]

4.1. The basic algorithm

The following basic algorithm (described in more detail in [Finkbeiner et al. 2015]) reduces the model-checking problem to the language emptiness problem of a Büchi automaton: the given Kripke structure satisfies the formula if and only if the language of the resulting automaton is empty.

The construction starts by negating the given formula, so that the resulting formula describes the existence of an error. Let the new formula have the form $Q_n\pi_n Q_2\pi_{n-1} \dots Q_1\pi_1. \psi$, where Q_1, Q_2, \dots, Q_n are trace quantifiers in $\{\exists, \forall\}$, and ψ is a quantifier-free formula over atomic propositions indexed by trace variables $\{\pi_1, \dots, \pi_n\}$. Similar to standard LTL model checking, we convert the LTL formula ψ into an equivalent Büchi automaton \mathcal{A}_0 over the alphabet $(2^{\text{AP}})^n$. Each letter is a tuple of n sets of atomic propositions, where the i th element of the tuple represents the atomic propositions of trace π_i .

The algorithm then eliminates the quantifiers one at a time, starting with the innermost quantifier. Before the elimination of the i th quantifier, the automaton \mathcal{A}_{i-1} over alphabet $(2^{\text{AP}})^{(n-i)}$ has been constructed, and now the first component of the tuple corresponds to π_i . We assume that Q_i is an existential quantifier; if Q_i is a universal quantifier, we turn it into an existential quantifier by complementing \mathcal{A}_{i-1} , based on the equivalence $\forall\pi. \phi = \neg\exists\neg\phi$.

We then combine \mathcal{A}_{i-1} with the Kripke structure \mathcal{K} to ensure that the first component is chosen consistently with some path in \mathcal{K} . Since Q_i is an existential quantifier, we eliminate the first component of the tuple by existential projection on the automaton. This results in the next automaton \mathcal{A}_i . After n such steps, all quantifiers have been eliminated and the language of the resulting automaton is over the one-letter alphabet consisting of the empty tuple. The HyperLTL formula is satisfied if and only if the language of automaton \mathcal{A}_n is empty.

4.2. Practical algorithms

Because of the expensive automata constructions, it is difficult to implement the basic algorithm efficiently. In the following, we briefly discuss some more practical approaches.

Self-composition. For the alternation-free fragment of HyperLTL, the model-checking problem can be solved by *self-composition* [Barthe et al. 2011]. If the system is given as a circuit C and the HyperLTL formula consists of an LTL formula ψ and a quantifier prefix consisting of k exclusively universal quantifiers (the case of exclusively existential quantifiers is analogous), then a new circuit can be constructed that consists of k copies of C and a monitor circuit C_ψ that recognizes violations of ψ . The HyperLTL model-checking problem then reduces to a standard model-checking problem of a trace property on the newly constructed circuit. This approach has been implemented in the HyperLTL model checker MCHyper [Finkbeiner et al. 2015].

Strategy-based verification. HyperLTL properties with a $\forall^*\exists^*$ quantifier alternation can be verified by finding a winning strategy for the “existential” player in a game against a “universal” player [Coenen et al. 2019]. In this game, the existential and universal players chose the existential and universal traces, respectively, by picking the successor state on paths through the Kripke structure, one position at a time. This strategy can be provided manually or synthesized by a game-solving algorithm. Strategy-based verification is incomplete because the existential player must choose the next state after observing only a finite prefix of the universal traces. The approach can be made complete by adding prophecy variables [Beutner and Finkbeiner 2022a]. Strategy-based verification was originally developed for finite-state hardware but has been extended to infinite-state software [Beutner and Finkbeiner 2022b].

Bounded Model Checking. In bounded model checking (BMC) for HyperLTL [Hsu et al. 2021], the system is unfolded up to a fixed depth. Pending obligations beyond that depth are either treated pessimistically (to show the satisfaction of a formula) or optimistically (to show the violation of a formula). While BMC for trace properties reduces to SAT solving, BMC for hyperproperties naturally reduces to QBF solving. As usual for bounded methods, BMC for HyperLTL is incomplete.

Language inclusion. The complementation operation in the basic algorithm, which is necessary for every quantifier alternation, can be replaced with language inclusion [Beutner and Finkbeiner 2023]. To check a HyperLTL formula $\phi = \forall\pi.\exists\pi'.\psi$, we first construct an automaton by intersecting a self-composition of the Kripke structure with the automaton for ψ over the alphabet $(2^{\text{AP}})^2$, and then eliminating the second component by existential projection, resulting in an automaton over the alphabet 2^{AP} . The Kripke structure satisfies ϕ iff its language is contained in the language of this automaton. Model checking via language inclusion is a complete verification method.

5. MONITORING

In monitoring, we determine, at runtime and in an online fashion, whether a system under observation satisfies a given HyperLTL formula. For trace properties, this is typically done by observing growing prefixes of the trace produced by a running system. Monitoring hyperproperties is more difficult because a violation of a hyperproperty generally involves a set of traces. Two principal ways in which the traces can be presented to the monitor are the parallel and sequential monitoring models [Finkbeiner et al. 2019] illustrated in Fig. 2. In the *parallel model*, the assumption is that the number of traces is fixed in advance. All traces then become available simultaneously, one position at a time from left to right. The parallel model occurs, for example, in secure-

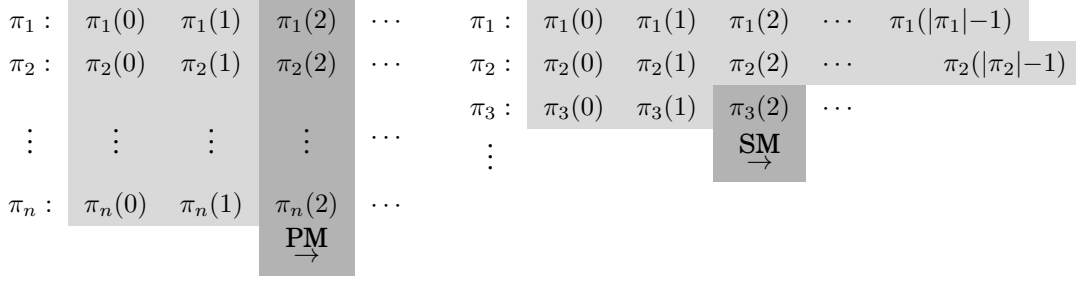


Fig. 2: Parallel monitoring (PM) vs. sequential monitoring (SM) of hyperproperties. In parallel monitoring, the number of traces is fixed in advance, and the monitor observes one position of all traces at a time; in sequential monitoring, the number of traces is a priori unbounded, and the monitor observes one position of one trace at a time.

multi-execution [Devriese and Piessens 2010], where several system executions are generated by providing different high-security inputs. In the *sequential model*, the traces become available one at a time. In this model, the number of traces is a priori unbounded. The sequential model applies, for example, when multiple sessions of a system are to be monitored one after the other in an online fashion. In the following, we focus on the sequential model.

5.1. Finite-trace HyperLTL

Since the monitoring verdict is based on observing finite traces, we adapt the semantics of HyperLTL to this setting. Let $\Pi_{fin} : \mathcal{V} \rightarrow (2^{AP})^+$ be a partial function mapping trace variables to finite traces. By $t \in \Pi_{fin}$ we denote that trace t is in the image of Π_{fin} . The satisfaction of a HyperLTL formula ϕ over a finite trace assignment Π_{fin} and a set of finite traces T , denoted by $(T, \Pi_{fin}, i) \models \phi$, is defined as follows [Brett et al. 2017; Hahn et al. 2019]:

$$\begin{aligned}
T, \Pi_{fin}, i \models a_\pi & \quad \text{iff} \quad a \in \Pi_{fin}(\pi)(i), \\
T, \Pi_{fin}, i \models \neg\psi & \quad \text{iff} \quad T, \Pi_{fin}, i \not\models \psi, \\
T, \Pi_{fin}, i \models \psi_1 \wedge \psi_2 & \quad \text{iff} \quad T, \Pi_{fin}, i \models \psi_1 \text{ and } T, \Pi_{fin}, i \models \psi_2, \\
T, \Pi_{fin}, i \models \mathbf{X}\psi & \quad \text{iff} \quad \forall t \in \Pi_{fin}. |t| > i + 1 \text{ and } T, \Pi_{fin}, i + 1 \models \psi, \\
T, \Pi_{fin}, i \models \psi_1 \mathbf{U} \psi_2 & \quad \text{iff} \quad \exists k \geq i. k < \min_{t \in \Pi_{fin}} |t| \wedge T, \Pi_{fin}, j \models \psi_2 \\
& \quad \wedge \forall i \geq j < k. T, \Pi_{fin}, j \models \psi_1, \\
T, \Pi_{fin}, i \models \exists\pi. \phi & \quad \text{iff} \quad \exists t \in T. T, \Pi_{fin}[\pi \mapsto t], i \models \phi, \\
T, \Pi_{fin}, i \models \forall\pi. \phi & \quad \text{iff} \quad \forall t \in T. T, \Pi_{fin}[\pi \mapsto t], i \models \phi.
\end{aligned}$$

Note that the finite traces may differ in length; since the temporal operators advance in all considered traces simultaneously, there is an implicit cut-off at the end of the shortest trace.

5.2. The basic algorithm

We describe a basic algorithm for sequentially monitoring HyperLTL formulas of the form $\forall\pi_1, \dots, \pi_n. \psi$, where ψ is an LTL formula expressing a safety property (for a more general approach see [Finkbeiner et al. 2019]). From ψ , we construct a deterministic finite-word automaton over alphabet $(2^{AP})^n$, which recognizes all sequences such that the traces π_1, \dots, π_n , obtained from the components of the tuples, violate ψ . The monitoring algorithm stores the set T of traces seen so far and a map S that assigns to each n -tuple of traces seen so far a state of the automaton. When a new trace π starts, S is

initialized with all n -tuples from traces in $T \cup \{\pi\}$ that contain the new trace π . When the current trace π progresses, we update the state assigned to each tuple with the successor state of the automaton. If the automaton reaches a final state, a violation is reported.

5.3. Practical algorithms

The currently available algorithms for monitoring hyperproperties can be grouped into two main approaches:

Combinatorial approaches. As combinatorial approaches [Agrawal and Bonakdarpour 2016; Finkbeiner et al. 2019] we understand algorithms that store the traces seen so far and explicitly iterate over the different combinations of the traces. Such approaches can be optimized with efficient data structures for storing sets of traces, such as prefix trees [Finkbeiner et al. 2020], and by reducing the number of tuples that need to be considered, for example by recognizing reflexive, symmetric, and transitive specifications. Such optimizations are implemented in the monitoring tool RVHyper [Finkbeiner et al. 2018].

Constraint-based approaches. Constraint-based approaches [Brett et al. 2017; Hahn et al. 2019] avoid storing the traces explicitly by translating the monitoring task into a constraint system. As new traces arrive, the constraint system is checked for satisfiability and rewritten according to the resulting requirements on future traces.

6. SYNTHESIS

In synthesis, we decide whether there exists an implementation that satisfies a given HyperLTL formula, and, if so, construct such an implementation. We refer to the decision part of synthesis as the *realizability* problem. We consider HyperLTL formulas where the atomic propositions $AP = I \cup O$ are partitioned into a set I of *inputs* and a set O of *outputs*. We are interested in finding implementations, called *strategies*, which observe the inputs and compute the outputs in such a way that the specification is satisfied.

Formally, a strategy $f: (2^I)^* \rightarrow 2^O$ is a function that maps sequences of input valuations 2^I to an output valuation 2^O . The behavior of a strategy $f: (2^I)^* \rightarrow 2^O$ is characterized by an infinite tree, called *computation tree*, that branches by the valuations of I and whose nodes $w \in (2^I)^*$ are labeled with the strategic choice $f(w)$. For an infinite word $w = w_0w_1w_2 \dots \in (2^I)^\omega$, the corresponding *trace* is defined as $(f(\epsilon) \cup w_0)(f(w_0) \cup w_1)(f(w_0w_1) \cup w_2) \dots \in (2^{I \cup O})^\omega$. The set of traces produced by f , written $\text{traces}(f)$, is $\{w \mid w \in f\}$. We define the satisfaction of a HyperLTL formula ϕ (over propositions $I \cup O$) on strategy f , written $f \models \phi$, as $\text{traces}(f) \models \phi$. A HyperLTL formula ϕ is *realizable* if there is a strategy $f: (2^I)^* \rightarrow 2^O$ that satisfies ϕ .

The realizability problem is undecidable. As we will see in the following, already the universal fragment suffices to encode the distributed synthesis problem, which is known to be undecidable [Pnueli and Rosner 1990]. The \exists^* and $\exists^*\forall^1$ fragments of HyperLTL are decidable [Finkbeiner et al. 2020b].

6.1. Distributed synthesis

Due to the expressiveness of HyperLTL, the HyperLTL synthesis problem subsumes various classic synthesis problems. For the fragment of HyperLTL with only a single, universal quantifier $\forall\pi$, ψ , the HyperLTL realizability problem is exactly the LTL realizability problem of ψ . With two universal quantifiers, we can express various LTL realizability problems with restrictions on the flow of information, such as synthesis under incomplete information [Kupferman and Vardi 1997], distributed synthesis [Pnueli and Rosner 1990], and fault-tolerant synthesis [Dimitrova and Finkbeiner 2009]. We

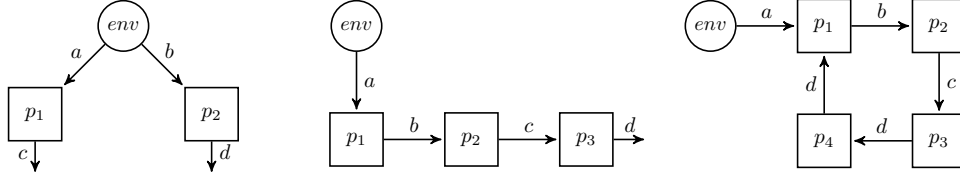


Fig. 3: Example architectures of distributed systems.

illustrate this in the following with the classic distributed synthesis problem (for more details and encodings of other synthesis problems, see [Finkbeiner et al. 2020b]).

The *distributed synthesis* problem introduces the concept of *architectures* as a constraint on the information flow. An architecture is a set of processes P , with distinct environment process $p_{env} \in P$, such that the processes produce outputs synchronously, but each process bases its decision only on the history of valuation of inputs that it observes.

Formally, a distributed architecture \mathcal{A} is a tuple $\langle P, p_{env}, \mathcal{I}, \mathcal{O} \rangle$ where P is a finite set of processes with distinguished environment process $p_{env} \in P$. The functions $\mathcal{I}: P \rightarrow 2^{AP}$ and $\mathcal{O}: P \rightarrow 2^{AP}$ define the inputs and outputs of processes. While processes may share the same inputs (in case of broadcasting), the outputs of processes must be pairwise disjoint, i.e., for all $p \neq p' \in P$ it holds that $\mathcal{O}(p) \cap \mathcal{O}(p') = \emptyset$. W.l.o.g. we assume that $\mathcal{I}(p_{env}) = \emptyset$. We denote by $P^- = P \setminus \{p_{env}\}$ the set of processes excluding the environment process. Figure 3 shows several example architectures.

To encode distributed realizability as a HyperLTL realizability problem, we first define the notion of *independence* as a hyperproperty:

$$D_{A \rightarrow C}^{\pi, \pi'} ::= \left(\bigvee_{a \in A} (a_\pi \leftrightarrow a_{\pi'}) \right) \mathbf{R} \left(\bigwedge_{c \in C} (c_\pi \leftrightarrow c_{\pi'}) \right)$$

$D_{A \rightarrow C}^{\pi, \pi'}$ requires that the valuations of propositions C on traces π and π' have to be equal until and including the point in time where there is a difference in the valuation of some proposition in A . Prefacing universal quantification, that is, the formula $\forall \pi \forall \pi'. D_{A \rightarrow C}^{\pi, \pi'}$ guarantees that every proposition $c \in C$ solely depends on propositions in A . Distributed realizability of $\langle \varphi, \mathcal{A} \rangle$ then corresponds to the realizability of the following HyperLTL formula over inputs $\mathcal{O}(p_{env})$ for P^- and outputs $\bigcup_{p \in P^-} \mathcal{O}(p)$:

$$\forall \pi \forall \pi'. \varphi[\pi] \wedge \bigwedge_{p \in P^-} D_{\mathcal{I}(p) \rightarrow \mathcal{O}(p)}^{\pi, \pi'}$$

6.2. Practical algorithms

Since the model-checking problem is decidable, the synthesis problem also becomes decidable if the set of potential solutions is finite. Bounded synthesis and controller synthesis ensure this by introducing a bound on the size of the implementation and by fixing the state graph of the implementation, respectively.

Bounded synthesis. In bounded synthesis [Finkbeiner and Schewe 2013], the size of the implementation is bounded by a constant, which can be increased incrementally. For universal HyperLTL formulas, this approach has been used to find finite generators of realizing strategies by encoding this search as a satisfiability problem for a decidable constraint system [Finkbeiner et al. 2020b]. In order to detect unrealizable specifications, the approach simultaneously searches for *counterexamples to realizabil-*

ity. For a universal HyperLTL formula $\phi = \forall\pi_1 \cdots \forall\pi_n. \psi$ over inputs I and outputs O , a counterexample to realizability is a set of input traces $T \subseteq (2^I)^\omega$ such that for every strategy $f: (2^I)^* \rightarrow 2^O$, the trace set that results from applying f to the input traces in T satisfies $\neg\phi = \exists\pi_1 \cdots \exists\pi_n. \neg\psi$.

Controller synthesis. The controller synthesis problem is a simpler variant of synthesis. In addition to the HyperLTL formula, we are given a *plant model*, which provides the state space of the implementation. The transitions of the plant model are partitioned into controllable and uncontrollable transitions. The goal of the controller synthesis problem is to eliminate a subset of the controllable transitions in such a way that the specification is satisfied. Controller synthesis is decidable for HyperLTL [Bonakdarpour and Finkbeiner 2020]. Controller synthesis has also been combined with synthesis from trace properties. In a first step, the plant is synthesized by finding a (most permissive) strategy for a safety trace property that defines the desired functionality of the system. In a second step, transitions that violate a given hyperproperty are removed from the plant by controller synthesis [Coenen et al. 2023].

7. THE HIERARCHY OF HYPERLOGICS

HyperLTL is the starting point of the hierarchy of *hyperlogics* discussed in the introduction and shown in Fig. 1. This hierarchy was initially studied in [Coenen et al. 2019]; it adds significant expressiveness to HyperLTL. Similar to the extension of LTL with trace quantification, the temporal logics HyperQPTL, HyperCTL*, and HyperQCTL* are obtained from the standard temporal logics QPTL, CTL*, and QCTL* by adding quantifiers over traces (or, in the case of CTL* and QCTL*, paths), so that the formula can refer to multiple traces or paths at the same time. The other logics in the hierarchy, FO[<, E], S1S[E], MPL[E], and MSO[E] are obtained from the standard first-order and second-order logics FO, S1S, MPL, and MSO by adding the equal-level predicate E (cf. [Thomas 2009; Finkbeiner and Zimmermann 2017]), which indicates that two points happen at the same time.

7.1. Linear-time hyperlogics

The logics shown on the left in Fig. 1 are linear-time logics.

FO[<, E]. FO[<, E] is a first-order logic for the specification of hyperproperties [Finkbeiner and Zimmermann 2017]. Kamp’s theorem [Kamp 1968] (in the formulation of [Gabbay et al. 1980]) states that linear-time temporal logic (LTL) is expressively equivalent to monadic first-order logic of order FO[<]. To express hyperproperties, we add to FO[<] the equal-level predicate E . Given a set of atomic propositions AP and a set V_1 of first-order variables, we define the syntax of FO[<, E] formulas as follows:

$$\begin{aligned} \tau &::= P_a(x) \mid x < y \mid x = y \mid E(x, y) \\ \phi &::= \tau \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \exists x. \phi, \end{aligned}$$

where $a \in AP$ and $x, y \in V_1$. While FO[<] formulas are interpreted over a trace t , we interpret an FO[<, E] formula ϕ over a set of traces T , writing $T \models \phi$ if T satisfies ϕ . We assign first-order variables with elements from the domain $T \times \mathbb{N}$. The $<$ relation is defined as the set $\{(t, n_1), (t, n_2) \in (T \times \mathbb{N})^2 \mid n_1 < n_2\}$ and the equal-level predicate is defined as $\{(t_1, n), (t_2, n) \in (T \times \mathbb{N})^2\}$.

FO[<, E] can express *promptness* requirements such as the existence of a common deadline over all traces by which a certain predicate a must become true on all traces:

$$\exists x. \forall y. E(x, y) \rightarrow P_a(y)$$

Promptness cannot be expressed in HyperLTL [Bozzelli et al. 2015].

HyperQPTL. HyperQPTL [Rabe 2016] captures the ω -regular hyperproperties [Finkbeiner et al. 2020a]. HyperQPTL extends HyperLTL with quantification over atomic propositions. To easily distinguish quantification over traces $\exists\pi, \forall\pi$ and quantification over propositions $\exists p, \forall p$, we use boldface for the latter. The formulas of HyperQPTL are generated by the following grammar:

$$\begin{aligned}\phi &::= \exists\pi. \phi \mid \forall\pi. \phi \mid \psi \mid \exists p. \phi \mid \forall p. \phi \\ \psi &::= a_\pi \mid p \mid \neg\psi \mid \psi \wedge \psi \mid \mathbf{X}\psi \mid \mathbf{F}\psi\end{aligned}$$

where $a, p \in \text{AP}$ and $\pi \in \mathcal{V}$. The semantics of HyperQPTL corresponds to the semantics of HyperLTL with additional rules for propositional quantification:

$$\begin{aligned}T, \Pi, i \models \exists q. \phi &\text{ iff } \exists t \in (2^{\{q\}})^\omega. T, \Pi[\pi_q \mapsto t], i \models \phi, \\ T, \Pi, i \models \forall q. \phi &\text{ iff } \forall t \in (2^{\{q\}})^\omega. T, \Pi[\pi_q \mapsto t], i \models \phi, \\ T, \Pi, i \models q &\text{ iff } q \in \Pi(\pi_q)(i).\end{aligned}$$

HyperQPTL can express all properties expressible in $\text{FO}[\langle, E]$, and, additionally, ω -regular hyperproperties like “on even positions, proposition a has the same value on all traces.” In the following HyperQPTL formula, even positions are identified by the existentially quantified proposition *even*:

$$\exists \text{even}. \forall \pi. \forall \pi'. \text{even} \wedge \mathbf{G}((\text{even} \leftrightarrow \mathbf{X} \neg \text{even}) \wedge (\text{even} \rightarrow (a_\pi \leftrightarrow a_{\pi'})))$$

Another class of properties that can be expressed in HyperQPTL are epistemic properties that express that an agent, who can only observe a subset of the atomic proposition has the *knowledge* that some trace property φ is true. This is expressed as the requirement that all traces that agree with the current trace on the observable propositions up to the current point in time satisfy φ . HyperQPTL subsumes epistemic temporal logic [Rabe 2016].

S1S[E]. S1S[E] is monadic second-order logic with one successor (S1S) extended with the equal-level predicate. Let $V_1 = \{x_1, x_2, \dots\}$ be a set of first-order variables, and $V_2 = \{X_1, X_2, \dots\}$ a set of second-order variables. The formulas ϕ of S1S[E] are generated by the following grammar:

$$\begin{aligned}\tau &::= x \mid \text{min}(x) \mid \text{Succ}(\tau) \\ \phi &::= \tau \in X \mid \tau = \tau \mid E(\tau, \tau) \mid \neg\phi \mid \phi \vee \phi \mid \exists x. \phi \mid \exists X. \phi,\end{aligned}$$

where $x \in V_1$ is a first-order variable, *Succ* denotes the successor relation, and *min*(x) indicates the minimal element of the trace identified by x . Furthermore, $E(\tau, \tau)$ is the equal-level predicate and $X \in V_2 \cup \{X_a \mid a \in \text{AP}\}$. We interpret S1S[E] formulas over a set of traces T . As for $\text{FO}[\langle, E]$, the domain of the first-order variables is $T \times \mathbb{N}$.

S1S[E] adds significant expressiveness to HyperQPTL. For example, it is possible to encode the existence of a terminating computation of a Turing machine [Coenen et al. 2019]. While the model-checking problem of HyperQPTL is still decidable (with a similar algorithm as for HyperLTL), the model-checking problem for S1S[E] is thus undecidable.

7.2. Branching-time hyperlogics

The logics shown on the right in Fig. 1 are branching-time logics. While the linear-time logics are interpreted over sets of traces, we interpret the branching-time logics over infinite trees. A Kripke structure satisfies a branching-time formula iff its unrolling into an infinite tree satisfies the formula.

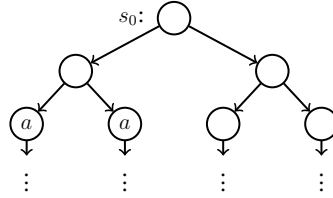
*HyperCTL**. Extending HyperLTL to branching time leads to the temporal logic HyperCTL*. HyperCTL* [Clarkson et al. 2014] has the same syntax as HyperLTL, except that the quantifiers refer to paths rather than traces, and that path quantifiers may occur in the scope of temporal modalities. Let $\pi \in \mathcal{V}$ be a *path variable* from an infinite supply of path variables \mathcal{V} and let $\exists\pi. \phi$ be the explicit existential *path quantification*. HyperCTL* formulas are generated by the following grammar:

$$\phi ::= a_\pi \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U} \phi \mid \exists\pi. \phi.$$

The semantics of a HyperCTL* formula is defined with respect to an infinite tree-shaped Kripke structure \mathcal{T} and a *path assignment* $\Pi : \mathcal{V} \rightarrow Paths(\mathcal{T})$, which is a partial mapping from path variables to paths in the tree. The satisfaction relation \models is given as follows:

$$\begin{aligned} \mathcal{T}, \Pi, i \models a_\pi & \quad \text{iff} \quad a \in L(\Pi(\pi)(i)), \\ \mathcal{T}, \Pi, i \models \neg\phi & \quad \text{iff} \quad \Pi, i \not\models_{\mathcal{T}} \phi, \\ \mathcal{T}, \Pi, i \models \phi_1 \wedge \phi_2 & \quad \text{iff} \quad \mathcal{T}, \Pi, i \models \phi_1 \text{ and } \mathcal{T}, \Pi, i \models \phi_2, \\ \mathcal{T}, \Pi, i \models \mathbf{X}\phi & \quad \text{iff} \quad \mathcal{T}, \Pi, i+1 \models \phi, \\ \mathcal{T}, \Pi, i \models \phi_1 \mathbf{U} \phi_2 & \quad \text{iff} \quad \exists j \geq i. \mathcal{T}, \Pi, j \models \phi_2, \\ & \quad \wedge \forall i \leq k < j. \mathcal{T}, \Pi, k \models_{\mathcal{T}} \phi_1, \\ \mathcal{T}, \Pi, i \models \exists\pi. \phi & \quad \text{iff} \quad \exists p \in Paths(\mathcal{T}). p[0, i] = \varepsilon[0, i], \\ & \quad \wedge \mathcal{T}, \Pi[\pi \mapsto p, \varepsilon \mapsto p], i \models \phi, \end{aligned}$$

where we use ε to denote the last path that was added to the path assignment Π . We say that a tree \mathcal{T} satisfies a HyperCTL* formula ϕ , written as $\mathcal{T} \models \phi$, if $\mathcal{T}, \emptyset, 0 \models \phi$, where \emptyset denotes the empty path assignment. HyperCTL* can express the flow of information that appears in different branches of the computation tree, as illustrated in the following example (taken from [Finkbeiner 2017]):



An observer who sees a can infer which branch was taken in the first nondeterministic choice, but not which branch was taken in the second nondeterministic choice. This is expressed by the HyperCTL* formula

$$\forall\pi. \mathbf{X} \forall\pi'. \mathbf{X} (a_\pi \leftrightarrow a_{\pi'}).$$

MPL[E]. Monadic path logic equipped with the equal-level predicate (MPL[E]) is the extension of $FO[<, E]$ with second-order quantification, where the second-order quantification is restricted to full paths in the tree.

Let $V_1 = \{x_1, x_2, \dots\}$ be a set of first-order variables, and $V_2 = \{X_1, X_2, \dots\}$ a set of second-order variables. The formulas ϕ of MPL[E] are generated by the following grammar:

$$\begin{aligned} \phi & ::= \psi \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \exists x. \phi \mid \exists X. \phi \\ \psi & ::= P_a(x) \mid x < y \mid x = y \mid x \in X \mid E(x, y), \end{aligned}$$

where $a \in AP$, $x, y \in V_1$, and $X \in V_2$. In the semantics of MPL[E], first-order variables are assigned to nodes in the tree and second-order variables to sets of nodes whose elements constitute exactly one path of the tree. P_a evaluates to the set of nodes

whose label contains a . $x < y$ indicates that x is an ancestor of y . Atomic formulas $x \in X$ and $x = y$ are interpreted as set membership and equality on nodes, respectively. The equal-level predicate $E(x, y)$ denotes that two nodes x and y are on the same level, i.e., have the same number of ancestors. $\text{MPL}[E]$ is strictly more expressive than HyperCTL^* . Similarly to HyperQPTL , $\text{MPL}[E]$ can express properties about knowledge [Coenen et al. 2019].

HyperQCTL and MSO[E].* HyperQCTL^* [Coenen et al. 2019] extends HyperCTL^* with quantification over atomic propositions. The formulas of HyperQCTL^* are generated by the following grammar:

$$\phi ::= a_\pi \mid \neg\phi \mid \phi \vee \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U} \phi \mid \exists\pi. \phi \mid \exists p. \phi$$

where $a, p \in \text{AP}$ and $\pi \in \mathcal{V}$. The semantics of HyperQCTL^* corresponds to the semantics of HyperCTL^* with an additional rule for propositional quantification: $\exists p. \phi$ is satisfied iff ϕ is satisfied after the labeling of the tree has been modified by (re-)assigning the valuation of p to the nodes of the tree. HyperQCTL^* has the same expressiveness as second-order monadic logic equipped with the equal-level predicate ($\text{MSO}[E]$), i.e., the extension of $\text{FO}[<, E]$ with second-order quantification [Coenen et al. 2019]. While the model-checking problem for $\text{MPL}[E]$ is still decidable [Finkbeiner 2021], the model-checking problems for HyperQCTL^* and $\text{MSO}[E]$ are undecidable [Coenen et al. 2019].

7.3. Further extensions

In addition to the logics discussed so far, HyperLTL has been extended in numerous further directions. In the following, we briefly discuss three major areas of ongoing research.

Probabilistic hyperproperties. To reason about probabilistic hyperproperties of Markov decision processes, HyperLTL has been extended with probabilistic operators and strategy quantifiers, resulting in the temporal logic PHL [Dimitrova et al. 2020]; similar extensions are the logics HyperPCTL [Ábrahám and Bonakdarpour 2018; Ábrahám et al. 2020] and HyperPCTL^* [Wang et al. 2021].

Infinite-state hyperproperties. HyperTSL [Coenen et al. 2023] is a temporal logic for the specification of hyperproperties of infinite-state software. It is based on temporal stream logic (TSL) [Finkbeiner et al. 2019], which extends linear temporal logic (LTL) with the concept of cells and uninterpreted functions and predicates. These mechanisms separate the temporal control flow from the concrete data in the system, which enables reasoning about infinite-state systems. *First-order HyperLTL* extends HyperLTL with first-order quantifiers over uninterpreted sorts [Finkbeiner et al. 2017]. For cyber-physical systems, HyperLTL has also been extended with *real-valued signals* [Nguyen et al. 2017; Wang et al. 2019].

Asynchronous hyperproperties. While HyperLTL has a strictly synchronous semantics, and a general logic for asynchronous hyperproperties has not been proposed yet, there are already two extensions of HyperLTL that go in this direction: the extension with *trajectory quantifiers* [Baumeister et al. 2021] considers multiple schedulings between the traces; *Stuttering HyperLTL* [Bozzelli et al. 2021] eliminates stuttering steps from the traces. A further big step towards asynchronous hyperproperties is the *temporal fixpoint calculus* H_μ [Gutsfeld et al. 2021], which integrates the quantifiers from HyperLTL into the modal μ -calculus. Many asynchronous hyperproperties can also be encoded in HyperATL^* [Beutner and Finkbeiner 2021], an extension of HyperLTL to

games with multiple coalitions of players. Asynchronous hyperproperties are increasingly supported by practical verification approaches such as bounded model checking [Hsu et al. 2023].

8. CONCLUDING REMARKS

HyperLTL and its extensions are a powerful unifying framework for reasoning about hyperproperties. For a wide range of algorithmic problems, including satisfiability, model checking, monitoring, and synthesis, the logics provide a common specification language and semantic foundation. It is encouraging that the model-checking problem remains decidable for a large portion of the hierarchy of hyperlogics, allowing for the development of automated verification techniques that cover broad classes of hyperproperties in a uniform manner.

Despite the significant progress over the past decade, much work remains to be done. Many of the currently available algorithms focus on finite-state hardware; reasoning about hyperproperties of infinite-state software remains a challenge. It is also clear that the expressiveness of logics like HyperLTL is dwarfed by the general power of hyperproperties. Many concepts that relate to concerns like privacy [Kalloniatis et al. 2008], explainability [Clinciu and Hastie 2019], or fairness [Mehrabani et al. 2021] are arguably covered by Clarkson and Schneider’s general definition of hyperproperties as properties of sets of traces; yet a specification language that would provide a unifying theory for all these notions remains out of reach for now.

ACKNOWLEDGMENTS

This work was partially supported by the European Research Council (ERC) grant HYPER (No. 101055412).

REFERENCES

- Erika Ábrahám, Ezio Bartocci, Borzoo Bonakdarpour, and Oyendrila Dobe. 2020. Probabilistic Hyperproperties with Nondeterminism. In *Automated Technology for Verification and Analysis*, Dang Van Hung and Oleg Sokolsky (Eds.). Springer International Publishing, Cham, 518–534.
- Erika Ábrahám and Borzoo Bonakdarpour. 2018. HyperPCTL: A Temporal Logic for Probabilistic Hyperproperties. In *Quantitative Evaluation of Systems - 15th International Conference, QEST 2018, Beijing, China, September 4-7, 2018, Proceedings (Lecture Notes in Computer Science)*, Annabelle McIver and András Horváth (Eds.), Vol. 11024. Springer, 20–35. DOI: http://dx.doi.org/10.1007/978-3-319-99154-2_2
- Karl Raymond Abrahamson. 1980. *Decidability and Expressiveness of Logics of Processes*. Ph.D. Dissertation. Seattle, WA, USA. AAI8109709.
- Shreya Agrawal and Borzoo Bonakdarpour. 2016. Runtime Verification of k -Safety Hyperproperties in HyperLTL. In *CSF 2016*. IEEE Computer Society, 239–252. DOI: <http://dx.doi.org/10.1109/CSF.2016.24>
- Gilles Barthe, Pedro R. D’Argenio, Bernd Finkbeiner, and Holger Hermanns. 2016. Facets of Software Doping. In *Leveraging Applications of Formal Methods, Verification and Validation: Discussion, Dissemination, Applications*, Tiziana Margaria and Bernhard Steffen (Eds.). Springer International Publishing, Cham, 601–608.
- Gilles Barthe, Pedro R. D’Argenio, and Tamara Rezk. 2011. Secure information flow by self-composition. *Math. Struct. Comput. Sci.* 21, 6 (2011), 1207–1252. DOI: <http://dx.doi.org/10.1017/S0960129511000193>
- Gilles Barthe, Thomas Espitau, Benjamin Grégoire, Justin Hsu, and Pierre-Yves Strub. 2018. Proving expected sensitivity of probabilistic programs. *Proc. ACM Program. Lang.* 2, POPL (2018), 57:1–57:29. DOI: <http://dx.doi.org/10.1145/3158145>
- Gilles Barthe, Boris Köpf, Federico Olmedo, and Santiago Zanella Béguelin. 2013. Probabilistic Relational Reasoning for Differential Privacy. *ACM Trans. Program. Lang. Syst.* 35, 3 (2013), 9:1–9:49. DOI: <http://dx.doi.org/10.1145/2492061>
- Jan Baumeister, Norine Coenen, Borzoo Bonakdarpour, Bernd Finkbeiner, and César Sánchez. 2021. A Temporal Logic for Asynchronous Hyperproperties. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I (Lecture Notes in Computer Science)*, Alexandra Silva and K. Rustan M. Leino (Eds.), Vol. 12759. Springer, 694–717. DOI: http://dx.doi.org/10.1007/978-3-030-81685-8_33

- Nick Benton. 2004. Simple relational correctness proofs for static analyses and program transformations. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*, Neil D. Jones and Xavier Leroy (Eds.). ACM, 14–25. DOI: <http://dx.doi.org/10.1145/964001.964003>
- Raven Beutner, David Carral, Bernd Finkbeiner, Jana Hofmann, and Markus Krötzsch. 2022. Deciding Hyperproperties Combined with Functional Specifications. In *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, Haifa, Israel, August 2 - 5, 2022*, Christel Baier and Dana Fisman (Eds.). ACM, 56:1–56:13. DOI: <http://dx.doi.org/10.1145/3531130.3533369>
- Raven Beutner and Bernd Finkbeiner. 2021. A Temporal Logic for Strategic Hyperproperties. In *32nd International Conference on Concurrency Theory, CONCUR 2021, August 24-27, 2021, Virtual Conference (LIPIcs)*, Serge Haddad and Daniele Varacca (Eds.), Vol. 203. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 24:1–24:19. DOI: <http://dx.doi.org/10.4230/LIPIcs.CONCUR.2021.24>
- Raven Beutner and Bernd Finkbeiner. 2022a. Prophecy Variables for Hyperproperty Verification. In *35th IEEE Computer Security Foundations Symposium, CSF 2022, Haifa, Israel, August 7-10, 2022*. IEEE, 471–485. DOI: <http://dx.doi.org/10.1109/CSF54842.2022.9919658>
- Raven Beutner and Bernd Finkbeiner. 2022b. Software Verification of Hyperproperties Beyond k-Safety. In *Computer Aided Verification - 34th International Conference, CAV 2022, Haifa, Israel, August 7-10, 2022, Proceedings, Part I (Lecture Notes in Computer Science)*, Sharon Shoham and Yakir Vizel (Eds.), Vol. 13371. Springer, 341–362. DOI: http://dx.doi.org/10.1007/978-3-031-13185-1_17
- Raven Beutner and Bernd Finkbeiner. 2023. AutoHyper: Explicit-State Model Checking for HyperLTL. In *Tools and Algorithms for the Construction and Analysis of Systems - 29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Paris, France, April 22–27, 2023, Proceedings, Part II (Lecture Notes in Computer Science)*, Sriram Sankaranarayanan and Natasha Sharygina (Eds.), Vol. 13994. Springer.
- Borzoo Bonakdarpour and Bernd Finkbeiner. 2018. The Complexity of Monitoring Hyperproperties. In *31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018*. IEEE Computer Society, 162–174. DOI: <http://dx.doi.org/10.1109/CSF.2018.00019>
- Borzoo Bonakdarpour and Bernd Finkbeiner. 2020. Controller Synthesis for Hyperproperties. In *33rd IEEE Computer Security Foundations Symposium, CSF 2020, Boston, MA, USA, June 22-26, 2020*. IEEE, 366–379. DOI: <http://dx.doi.org/10.1109/CSF49147.2020.00033>
- Laura Bozzelli, Bastien Maubert, and Sophie Pinchinat. 2015. Unifying Hyper and Epistemic Temporal Logics. In *FoSSaCS 2015 (Lecture Notes in Computer Science)*, Andrew M. Pitts (Ed.), Vol. 9034. Springer, 167–182. DOI: http://dx.doi.org/10.1007/978-3-662-46678-0_11
- Laura Bozzelli, Adriano Peron, and César Sánchez. 2021. Asynchronous Extensions of HyperLTL. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*. IEEE, 1–13. DOI: <http://dx.doi.org/10.1109/LICS52264.2021.9470583>
- Noel Brett, Umair Siddique, and Borzoo Bonakdarpour. 2017. Rewriting-Based Runtime Verification for Alternation-Free HyperLTL. In *Tools and Algorithms for the Construction and Analysis of Systems, Axel Legay and Tiziana Margaria (Eds.)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 77–93.
- Swarat Chaudhuri, Sumit Gulwani, and Roberto Lubliner. 2012. Continuity and robustness of programs. *Commun. ACM* 55, 8 (2012), 107–115. DOI: <http://dx.doi.org/10.1145/2240236.2240262>
- M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez. 2014. Temporal Logics for Hyperproperties. In *Proceedings of the 3rd Conference on Principles of Security and Trust (POST)*. 265–284.
- Michael R. Clarkson and Fred B. Schneider. 2008. Hyperproperties. In *2008 21st IEEE Computer Security Foundations Symposium*. 51–65. DOI: <http://dx.doi.org/10.1109/CSF.2008.7>
- Miruna A. Clinciu and Helen F. Hastie. 2019. A Survey of Explainable AI Terminology. In *Proceedings of the 1st Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence*. 8–13. 1st Workshop on Interactive Natural Language Technology for Explainable Artificial Intelligence, NL4XAI 2019.
- Norine Coenen, Bernd Finkbeiner, Christopher Hahn, and Jana Hofmann. 2019. The Hierarchy of Hyperlogics. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*. IEEE, 1–13. DOI: <http://dx.doi.org/10.1109/LICS.2019.8785713>
- Norine Coenen, Bernd Finkbeiner, Jana Hofmann, and Julia Tillman. 2023. Smart Contract Synthesis Modulo Hyperproperties. In *2023 36th Computer Security Foundations Symposium (CSF) (CSF)*. IEEE Computer Society, Los Alamitos, CA, USA, 79–94. DOI: <http://dx.doi.org/10.1109/CSF57540.2023.00006>
- Norine Coenen, Bernd Finkbeiner, César Sánchez, and Leander Tentrup. 2019. Verifying Hyperliveness. In *Computer Aided Verification - 31st International Conference, CAV 2019, Part I (Lecture*

- Notes in Computer Science*), Isil Dillig and Serdar Tasiran (Eds.), Vol. 11561. Springer, 121–139. DOI: http://dx.doi.org/10.1007/978-3-030-25540-4_7
- Dominique Devriese and Frank Piessens. 2010. Noninterference through Secure Multi-execution. In *Proceedings of SP*. IEEE Computer Society, 109–124. DOI: <http://dx.doi.org/10.1109/SP.2010.15>
- Rayna Dimitrova and Bernd Finkbeiner. 2009. Synthesis of Fault-Tolerant Distributed Systems. In *Proceedings of ATVA (LNCS)*, Vol. 5799. Springer, 321–336. DOI: http://dx.doi.org/10.1007/978-3-642-04761-9_24
- Rayna Dimitrova, Bernd Finkbeiner, and Hazem Torfah. 2020. Probabilistic Hyperproperties of Markov Decision Processes. In *Automated Technology for Verification and Analysis*, Dang Van Hung and Oleg Sokolsky (Eds.). Springer International Publishing, Cham, 484–500.
- Geir Dullerud and Fernando Paganini. 2000. *A Course in Robust Control Theory a convex approach*. DOI: <http://dx.doi.org/10.1007/978-1-4757-3290-0>
- R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. 2003. *Reasoning about Knowledge*. MIT Press, Cambridge, Mass.
- Bernd Finkbeiner. 2017. Temporal Hyperproperties. *Bulletin of the EATCS* 123 (2017).
- Bernd Finkbeiner. 2021. Model Checking Algorithms for Hyperproperties (Invited Paper). In *Verification, Model Checking, and Abstract Interpretation - 22nd International Conference, VMCAI 2021, Copenhagen, Denmark, January 17-19, 2021, Proceedings (Lecture Notes in Computer Science)*, Fritz Henglein, Sharon Shoham, and Yakir Vizel (Eds.), Vol. 12597. Springer, 3–16. DOI: http://dx.doi.org/10.1007/978-3-030-67067-2_1
- Bernd Finkbeiner and Christopher Hahn. 2016. Deciding Hyperproperties. In *CONCUR 2016 (LIPIcs)*, José Desharnais and Radha Jagadeesan (Eds.), Vol. 59. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 13:1–13:14. DOI: <http://dx.doi.org/10.4230/LIPIcs.CONCUR.2016.13>
- Bernd Finkbeiner, Christopher Hahn, and Tobias Hans. 2018. MGHyper: Checking Satisfiability of HyperLTL Formulas Beyond the $\exists^*\forall^*$ Fragment. In *Automated Technology for Verification and Analysis - 16th International Symposium, ATVA 2018 (Lecture Notes in Computer Science)*, Shuvendu K. Lahiri and Chao Wang (Eds.), Vol. 11138. Springer, 521–527. DOI: http://dx.doi.org/10.1007/978-3-030-01090-4_31
- Bernd Finkbeiner, Christopher Hahn, Jana Hofmann, and Leander Tentrup. 2020a. Realizing ω -regular Hyperproperties. In *Computer Aided Verification*. Springer International Publishing, Cham, 40–63.
- Bernd Finkbeiner, Christopher Hahn, Philip Lukert, Marvin Stenger, and Leander Tentrup. 2020b. Synthesis from hyperproperties. *Acta Informatica* 57, 1-2 (2020), 137–163. DOI: <http://dx.doi.org/10.1007/s00236-019-00358-2>
- Bernd Finkbeiner, Christopher Hahn, and Marvin Stenger. 2017. EAHyper: Satisfiability, Implication, and Equivalence Checking of Hyperproperties. In *Proceedings of CAV (LNCS)*, Vol. 10427. Springer, 564–570. DOI: http://dx.doi.org/10.1007/978-3-319-63390-9_29
- Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. 2018. RVHyper: A Runtime Verification Tool for Temporal Hyperproperties. In *Proceedings of TACAS (LNCS)*, Vol. 10806. Springer, 194–200. DOI: http://dx.doi.org/10.1007/978-3-319-89963-3_11
- Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. 2019. Monitoring hyperproperties. *Formal Methods Syst. Des.* 54, 3 (2019), 336–363. DOI: <http://dx.doi.org/10.1007/s10703-019-00334-z>
- Bernd Finkbeiner, Christopher Hahn, Marvin Stenger, and Leander Tentrup. 2020. Efficient monitoring of hyperproperties using prefix trees. *Int. J. Softw. Tools Technol. Transf.* 22, 6 (2020), 729–740. DOI: <http://dx.doi.org/10.1007/s10009-020-00552-5>
- Bernd Finkbeiner, Felix Klein, Ruzica Piskac, and Mark Santolucito. 2019. Temporal Stream Logic: Synthesis Beyond the Booleans. In *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part I (Lecture Notes in Computer Science)*, Isil Dillig and Serdar Tasiran (Eds.), Vol. 11561. Springer, 609–629. DOI: http://dx.doi.org/10.1007/978-3-030-25540-4_35
- Bernd Finkbeiner, Christian Müller, Helmut Seidl, and Eugen Zalinescu. 2017. Verifying Security Policies in Multi-agent Workflows with Loops. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 633–645. DOI: <http://dx.doi.org/10.1145/3133956.3134080>
- Bernd Finkbeiner, Markus N. Rabe, and César Sánchez. 2015. Algorithms for Model Checking HyperLTL and HyperCTL*. In *CAV 2015 (Part I) (Lecture Notes in Computer Science)*, Daniel Kroening and Corina S. Pasareanu (Eds.), Vol. 9206. Springer, 30–48. DOI: http://dx.doi.org/10.1007/978-3-319-21690-4_3
- Bernd Finkbeiner and Sven Schewe. 2013. Bounded synthesis. *Int. J. Softw. Tools Technol. Transf.* 15, 5-6 (2013), 519–539. DOI: <http://dx.doi.org/10.1007/s10009-012-0228-z>

- Bernd Finkbeiner and Martin Zimmermann. 2017. The First-Order Logic of Hyperproperties. In *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*. 30:1–30:14. DOI: <http://dx.doi.org/10.4230/LIPIcs.STACS.2017.30>
- Marie Fortin, Louwe B. Kuijjer, Patrick Totzke, and Martin Zimmermann. 2021. HyperLTL Satisfiability is Σ_1^1 -complete, HyperCTL* Satisfiability is Σ_2^1 -complete. In *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021, Tallinn, Estonia (LIPIcs)*, Filippo Bonchi and Simon J. Puglisi (Eds.), Vol. 202. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 47:1–47:19.
- Tim French. 2001. Decidability of Quantified Propositional Branching Time Logics. In *AI 2001: Advances in Artificial Intelligence, 14th Australian Joint Conference on Artificial Intelligence, Adelaide, Australia, December 10-14, 2001, Proceedings*. 165–176. DOI: http://dx.doi.org/10.1007/3-540-45656-2_15
- Dov M. Gabbay, Amir Pnueli, Saharon Shelah, and Jonathan Stavi. 1980. On the Temporal Basis of Fairness. In *POPL 1980*, Paul W. Abrahams, Richard J. Lipton, and Stephen R. Bourne (Eds.). ACM Press, 163–173. DOI: <http://dx.doi.org/10.1145/567446.567462>
- Joseph A. Goguen and José Meseguer. 1982. Security Policies and Security Models. In *IEEE Symposium on Security and Privacy*. 11–20.
- Jens Oliver Gutsfeld, Markus Müller-Olm, and Christoph Ohrem. 2021. Automata and Fixpoints for Asynchronous Hyperproperties. *Proc. ACM Program. Lang.* 5, POPL, Article 38 (Jan. 2021), 29 pages. DOI: <http://dx.doi.org/10.1145/3434319>
- Christopher Hahn, Marvin Stenger, and Leander Tentrup. 2019. Constraint-Based Monitoring of Hyperproperties. In *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, April 6-11, 2019, Proceedings, Part II (Lecture Notes in Computer Science)*, Tomáš Vojnar and Lijun Zhang (Eds.), Vol. 11428. Springer, 115–131. DOI: http://dx.doi.org/10.1007/978-3-030-17465-1_7
- J. Y. Halpern and Y. Moses. 1990. Knowledge and Common Knowledge in a Distributed Environment. *J. ACM* 37, 3 (1990), 549–587.
- Tzu-Han Hsu, Borzoo Bonakdarpour, Bernd Finkbeiner, and César Sánchez. 2023. Bounded Model Checking for Asynchronous Hyperproperties. In *Tools and Algorithms for the Construction and Analysis of Systems - 29th International Conference, TACAS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22–27, 2023, Proceedings, Part II (Lecture Notes in Computer Science)*, Sriram Sankaranarayanan and Natasha Sharygina (Eds.), Vol. 13994. Springer.
- Tzu-Han Hsu, César Sánchez, and Borzoo Bonakdarpour. 2021. Bounded Model Checking for Hyperproperties. In *Tools and Algorithms for the Construction and Analysis of Systems, Jan Friso Groote and Kim Guldstrand Larsen (Eds.)*. Springer International Publishing, Cham, 94–112.
- Christos Kalloniatis, Evangelia Kavakli, and Stefanos Gritzalis. 2008. Addressing privacy requirements in system design: the PriS method. *Requir. Eng.* 13, 3 (2008), 241–255. DOI: <http://dx.doi.org/10.1007/s00766-008-0067-3>
- Hans W. Kamp. 1968. *Tense Logic and the Theory of Linear Order*. Ph.D. Dissertation. Computer Science Department, University of California at Los Angeles, USA.
- Yonit Kesten and Amir Pnueli. 1995. A Complete Proof Systems for QPTL. In *Proceedings, 10th Annual IEEE Symposium on Logic in Computer Science, San Diego, California, USA, June 26-29, 1995*. 2–12. DOI: <http://dx.doi.org/10.1109/LICS.1995.523239>
- Orna Kupferman and Moshe Y. Vardi. 1997. Synthesis with incomplete information. In *Temporal Logic, 2nd International Conference, ICTL '97, Manchester, England, July 14-18, 1997, Proceedings*.
- François Laroussinie and Nicolas Markey. 2014. Quantified CTL: Expressiveness and Complexity. *Logical Methods in Computer Science* 10, 4 (2014). DOI: [http://dx.doi.org/10.2168/LMCS-10\(4:17\)2014](http://dx.doi.org/10.2168/LMCS-10(4:17)2014)
- Heiko Mantel and Alexander Reinhard. 2007. Controlling the What and Where of Declassification in Language-Based Security. In *Programming Languages and Systems*, Rocco De Nicola (Ed.). Lecture Notes in Computer Science, Vol. 4421. Springer Berlin Heidelberg, 141–156. DOI: http://dx.doi.org/10.1007/978-3-540-71316-6_11
- Corto Mascle and Martin Zimmermann. 2020. The Keys to Decidable HyperLTL Satisfiability: Small Models or Very Simple Formulas. In *28th EACSL Annual Conference on Computer Science Logic, CSL 2020, January 13-16, 2020, Barcelona, Spain (LIPIcs)*, Maribel Fernández and Anca Muscholl (Eds.), Vol. 152. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 29:1–29:16. DOI: <http://dx.doi.org/10.4230/LIPIcs.CSL.2020.29>
- Daryl McCullough. 1988. Noninterference and the composability of security properties. In *Proc. IEEE Symposium on Security and Privacy*. 177–186.

- John McLean. 1994. A General Theory of Composition for Trace Sets Closed Under Selective Interleaving Functions. In *Proc. IEEE Symposium on Security and Privacy*. 79–93.
- Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A Survey on Bias and Fairness in Machine Learning. *ACM Comput. Surv.* 54, 6, Article 115 (July 2021), 35 pages. DOI : <http://dx.doi.org/10.1145/3457607>
- Luan Viet Nguyen, James Kapinski, Xiaoqing Jin, Jyotirmoy V. Deshmukh, and Taylor T. Johnson. 2017. Hyperproperties of Real-Valued Signals. In *Proceedings of the 15th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE '17)*. Association for Computing Machinery, New York, NY, USA, 104–113. DOI : <http://dx.doi.org/10.1145/3127041.3127058>
- Amir Pnueli. 1977. The Temporal Logic of Programs. In *FOCS 1977*. 46–57.
- Amir Pnueli and Roni Rosner. 1990. 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*. 746–757. DOI : <http://dx.doi.org/doi:10.1109/FSCS.1990.89597>
- Markus N. Rabe. 2016. *A Temporal Logic Approach to Information-flow Control*. Ph.D. Dissertation. Saarland University.
- John Alan Robinson and Andrei Voronkov (Eds.). 2001. *Handbook of Automated Reasoning*. Elsevier and MIT Press. <https://www.sciencedirect.com/book/9780444508133/handbook-of-automated-reasoning>
- Andrei Sabelfeld and David Sands. 2009. Declassification: Dimensions and principles. *Journal of Computer Security* 17, 5 (2009), 517–548. DOI : <http://dx.doi.org/10.3233/JCS-2009-0352>
- Aravinda Prasad Sistla. 1983. *Theoretical Issues in the Design and Verification of Distributed Systems*. Ph.D. Dissertation.
- A. P. Sistla and E. M. Clarke. 1985. The Complexity of Propositional Linear Temporal Logics. *J. ACM* 32, 3 (jul 1985), 733–749. DOI : <http://dx.doi.org/10.1145/3828.3837>
- Wolfgang Thomas. 2009. Path Logics with Synchronization. In *Perspectives in Concurrency Theory*, Kamal Lodaya, Madhavan Mukund, and R. Ramanujam (Eds.). IARCS-Universities, Universities Press, 469–481.
- Yu Wang, Siddhartha Nalluri, Borzoo Bonakdarpour, and Miroslav Pajic. 2021. Statistical Model Checking for Hyperproperties. In *34th IEEE Computer Security Foundations Symposium, CSF 2021, Dubrovnik, Croatia, June 21-25, 2021*. IEEE, 1–16. DOI : <http://dx.doi.org/10.1109/CSF51468.2021.00009>
- Yu Wang, Mojtaba Zarei, Borzoo Bonakdarpour, and Miroslav Pajic. 2019. Statistical Verification of Hyperproperties for Cyber-Physical Systems. *ACM Trans. Embed. Comput. Syst.* 18, 5s (2019), 92:1–92:23. DOI : <http://dx.doi.org/10.1145/3358232>
- S. Zdancewic and A. C. Myers. 2003. Observational Determinism for Concurrent Program Security. In *Proc. of CSFW03*.